



MapReduce Service

Component Operation Guide

Date **2024-07-19**

Contents

1 Using Alluxio.....	1
1.1 Configuring an Underlying Storage System.....	1
1.2 Accessing Alluxio Using a Data Application.....	2
1.3 Common Operations of Alluxio.....	5
2 Using CarbonData (for Versions Earlier Than MRS 3.x).....	8
2.1 Using CarbonData from Scratch.....	8
2.2 About CarbonData Table.....	10
2.3 Creating a CarbonData Table.....	11
2.4 Deleting a CarbonData Table.....	13
3 Using CarbonData (for MRS 3.x or Later).....	15
3.1 Overview.....	15
3.1.1 CarbonData Overview.....	15
3.1.2 Main Specifications of CarbonData.....	18
3.2 Configuration Reference.....	20
3.3 CarbonData Operation Guide.....	33
3.3.1 CarbonData Quick Start.....	33
3.3.2 CarbonData Table Management.....	36
3.3.2.1 About CarbonData Table.....	36
3.3.2.2 Creating a CarbonData Table.....	38
3.3.2.3 Deleting a CarbonData Table.....	40
3.3.2.4 Modify the CarbonData Table.....	40
3.3.3 CarbonData Table Data Management.....	41
3.3.3.1 Loading Data.....	41
3.3.3.2 Deleting Segments.....	41
3.3.3.3 Combining Segments.....	43
3.3.4 CarbonData Data Migration.....	46
3.3.5 Migrating Data on CarbonData from Spark 1.5 to Spark2x.....	48
3.4 CarbonData Performance Tuning.....	49
3.4.1 Tuning Guidelines.....	50
3.4.2 Suggestions for Creating CarbonData Tables.....	52
3.4.3 Configurations for Performance Tuning.....	54
3.5 CarbonData Access Control.....	57

3.6 CarbonData Syntax Reference.....	59
3.6.1 DDL.....	59
3.6.1.1 CREATE TABLE.....	59
3.6.1.2 CREATE TABLE As SELECT.....	62
3.6.1.3 DROP TABLE.....	63
3.6.1.4 SHOW TABLES.....	64
3.6.1.5 ALTER TABLE COMPACTION.....	65
3.6.1.6 TABLE RENAME.....	66
3.6.1.7 ADD COLUMNS.....	67
3.6.1.8 DROP COLUMNS.....	68
3.6.1.9 CHANGE DATA TYPE.....	69
3.6.1.10 REFRESH TABLE.....	70
3.6.1.11 REGISTER INDEX TABLE.....	71
3.6.2 DML.....	72
3.6.2.1 LOAD DATA.....	73
3.6.2.2 UPDATE CARBON TABLE.....	78
3.6.2.3 DELETE RECORDS from CARBON TABLE.....	79
3.6.2.4 INSERT INTO CARBON TABLE.....	80
3.6.2.5 DELETE SEGMENT by ID.....	82
3.6.2.6 DELETE SEGMENT by DATE.....	82
3.6.2.7 SHOW SEGMENTS.....	83
3.6.2.8 CREATE SECONDARY INDEX.....	84
3.6.2.9 SHOW SECONDARY INDEXES.....	85
3.6.2.10 DROP SECONDARY INDEX.....	86
3.6.2.11 CLEAN FILES.....	87
3.6.2.12 SET/RESET.....	88
3.6.3 Operation Concurrent Execution.....	91
3.6.4 API.....	95
3.6.5 Spatial Indexes.....	96
3.7 CarbonData Troubleshooting.....	111
3.7.1 Filter Result Is not Consistent with Hive when a Big Double Type Value Is Used in Filter.....	111
3.7.2 Query Performance Deterioration.....	111
3.8 CarbonData FAQ.....	112
3.8.1 Why Is Incorrect Output Displayed When I Perform Query with Filter on Decimal Data Type Values?.....	112
3.8.2 How to Avoid Minor Compaction for Historical Data?.....	113
3.8.3 How to Change the Default Group Name for CarbonData Data Loading?.....	113
3.8.4 Why Does INSERT INTO CARBON TABLE Command Fail?.....	114
3.8.5 Why Is the Data Logged in Bad Records Different from the Original Input Data with Escape Characters?.....	114
3.8.6 Why Data Load Performance Decreases due to Bad Records?.....	115
3.8.7 Why INSERT INTO/LOAD DATA Task Distribution Is Incorrect and the Opened Tasks Are Less Than the Available Executors when the Number of Initial Executors Is Zero?.....	115

3.8.8 Why Does CarbonData Require Additional Executors Even Though the Parallelism Is Greater Than the Number of Blocks to Be Processed?.....	116
3.8.9 Why Data loading Fails During off heap?.....	116
3.8.10 Why Do I Fail to Create a Hive Table?.....	117
3.8.11 Why CarbonData tables created in V100R002C50RC1 not reflecting the privileges provided in Hive Privileges for non-owner?.....	117
3.8.12 How Do I Logically Split Data Across Different Namespaces?.....	118
3.8.13 Why Missing Privileges Exception is Reported When I Perform Drop Operation on Databases?....	119
3.8.14 Why the UPDATE Command Cannot Be Executed in Spark Shell?.....	119
3.8.15 How Do I Configure Unsafe Memory in CarbonData?.....	120
3.8.16 Why Exception Occurs in CarbonData When Disk Space Quota is Set for Storage Directory in HDFS?.....	120
3.8.17 Why Does Data Query or Loading Fail and "org.apache.carbondata.core.memory.MemoryException: Not enough memory" Is Displayed?.....	121
3.8.18 Why Do Files of a Carbon Table Exist in the Recycle Bin Even If the drop table Command Is Not Executed When Mis-deletion Prevention Is Enabled?.....	122
4 Using ClickHouse.....	123
4.1 Using ClickHouse from Scratch.....	123
4.2 ClickHouse Table Engine Overview.....	126
4.3 Creating a ClickHouse Table.....	133
4.4 ClickHouse Data Type.....	138
4.5 Common ClickHouse SQL Syntax.....	142
4.5.1 CREATE DATABASE: Creating a Database.....	142
4.5.2 CREATE TABLE: Creating a Table.....	143
4.5.3 INSERT INTO: Inserting Data into a Table.....	144
4.5.4 SELECT: Querying Table Data.....	144
4.5.5 ALTER TABLE: Modifying a Table Structure.....	145
4.5.6 DESC: Querying a Table Structure.....	146
4.5.7 DROP: Deleting a Table.....	146
4.5.8 SHOW: Displaying Information About Databases and Tables.....	147
4.6 Migrating ClickHouse Data.....	147
4.6.1 Using ClickHouse to Import and Export Data.....	147
4.6.2 Synchronizing Kafka Data to ClickHouse.....	149
4.6.3 Using the ClickHouse Data Migration Tool.....	152
4.7 User Management and Authentication.....	155
4.7.1 ClickHouse User and Permission Management.....	155
4.7.2 Interconnecting ClickHouse With OpenLDAP for Authentication.....	161
4.8 Backing Up and Restoring ClickHouse Data Using a Data File.....	164
4.9 ClickHouse Log Overview.....	166
4.10 ClickHouse Performance Tuning.....	169
4.10.1 Solution to the "Too many parts" Error in Data Tables.....	169
4.10.2 Accelerating Merge Operations.....	170
4.10.3 Accelerating TTL Operations.....	171

4.11 ClickHouse FAQ.....	172
4.11.1 How Do I Do If the Disk Status Displayed in the System.disks Table Is fault or abnormal?.....	172
4.11.2 How Do I Migrate Data from Hive/HDFS to ClickHouse?.....	172
4.11.3 How Do I Migrate Data from OBS/S3 to ClickHouse?.....	173
4.11.4 An Error Is Reported in Logs When the Auxiliary ZooKeeper or Replica Data Is Used to Synchronize Table Data.....	174
4.11.5 How Do I Grant the Select Permission at the Database Level to ClickHouse Users?.....	174
5 Using DBService.....	176
5.1 DBService Log Overview.....	176
6 Using Flink.....	180
6.1 Using Flink from Scratch.....	180
6.2 Viewing Flink Job Information.....	190
6.3 Configuring Flink Service Parameters.....	191
6.4 Configuring Flink Security Features.....	217
6.4.1 Security Features.....	217
6.4.2 Authentication and Encryption.....	221
6.4.3 Configuring Kafka.....	227
6.4.4 Configuring Pipeline.....	229
6.5 Configuring and Developing a Flink Visualization Job.....	230
6.5.1 Introduction to Flink Web UI.....	230
6.5.2 Flink Web UI Permission Management.....	234
6.5.3 Creating a FlinkServer Role.....	234
6.5.4 Accessing the Flink Web UI.....	235
6.5.5 Creating an Application.....	236
6.5.6 Creating a Cluster Connection.....	236
6.5.7 Creating a Data Connection.....	238
6.5.8 Creating a Stream Table.....	239
6.5.9 Creating a Job.....	241
6.6 Flink Log Overview.....	244
6.7 Flink Performance Tuning.....	247
6.7.1 Memory Configuration Optimization.....	247
6.7.2 Configuring DOP.....	248
6.7.3 Configuring Process Parameters.....	249
6.7.4 Optimizing the Design of Partitioning Method.....	250
6.7.5 Configuring the Netty Network Communication.....	251
6.7.6 Experience Summary.....	251
6.8 Common Flink Shell Commands.....	252
6.9 Reference.....	257
6.9.1 Example of Issuing a Certificate.....	258
6.10 Flink Restart Policy.....	262
7 Using Flume.....	264
7.1 Using Flume from Scratch.....	264

7.2 Overview.....	270
7.3 Installing the Flume Client.....	273
7.3.1 Installing the Flume Client on Clusters of Versions Earlier Than MRS 3.x.....	273
7.3.2 Installing the Flume Client on MRS 3.x or Later Clusters.....	277
7.4 Viewing Flume Client Logs.....	279
7.5 Stopping or Uninstalling the Flume Client.....	280
7.6 Using the Encryption Tool of the Flume Client.....	281
7.7 Flume Service Configuration Guide.....	281
7.8 Flume Configuration Parameter Description.....	314
7.9 Using Environment Variables in the properties.properties File.....	329
7.10 Non-Encrypted Transmission.....	331
7.10.1 Configuring Non-encrypted Transmission.....	331
7.10.2 Typical Scenario: Collecting Local Static Logs and Uploading Them to Kafka.....	336
7.10.3 Typical Scenario: Collecting Local Static Logs and Uploading Them to HDFS.....	338
7.10.4 Typical Scenario: Collecting Local Dynamic Logs and Uploading Them to HDFS.....	341
7.10.5 Typical Scenario: Collecting Logs from Kafka and Uploading Them to HDFS.....	344
7.10.6 Typical Scenario: Collecting Logs from Kafka and Uploading Them to HDFS Through the Flume Client.....	347
7.10.7 Typical Scenario: Collecting Local Static Logs and Uploading Them to HBase.....	352
7.11 Encrypted Transmission.....	360
7.11.1 Configuring the Encrypted Transmission.....	360
7.11.2 Typical Scenario: Collecting Local Static Logs and Uploading Them to HDFS.....	371
7.12 Viewing Flume Client Monitoring Information.....	386
7.13 Connecting Flume to Kafka in Security Mode.....	386
7.14 Connecting Flume with Hive in Security Mode.....	387
7.15 Configuring the Flume Service Model.....	390
7.15.1 Overview.....	390
7.15.2 Service Model Configuration Guide.....	390
7.16 Introduction to Flume Logs.....	396
7.17 Flume Client Cgroup Usage Guide.....	399
7.18 Secondary Development Guide for Flume Third-Party Plug-ins.....	400
7.19 Common Issues About Flume.....	401
8 Using HBase.....	403
8.1 Using HBase from Scratch.....	403
8.2 Using an HBase Client.....	407
8.3 Creating HBase Roles.....	409
8.4 Configuring HBase Replication.....	412
8.5 Configuring HBase Parameters.....	422
8.6 Enabling Cross-Cluster Copy.....	423
8.7 Using the ReplicationSyncUp Tool.....	424
8.8 Using HIndex.....	426
8.8.1 Introduction to HIndex.....	426

8.8.2 Loading Index Data in Batches.....	437
8.8.3 Using an Index Generation Tool.....	439
8.8.4 Migrating Index Data.....	442
8.9 Configuring HBase DR.....	444
8.10 Configuring HBase Data Compression and Encoding.....	454
8.11 Performing an HBase DR Service Switchover.....	456
8.12 Performing an HBase DR Active/Standby Cluster Switchover.....	458
8.13 Community BulkLoad Tool.....	460
8.14 Configuring the MOB.....	460
8.15 Configuring Secure HBase Replication.....	462
8.16 Configuring Region In Transition Recovery Chore Service.....	463
8.17 Using a Secondary Index.....	464
8.18 HBase Log Overview.....	465
8.19 HBase Performance Tuning.....	468
8.19.1 Improving the BulkLoad Efficiency.....	469
8.19.2 Improving Put Performance.....	469
8.19.3 Optimizing Put and Scan Performance.....	470
8.19.4 Improving Real-time Data Write Performance.....	474
8.19.5 Improving Real-time Data Read Performance.....	483
8.19.6 Optimizing JVM Parameters.....	490
8.20 Common Issues About HBase.....	491
8.20.1 Why Does a Client Keep Failing to Connect to a Server for a Long Time?.....	491
8.20.2 Operation Failures Occur in Stopping BulkLoad On the Client.....	493
8.20.3 Why May a Table Creation Exception Occur When HBase Deletes or Creates the Same Table Consecutively?.....	493
8.20.4 Why Other Services Become Unstable If HBase Sets up A Large Number of Connections over the Network Port?.....	494
8.20.5 Why Does the HBase BulkLoad Task (One Table Has 26 TB Data) Consisting of 210,000 Map Tasks and 10,000 Reduce Tasks Fail?.....	495
8.20.6 How Do I Restore a Region in the RIT State for a Long Time?.....	495
8.20.7 Why Does HMaster Exits Due to Timeout When Waiting for the Namespace Table to Go Online?	496
8.20.8 Why Does SocketTimeoutException Occur When a Client Queries HBase?.....	497
8.20.9 Why Modified and Deleted Data Can Still Be Queried by Using the Scan Command?.....	498
8.20.10 Why "java.lang.UnsatisfiedLinkError: Permission denied" exception thrown while starting HBase shell?.....	499
8.20.11 When does the RegionServers listed under "Dead Region Servers" on HMaster WebUI gets cleared?.....	499
8.20.12 Why Are Different Query Results Returned After I Use Same Query Criteria to Query Data Successfully Imported by HBase bulkload?.....	500
8.20.13 What Should I Do If I Fail to Create Tables Due to the FAILED_OPEN State of Regions?.....	500
8.20.14 How Do I Delete Residual Table Names in the /hbase/table-lock Directory of ZooKeeper?.....	501
8.20.15 Why Does HBase Become Faulty When I Set a Quota for the Directory Used by HBase in HDFS?	501

8.20.16 Why HMaster Times Out While Waiting for Namespace Table to be Assigned After Rebuilding Meta Using OfflineMetaRepair Tool and Startups Failed.....	503
8.20.17 Why Messages Containing FileNotFoundException and no lease Are Frequently Displayed in the HMaster Logs During the WAL Splitting Process?.....	503
8.20.18 Why Does the ImportTsv Tool Display "Permission denied" When the Same Linux User as and a Different Kerberos User from the Region Server Are Used?.....	505
8.20.19 Insufficient Rights When a Tenant Accesses Phoenix.....	506
8.20.20 What Can I Do When HBase Fails to Recover a Task and a Message Is Displayed Stating "Rollback recovery failed"?.....	506
8.20.21 How Do I Fix Region Overlapping?.....	507
8.20.22 Why Does RegionServer Fail to Be Started When GC Parameters Xms and Xmx of HBase RegionServer Are Set to 31 GB?.....	508
8.20.23 Why Does the LoadIncrementalHFiles Tool Fail to Be Executed and "Permission denied" Is Displayed When Nodes in a Cluster Are Used to Import Data in Batches?.....	508
8.20.24 Why Is the Error Message "import argparse" Displayed When the Phoenix sqlline Script Is Used?.....	510
8.20.25 How Do I Deal with the Restrictions of the Phoenix BulkLoad Tool?.....	510
8.20.26 Why a Message Is Displayed Indicating that the Permission is Insufficient When CTBase Connects to the Ranger Plug-ins?.....	511
9 Using HDFS.....	513
9.1 Using Hadoop from Scratch.....	513
9.2 Configuring Memory Management.....	516
9.3 Creating an HDFS Role.....	518
9.4 Using the HDFS Client.....	520
9.5 Running the DistCp Command.....	522
9.6 Overview of HDFS File System Directories.....	526
9.7 Changing the DataNode Storage Directory.....	534
9.8 Configuring HDFS Directory Permission.....	538
9.9 Configuring NFS.....	538
9.10 Planning HDFS Capacity.....	539
9.11 Configuring ulimit for HBase and HDFS.....	543
9.12 Balancing DataNode Capacity.....	544
9.13 Configuring Replica Replacement Policy for Heterogeneous Capacity Among DataNodes.....	549
9.14 Configuring the Number of Files in a Single HDFS Directory	550
9.15 Configuring the Recycle Bin Mechanism.....	551
9.16 Setting Permissions on Files and Directories.....	552
9.17 Setting the Maximum Lifetime and Renewal Interval of a Token.....	553
9.18 Configuring the Damaged Disk Volume.....	554
9.19 Configuring Encrypted Channels.....	554
9.20 Reducing the Probability of Abnormal Client Application Operation When the Network Is Not Stable	556
9.21 Configuring the NameNode Blacklist.....	556
9.22 Optimizing HDFS NameNode RPC QoS.....	559
9.23 Optimizing HDFS DataNode RPC QoS.....	562

9.24 Configuring Reserved Percentage of Disk Usage on DataNodes.....	562
9.25 Configuring HDFS NodeLabel.....	563
9.26 Configuring HDFS Mover.....	569
9.27 Using HDFS AZ Mover.....	571
9.28 Configuring HDFS DiskBalancer.....	572
9.29 Configuring the Observer NameNode to Process Read Requests.....	575
9.30 Performing Concurrent Operations on HDFS Files.....	576
9.31 Introduction to HDFS Logs.....	579
9.32 HDFS Performance Tuning.....	583
9.32.1 Improving Write Performance.....	583
9.32.2 Improving Read Performance Using Client Metadata Cache.....	584
9.32.3 Improving the Connection Between the Client and NameNode Using Current Active Cache.....	586
9.33 FAQ.....	587
9.33.1 NameNode Startup Is Slow.....	588
9.33.2 DataNode Is Normal but Cannot Report Data Blocks.....	588
9.33.3 HDFS WebUI Cannot Properly Update Information About Damaged Data.....	589
9.33.4 Why Does the Distcp Command Fail in the Secure Cluster, Causing an Exception?.....	590
9.33.5 Why Does DataNode Fail to Start When the Number of Disks Specified by dfs.datanode.data.dir Equals dfs.datanode.failed.volumes.tolerated?.....	590
9.33.6 Failed to Calculate the Capacity of a DataNode when Multiple data.dir Directories Are Configured in a Disk Partition.....	591
9.33.7 Standby NameNode Fails to Be Restarted When the System Is Powered off During Metadata (Namespace) Storage.....	592
9.33.8 Why Data in the Buffer Is Lost If a Power Outage Occurs During Storage of Small Files.....	593
9.33.9 Why Does Array Border-crossing Occur During FileInputFormat Split?.....	593
9.33.10 Why Is the Storage Type of File Copies DISK When the Tiered Storage Policy Is LAZY_PERSIST?.....	594
9.33.11 The HDFS Client Is Unresponsive When the NameNode Is Overloaded for a Long Time.....	594
9.33.12 Can I Delete or Modify the Data Storage Directory in DataNode?.....	595
9.33.13 Blocks Miss on the NameNode UI After the Successful Rollback.....	596
9.33.14 Why Is "java.net.SocketException: No buffer space available" Reported When Data Is Written to HDFS.....	597
9.33.15 Why are There Two Standby NameNodes After the active NameNode Is Restarted?.....	598
9.33.16 When Does a Balance Process in HDFS, Shut Down and Fail to be Executed Again?.....	600
9.33.17 "This page can't be displayed" Is Displayed When Internet Explorer Fails to Access the Native HDFS UI.....	600
9.33.18 NameNode Fails to Be Restarted Due to EditLog Discontinuity.....	601
10 Using Hive.....	603
10.1 Using Hive from Scratch.....	603
10.2 Configuring Hive Parameters.....	607
10.3 Hive SQL.....	609
10.4 Permission Management.....	612
10.4.1 Hive Permission.....	612
10.4.2 Creating a Hive Role.....	615

10.4.3 Configuring Permissions for Hive Tables, Columns, or Databases.....	620
10.4.4 Configuring Permissions to Use Other Components for Hive.....	624
10.5 Using a Hive Client.....	628
10.6 Using HDFS Colocation to Store Hive Tables.....	631
10.7 Using the Hive Column Encryption Function.....	633
10.8 Customizing Row Separators.....	634
10.9 Configuring Hive on HBase in Across Clusters with Mutual Trust Enabled.....	635
10.10 Deleting Single-Row Records from Hive on HBase.....	636
10.11 Configuring HTTPS/HTTP-based REST APIs.....	636
10.12 Enabling or Disabling the Transform Function.....	637
10.13 Access Control of a Dynamic Table View on Hive.....	638
10.14 Specifying Whether the ADMIN Permissions Is Required for Creating Temporary Functions.....	639
10.15 Using Hive to Read Data in a Relational Database.....	640
10.16 Supporting Traditional Relational Database Syntax in Hive.....	641
10.17 Creating User-Defined Hive Functions.....	643
10.18 Enhancing beeline Reliability.....	645
10.19 Viewing Table Structures Using the show create Statement as Users with the select Permission...	647
10.20 Writing a Directory into Hive with the Old Data Removed to the Recycle Bin.....	647
10.21 Inserting Data to a Directory That Does Not Exist.....	648
10.22 Creating Databases and Creating Tables in the Default Database Only as the Hive Administrator	649
10.23 Disabling of Specifying the location Keyword When Creating an Internal Hive Table.....	650
10.24 Enabling the Function of Creating a Foreign Table in a Directory That Can Only Be Read.....	651
10.25 Authorizing Over 32 Roles in Hive.....	652
10.26 Restricting the Maximum Number of Maps for Hive Tasks.....	653
10.27 HiveServer Lease Isolation.....	653
10.28 Hive Supporting Transactions.....	654
10.29 Switching the Hive Execution Engine to Tez.....	660
10.30 Hive Materialized View.....	661
10.31 Hive Log Overview.....	664
10.32 Hive Performance Tuning.....	668
10.32.1 Creating Table Partitions.....	668
10.32.2 Optimizing Join.....	669
10.32.3 Optimizing Group By.....	672
10.32.4 Optimizing Data Storage.....	672
10.32.5 Optimizing SQL Statements.....	673
10.32.6 Optimizing the Query Function Using Hive CBO.....	674
10.33 Common Issues About Hive.....	676
10.33.1 How Do I Delete UDFs on Multiple HiveServers at the Same Time?.....	676
10.33.2 Why Cannot the DROP operation Be Performed on a Backed-up Hive Table?.....	677
10.33.3 How to Perform Operations on Local Files with Hive User-Defined Functions.....	678
10.33.4 How Do I Forcibly Stop MapReduce Jobs Executed by Hive?.....	678
10.33.5 Table Creation Fails Because Hive Complex Fields' Names Contain Special Characters.....	679

10.33.6 How Do I Monitor the Hive Table Size?.....	679
10.33.7 How Do I Prevent Key Directories from Data Loss Caused by Misoperations of the insert overwrite Statement?.....	680
10.33.8 Why Is Hive on Spark Task Freezing When HBase Is Not Installed?.....	680
10.33.9 Error Reported When the WHERE Condition Is Used to Query Tables with Excessive Partitions in FusionInsight Hive.....	681
10.33.10 Why Cannot I Connect to HiveServer When I Use IBM JDK to Access the Beeline Client?.....	682
10.33.11 Description of Hive Table Location (Either Be an OBS or HDFS Path).....	682
10.33.12 Why Cannot Data Be Queried After the MapReduce Engine Is Switched After the Tez Engine Is Used to Execute Union-related Statements?.....	683
10.33.13 Why Does Hive Not Support Concurrent Data Writing to the Same Table or Partition?.....	683
10.33.14 Why Does Hive Not Support Vectorized Query?.....	683
10.33.15 Why Does Metadata Still Exist When the HDFS Data Directory of the Hive Table Is Deleted by Mistake?.....	684
10.33.16 How Do I Disable the Logging Function of Hive?.....	684
10.33.17 Why Hive Tables in the OBS Directory Fail to Be Deleted?.....	685
10.33.18 Hive Configuration Problems.....	685
11 Using Hudi.....	687
11.1 Getting Started.....	687
11.2 Basic Operations.....	690
11.2.1 Hudi Table Schema.....	690
11.2.2 Write.....	691
11.2.2.1 Before You Start.....	691
11.2.2.2 Batch Write.....	691
11.2.2.3 Stream Write.....	694
11.2.2.4 Synchronizing Hudi Table Data to Hive.....	696
11.2.3 Read.....	698
11.2.3.1 Overview.....	698
11.2.3.2 Reading COW Table Views.....	699
11.2.3.3 Reading MOR Table Views.....	700
11.2.4 Data Management and Maintenance.....	701
11.2.4.1 Clustering.....	701
11.2.4.2 Cleaning.....	703
11.2.4.3 Compaction.....	704
11.2.4.4 Savepoint.....	705
11.2.4.5 Single-Table Concurrency Control.....	705
11.2.5 Using the Hudi Client.....	707
11.2.5.1 Operating a Hudi Table Using hudi-cli.sh.....	707
11.2.6 Configuration Reference.....	709
11.2.6.1 Write Configuration.....	709
11.2.6.2 Configuration of Hive Table Synchronization.....	712
11.2.6.3 Index Configuration.....	713
11.2.6.4 Storage Configuration.....	716

11.2.6.5	Compaction and Cleaning Configurations.....	717
11.2.6.6	Single-Table Concurrency Control Configuration.....	721
11.3	Hudi Performance Tuning.....	722
11.4	Common Issues About Hudi.....	723
11.4.1	Data Write.....	723
11.4.1.1	Parquet/Avro schema Is Reported When Updated Data Is Written.....	723
11.4.1.2	UnsupportedOperationException Is Reported When Updated Data Is Written.....	723
11.4.1.3	SchemaCompatabilityException Is Reported When Updated Data Is Written.....	724
11.4.1.4	What Should I Do If Hudi Consumes Much Space in a Temporary Folder During Upsert?.....	724
11.4.1.5	Hudi Fails to Write Decimal Data with Lower Precision.....	724
11.4.2	Data Collection.....	725
11.4.2.1	IllegalArgumentException Is Reported When Kafka Is Used to Collect Data.....	725
11.4.2.2	HoodieException Is Reported When Data Is Collected.....	725
11.4.2.3	HoodieKeyException Is Reported When Data Is Collected.....	726
11.4.3	Hive Synchronization.....	726
11.4.3.1	SQLException Is Reported During Hive Data Synchronization.....	726
11.4.3.2	HoodieHiveSyncException Is Reported During Hive Data Synchronization.....	726
11.4.3.3	SemanticException Is Reported During Hive Data Synchronization.....	727
12	Using Hue (Versions Earlier Than MRS 3.x).....	728
12.1	Using Hue from Scratch.....	728
12.2	Accessing the Hue Web UI.....	729
12.3	Hue Common Parameters.....	730
12.4	Using HiveQL Editor on the Hue Web UI.....	731
12.5	Using the Metadata Browser on the Hue Web UI.....	733
12.6	Using File Browser on the Hue Web UI.....	737
12.7	Using Job Browser on the Hue Web UI.....	740
13	Using Hue (MRS 3.x or Later).....	742
13.1	Using Hue from Scratch.....	742
13.2	Accessing the Hue Web UI.....	743
13.3	Hue Common Parameters.....	744
13.4	Using HiveQL Editor on the Hue Web UI.....	745
13.5	Using the SparkSql Editor on the Hue Web UI.....	747
13.6	Using the Metadata Browser on the Hue Web UI.....	749
13.7	Using File Browser on the Hue Web UI.....	750
13.8	Using Job Browser on the Hue Web UI.....	752
13.9	Using HBase on the Hue Web UI.....	754
13.10	Typical Scenarios.....	754
13.10.1	HDFS on Hue.....	755
13.10.2	Hive on Hue.....	758
13.10.3	Oozie on Hue.....	760
13.11	Hue Log Overview.....	761
13.12	Common Issues About Hue.....	764

13.12.1 Why Do HQL Statements Fail to Execute in Hue Using Internet Explorer?.....	764
13.12.2 Why Does the use database Statement Become Invalid in Hive?.....	765
13.12.3 Why Do HDFS Files Fail to Access Through the Hue Web UI?.....	765
13.12.4 Why Do Large Files Fail to Upload on the Hue Page.....	766
13.12.5 Why Is the Hue Native Page Cannot Be Properly Displayed If the Hive Service Is Not Installed in a Cluster?.....	767
14 Using Impala.....	768
14.1 Using Impala from Scratch.....	768
14.2 Common Impala Parameters.....	771
14.3 Accessing the Impala Web UI.....	772
14.4 Using Impala to Operate Kudu.....	773
14.5 Interconnecting Impala with External LDAP.....	774
14.6 Enabling and Configuring a Dynamic Resource Pool for Impala.....	775
15 Using Kafka.....	779
15.1 Using Kafka from Scratch.....	779
15.2 Managing Kafka Topics.....	781
15.3 Querying Kafka Topics.....	785
15.4 Managing Kafka User Permissions.....	786
15.5 Managing Messages in Kafka Topics.....	790
15.6 Synchronizing Binlog-based MySQL Data to the MRS Cluster.....	791
15.7 Creating a Kafka Role.....	797
15.8 Kafka Common Parameters.....	798
15.9 Safety Instructions on Using Kafka.....	803
15.10 Kafka Specifications.....	806
15.11 Using the Kafka Client.....	807
15.12 Configuring Kafka HA and High Reliability Parameters.....	808
15.13 Changing the Broker Storage Directory.....	814
15.14 Checking the Consumption Status of Consumer Group.....	816
15.15 Kafka Balancing Tool Instructions.....	817
15.16 Balancing Data After Kafka Node Scale-Out.....	821
15.17 Kafka Token Authentication Mechanism Tool Usage.....	823
15.18 Introduction to Kafka Logs.....	824
15.19 Performance Tuning.....	827
15.19.1 Kafka Performance Tuning.....	827
15.20 Kafka Feature Description.....	828
15.21 Migrating Data Between Kafka Nodes.....	831
15.22 Common Issues About Kafka.....	833
15.22.1 How Do I Solve the Problem that Kafka Topics Cannot Be Deleted?.....	833
16 Using KafkaManager.....	835
16.1 Introduction to KafkaManager.....	835
16.2 Accessing the KafkaManager Web UI.....	835
16.3 Managing Kafka Clusters.....	836

16.4 Kafka Cluster Monitoring Management.....	839
17 Using Loader.....	848
17.1 Using Loader from Scratch.....	848
17.2 How to Use Loader.....	849
17.3 Common Loader Parameters.....	850
17.4 Creating a Loader Role.....	852
17.5 Loader Link Configuration.....	855
17.6 Managing Loader Links (Versions Earlier Than MRS 3.x).....	857
17.7 Managing Loader Links (MRS 3.x and Later Versions).....	858
17.8 Source Link Configurations of Loader Jobs.....	863
17.9 Destination Link Configurations of Loader Jobs.....	866
17.10 Managing Loader Jobs.....	869
17.11 Preparing a Driver for MySQL Database Link.....	872
17.12 Importing Data.....	874
17.12.1 Overview.....	874
17.12.2 Importing Data Using Loader.....	876
17.12.3 Typical Scenario: Importing Data from an SFTP Server to HDFS or OBS.....	893
17.12.4 Typical Scenario: Importing Data from an SFTP Server to HBase.....	900
17.12.5 Typical Scenario: Importing Data from an SFTP Server to Hive.....	908
17.12.6 Typical Scenario: Importing Data from an FTP Server to HBase.....	915
17.12.7 Typical Scenario: Importing Data from a Relational Database to HDFS or OBS.....	922
17.12.8 Typical Scenario: Importing Data from a Relational Database to HBase.....	929
17.12.9 Typical Scenario: Importing Data from a Relational Database to Hive.....	936
17.12.10 Typical Scenario: Importing Data from HDFS or OBS to HBase.....	942
17.12.11 Typical Scenario: Importing Data from a Relational Database to ClickHouse.....	946
17.12.12 Typical Scenario: Importing Data from HDFS to ClickHouse.....	951
17.13 Exporting Data.....	955
17.13.1 Overview.....	955
17.13.2 Using Loader to Export Data.....	957
17.13.3 Typical Scenario: Exporting Data from HDFS or OBS to an SFTP Server.....	969
17.13.4 Typical Scenario: Exporting Data from HBase to an SFTP Server.....	976
17.13.5 Typical Scenario: Exporting Data from Hive to an SFTP Server.....	981
17.13.6 Typical Scenario: Exporting Data from HDFS or OBS to a Relational Database.....	986
17.13.7 Typical Scenario: Exporting Data from HBase to a Relational Database.....	992
17.13.8 Typical Scenario: Exporting Data from Hive to a Relational Database.....	997
17.13.9 Typical Scenario: Importing Data from HBase to HDFS or OBS.....	1002
17.14 Managing Jobs.....	1005
17.14.1 Migrating Loader Jobs in Batches.....	1006
17.14.2 Deleting Loader Jobs in Batches.....	1007
17.14.3 Importing Loader Jobs in Batches.....	1007
17.14.4 Exporting Loader Jobs in Batches.....	1008
17.14.5 Viewing Historical Job Information.....	1009

17.15 Operator Help.....	1010
17.15.1 Overview.....	1011
17.15.2 Input Operators.....	1014
17.15.2.1 CSV File Input.....	1014
17.15.2.2 Fixed File Input.....	1016
17.15.2.3 Table Input.....	1018
17.15.2.4 HBase Input.....	1020
17.15.2.5 HTML Input.....	1023
17.15.2.6 Hive input.....	1026
17.15.2.7 Spark Input.....	1028
17.15.3 Conversion Operators.....	1030
17.15.3.1 Long Date Conversion.....	1030
17.15.3.2 Null Value Conversion.....	1032
17.15.3.3 Constant Field Addition.....	1033
17.15.3.4 Random Value Conversion.....	1035
17.15.3.5 Concat Fields.....	1036
17.15.3.6 Extract Fields.....	1038
17.15.3.7 Modulo Integer.....	1040
17.15.3.8 String Cut.....	1041
17.15.3.9 EL Operation.....	1043
17.15.3.10 String Operations.....	1045
17.15.3.11 String Reverse.....	1047
17.15.3.12 String Trim.....	1048
17.15.3.13 Filter Rows.....	1050
17.15.3.14 Update Fields Operator.....	1051
17.15.4 Output Operators.....	1053
17.15.4.1 Hive output.....	1053
17.15.4.2 Spark Output.....	1056
17.15.4.3 Table Output	1059
17.15.4.4 File Output.....	1061
17.15.4.5 HBase Output.....	1063
17.15.4.6 ClickHouse Output.....	1065
17.15.5 Associating, Editing, Importing, or Exporting the Field Configuration of an Operator.....	1068
17.15.6 Using Macro Definitions in Configuration Items	1071
17.15.7 Operator Data Processing Rules.....	1072
17.16 Client Tools.....	1076
17.16.1 Running a Loader Job by Using Commands.....	1077
17.16.2 loader-tool Usage Guide.....	1081
17.16.3 loader-tool Usage Example.....	1090
17.16.4 schedule-tool Usage Guide.....	1093
17.16.5 schedule-tool Usage Example.....	1097
17.16.6 Using loader-backup to Back Up Job Data.....	1100

17.16.7 Open Source sqoop-shell Tool Usage Guide.....	1103
17.16.8 Example for Using the Open-Source sqoop-shell Tool (SFTP-HDFS).....	1115
17.16.9 Example for Using the Open-Source sqoop-shell Tool (Oracle-HBase).....	1125
17.17 Loader Log Overview.....	1135
17.18 Example: Using Loader to Import Data from OBS to HDFS.....	1138
17.19 Common Issues About Loader.....	1139
17.19.1 Data Cannot Be Saved in Internet Explorer 10 or 11.....	1139
17.19.2 Differences Among Connectors Used During the Process of Importing Data from the Oracle Database to HDFS.....	1140
18 Using Kudu.....	1142
18.1 Using Kudu from Scratch.....	1142
18.2 Accessing the Kudu Web UI.....	1143
19 Using MapReduce.....	1145
19.1 Configuring the Log Archiving and Clearing Mechanism.....	1145
19.2 Reducing Client Application Failure Rate.....	1147
19.3 Transmitting MapReduce Tasks from Windows to Linux.....	1148
19.4 Configuring the Distributed Cache.....	1149
19.5 Configuring the MapReduce Shuffle Address.....	1151
19.6 Configuring the Cluster Administrator List.....	1152
19.7 Introduction to MapReduce Logs.....	1153
19.8 MapReduce Performance Tuning.....	1156
19.8.1 Optimization Configuration for Multiple CPU Cores.....	1156
19.8.2 Determining the Job Baseline.....	1161
19.8.3 Streamlining Shuffle.....	1163
19.8.4 AM Optimization for Big Tasks.....	1167
19.8.5 Speculative Execution.....	1168
19.8.6 Using Slow Start.....	1169
19.8.7 Optimizing Performance for Committing MR Jobs.....	1169
19.9 Common Issues About MapReduce.....	1170
19.9.1 Why Does a MapReduce Task Stay Unchanged for a Long Time?.....	1170
19.9.2 Why the Client Hangs During Job Running?.....	1171
19.9.3 Why Cannot HDFS_DELEGATION_TOKEN Be Found in the Cache?.....	1171
19.9.4 How Do I Set the Task Priority When Submitting a MapReduce Task?.....	1172
19.9.5 Why Physical Memory Overflow Occurs If a MapReduce Task Fails?.....	1172
19.9.6 After the Address of MapReduce JobHistoryServer Is Changed, Why the Wrong Page is Displayed When I Click the Tracking URL on the ResourceManager WebUI?.....	1173
19.9.7 MapReduce Job Failed in Multiple NameService Environment.....	1174
19.9.8 Why a Fault MapReduce Node Is Not Blacklisted?.....	1175
20 Using OpenTSDB.....	1176
20.1 Using an MRS Client to Operate OpenTSDB Metric Data.....	1176
20.2 Running the curl Command to Operate OpenTSDB.....	1179
21 Using Oozie.....	1181

21.1 Using Oozie from Scratch.....	1181
21.2 Using the Oozie Client.....	1182
21.3 Using Oozie Client to Submit an Oozie Job.....	1184
21.3.1 Submitting a Hive Job.....	1184
21.3.2 Submitting a Spark2x Job.....	1186
21.3.3 Submitting a Loader Job.....	1188
21.3.4 Submitting a DistCp Job.....	1190
21.3.5 Submitting Other Jobs.....	1192
21.4 Using Hue to Submit an Oozie Job.....	1195
21.4.1 Creating a Workflow.....	1195
21.4.2 Submitting a Workflow Job.....	1196
21.4.2.1 Submitting a Hive2 Job.....	1197
21.4.2.2 Submitting a Spark2x Job.....	1198
21.4.2.3 Submitting a Java Job.....	1200
21.4.2.4 Submitting a Loader Job.....	1200
21.4.2.5 Submitting a MapReduce Job.....	1201
21.4.2.6 Submitting a Sub-workflow Job.....	1202
21.4.2.7 Submitting a Shell Job.....	1203
21.4.2.8 Submitting an HDFS Job.....	1205
21.4.2.9 Submitting a Streaming Job.....	1205
21.4.2.10 Submitting a DistCp Job.....	1206
21.4.2.11 Example of Mutual Trust Operations.....	1208
21.4.2.12 Submitting an SSH Job.....	1208
21.4.2.13 Submitting a Hive Script.....	1209
21.4.3 Submitting a Coordinator Periodic Scheduling Job.....	1210
21.4.4 Submitting a Bundle Batch Processing Job.....	1211
21.4.5 Querying Job Execution Results.....	1212
21.5 Oozie Log Overview.....	1213
21.6 Common Issues About Oozie.....	1216
21.6.1 Oozie Scheduled Tasks Are Not Executed on Time.....	1216
21.6.2 Why Update of the share lib Directory of Oozie on HDFS Does Not Take Effect?.....	1216
21.6.3 Common Oozie Troubleshooting Methods.....	1216
22 Using Presto.....	1218
22.1 Accessing the Presto Web UI.....	1218
22.2 Using a Client to Execute Query Statements.....	1219
22.3 Presto FAQ.....	1221
22.3.1 How Do I Configure Multiple Hive Connections for Presto?.....	1221
23 Using Ranger (MRS 1.9.2).....	1223
23.1 Creating a Ranger Cluster.....	1223
23.2 Accessing the Ranger Web UI and Synchronizing Unix Users to the Ranger Web UI.....	1224
23.3 Configuring Hive/Impala Access Permissions in Ranger.....	1226
23.4 Configuring HBase Access Permissions in Ranger.....	1231

24 Using Ranger (MRS 3.x)	1237
24.1 Logging In to the Ranger Web UI	1237
24.2 Enabling Ranger Authentication	1239
24.3 Configuring Component Permission Policies	1240
24.4 Viewing Ranger Audit Information	1242
24.5 Configuring a Security Zone	1243
24.6 Viewing Ranger Permission Information	1246
24.7 Adding a Ranger Access Permission Policy for HDFS	1248
24.8 Adding a Ranger Access Permission Policy for HBase	1252
24.9 Adding a Ranger Access Permission Policy for Hive	1256
24.10 Adding a Ranger Access Permission Policy for Yarn	1266
24.11 Adding a Ranger Access Permission Policy for Spark2x	1269
24.12 Adding a Ranger Access Permission Policy for Kafka	1278
24.13 Adding a Ranger Access Permission Policy for Storm	1287
24.14 Ranger Log Overview	1290
24.15 Common Issues About Ranger	1293
24.15.1 Why Ranger Startup Fails During the Cluster Installation?	1293
24.15.2 How Do I Determine Whether the Ranger Authentication Is Used for a Service?	1293
24.15.3 Why Cannot a New User Log In to Ranger After Changing the Password?	1294
24.15.4 When an HBase Policy Is Added or Modified on Ranger, Wildcard Characters Cannot Be Used to Search for Existing HBase Tables	1294
25 Using Spark	1296
25.1 Precautions	1296
25.2 Getting Started with Spark	1296
25.3 Getting Started with Spark SQL	1298
25.4 Using the Spark Client	1300
25.5 Accessing the Spark Web UI	1301
25.6 Interconnecting Spark with OpenTSDB	1303
25.6.1 Creating a Table and Associating It with OpenTSDB	1303
25.6.2 Inserting Data to the OpenTSDB Table	1304
25.6.3 Querying an OpenTSDB Table	1304
25.6.4 Modifying the Default Configuration Data	1305
26 Using Spark2x	1307
26.1 Precautions	1307
26.2 Basic Operation	1307
26.2.1 Getting Started	1307
26.2.2 Configuring Parameters Rapidly	1310
26.2.3 Common Parameters	1319
26.2.4 Spark on HBase Overview and Basic Applications	1343
26.2.5 Spark on HBase V2 Overview and Basic Applications	1345
26.2.6 SparkSQL Permission Management(Security Mode)	1347
26.2.6.1 Spark SQL Permissions	1347

26.2.6.2	Creating a Spark SQL Role.....	1352
26.2.6.3	Configuring Permissions for SparkSQL Tables, Columns, and Databases.....	1356
26.2.6.4	Configuring Permissions for SparkSQL to Use Other Components.....	1358
26.2.6.5	Configuring the Client and Server.....	1360
26.2.7	Scenario-Specific Configuration.....	1362
26.2.7.1	Configuring Multi-active Instance Mode.....	1362
26.2.7.2	Configuring the Multi-tenant Mode.....	1363
26.2.7.3	Configuring the Switchover Between the Multi-active Instance Mode and the Multi-tenant Mode.....	1365
26.2.7.4	Configuring the Size of the Event Queue.....	1366
26.2.7.5	Configuring Executor Off-Heap Memory.....	1367
26.2.7.6	Enhancing Stability in a Limited Memory Condition.....	1368
26.2.7.7	Viewing Aggregated Container Logs on the Web UI.....	1369
26.2.7.8	Configuring Environment Variables in Yarn-Client and Yarn-Cluster Modes.....	1371
26.2.7.9	Configuring the Default Number of Data Blocks Divided by SparkSQL.....	1373
26.2.7.10	Configuring the Compression Format of a Parquet Table.....	1373
26.2.7.11	Configuring the Number of Lost Executors Displayed in WebUI.....	1374
26.2.7.12	Setting the Log Level Dynamically.....	1374
26.2.7.13	Configuring Whether Spark Obtains HBase Tokens.....	1376
26.2.7.14	Configuring LIFO for Kafka.....	1377
26.2.7.15	Configuring Reliability for Connected Kafka.....	1379
26.2.7.16	Configuring Streaming Reading of Driver Execution Results.....	1380
26.2.7.17	Filtering Partitions without Paths in Partitioned Tables.....	1382
26.2.7.18	Configuring Spark2x Web UI ACLs.....	1383
26.2.7.19	Configuring Vector-based ORC Data Reading.....	1385
26.2.7.20	Broaden Support for Hive Partition Pruning Predicate Pushdown.....	1387
26.2.7.21	Hive Dynamic Partition Overwriting Syntax.....	1387
26.2.7.22	Configuring the Column Statistics Histogram to Enhance the CBO Accuracy.....	1388
26.2.7.23	Configuring Local Disk Cache for JobHistory.....	1391
26.2.7.24	Configuring Spark SQL to Enable the Adaptive Execution Feature.....	1391
26.2.7.25	Configuring Event Log Rollover.....	1394
26.2.8	Adapting to the Third-party JDK When Ranger Is Used.....	1396
26.3	Spark2x Logs.....	1396
26.4	Obtaining Container Logs of a Running Spark Application.....	1400
26.5	Small File Combination Tools.....	1400
26.6	Using CarbonData for First Query.....	1403
26.7	Spark2x Performance Tuning.....	1404
26.7.1	Spark Core Tuning.....	1405
26.7.1.1	Data Serialization.....	1405
26.7.1.2	Optimizing Memory Configuration.....	1406
26.7.1.3	Setting the DOP.....	1406
26.7.1.4	Using Broadcast Variables.....	1407
26.7.1.5	Using the external shuffle service to improve performance.....	1407

26.7.1.6 Configuring Dynamic Resource Scheduling in Yarn Mode.....	1408
26.7.1.7 Configuring Process Parameters.....	1409
26.7.1.8 Designing the Direction Acyclic Graph (DAG).....	1411
26.7.1.9 Experience.....	1413
26.7.2 Spark SQL and DataFrame Tuning.....	1415
26.7.2.1 Optimizing the Spark SQL Join Operation.....	1415
26.7.2.2 Improving Spark SQL Calculation Performance Under Data Skew.....	1417
26.7.2.3 Optimizing Spark SQL Performance in the Small File Scenario.....	1419
26.7.2.4 Optimizing the INSERT...SELECT Operation.....	1420
26.7.2.5 Multiple JDBC Clients Concurrently Connecting to JDBCServer.....	1421
26.7.2.6 Optimizing Memory when Data Is Inserted into Dynamic Partitioned Tables.....	1421
26.7.2.7 Optimizing Small Files.....	1422
26.7.2.8 Optimizing the Aggregate Algorithms.....	1423
26.7.2.9 Optimizing Datasource Tables.....	1424
26.7.2.10 Merging CBO.....	1425
26.7.2.11 Optimizing SQL Query of Data of Multiple Sources.....	1427
26.7.2.12 SQL Optimization for Multi-level Nesting and Hybrid Join.....	1430
26.7.3 Spark Streaming Tuning.....	1432
26.8 Common Issues About Spark2x.....	1433
26.8.1 Spark Core.....	1433
26.8.1.1 How Do I View Aggregated Spark Application Logs?.....	1434
26.8.1.2 Why Cannot Exit the Driver Process?.....	1434
26.8.1.3 Why Does FetchFailedException Occur When the Network Connection Is Timed out.....	1434
26.8.1.4 How to Configure Event Queue Size If Event Queue Overflows?.....	1435
26.8.1.5 What Can I Do If the getApplicationReport Exception Is Recorded in Logs During Spark Application Execution and the Application Does Not Exit for a Long Time?.....	1436
26.8.1.6 What Can I Do If "Connection to ip:port has been quiet for xxx ms while there are outstanding requests" Is Reported When Spark Executes an Application and the Application Ends?.....	1437
26.8.1.7 Why Do Executors Fail to be Removed After the NodeManager Is Shut Down?.....	1439
26.8.1.8 What Can I Do If the Message "Password cannot be null if SASL is enabled" Is Displayed?.....	1439
26.8.1.9 What Should I Do If the Message "Failed to CREATE_FILE" Is Displayed in the Restarted Tasks When Data Is Inserted Into the Dynamic Partition Table?.....	1440
26.8.1.10 Why Tasks Fail When Hash Shuffle Is Used?.....	1440
26.8.1.11 What Can I Do If the Error Message "DNS query failed" Is Displayed When I Access the Aggregated Logs Page of Spark Applications?.....	1441
26.8.1.12 What Can I Do If Shuffle Fetch Fails Due to the "Timeout Waiting for Task" Exception?.....	1442
26.8.1.13 Why Does the Stage Retry due to the Crash of the Executor?.....	1442
26.8.1.14 Why Do the Executors Fail to Register Shuffle Services During the Shuffle of a Large Amount of Data?.....	1443
26.8.1.15 Why Does the Out of Memory Error Occur in NodeManager During the Execution of Spark Applications.....	1444
26.8.1.16 Why Does the Realm Information Fail to Be Obtained When SparkBench is Run on HiBench for the Cluster in Security Mode?.....	1446
26.8.2 Spark SQL and DataFrame.....	1446

26.8.2.1 What Do I have to Note When Using Spark SQL ROLLUP and CUBE?.....	1446
26.8.2.2 Why Spark SQL Is Displayed as a Temporary Table in Different Databases?.....	1447
26.8.2.3 How to Assign a Parameter Value in a Spark Command?.....	1448
26.8.2.4 What Directory Permissions Do I Need to Create a Table Using SparkSQL?.....	1449
26.8.2.5 Why Do I Fail to Delete the UDF Using Another Service?.....	1450
26.8.2.6 Why Cannot I Query Newly Inserted Data in a Parquet Hive Table Using SparkSQL?.....	1450
26.8.2.7 How to Use Cache Table?.....	1451
26.8.2.8 Why Are Some Partitions Empty During Repartition?.....	1451
26.8.2.9 Why Does 16 Terabytes of Text Data Fails to Be Converted into 4 Terabytes of Parquet Data?.....	1452
26.8.2.10 Why the Operation Fails When the Table Name Is TABLE?.....	1454
26.8.2.11 Why Is a Task Suspended When the ANALYZE TABLE Statement Is Executed and Resources Are Insufficient?.....	1454
26.8.2.12 If I Access a parquet Table on Which I Do not Have Permission, Why a Job Is Run Before "Missing Privileges" Is Displayed?.....	1455
26.8.2.13 Why Do I Fail to Modify MetaData by Running the Hive Command?.....	1455
26.8.2.14 Why Is "RejectedExecutionException" Displayed When I Exit Spark SQL?.....	1455
26.8.2.15 How Do I Do If I Incidentally Kill the JDBCServer Process During Health Check?.....	1456
26.8.2.16 Why No Result Is found When 2016-6-30 Is Set in the Date Field as the Filter Condition?.....	1456
26.8.2.17 Why Does the "--hivevar" Option I Specified in the Command for Starting spark-beeline Fail to Take Effect?.....	1457
26.8.2.18 Why Does the "Permission denied" Exception Occur When I Create a Temporary Table or View in Spark-beeline?.....	1458
26.8.2.19 Why Is the "Code of method ... grows beyond 64 KB" Error Message Displayed When I Run Complex SQL Statements?.....	1458
26.8.2.20 Why Is Memory Insufficient if 10 Terabytes of TPCDS Test Suites Are Consecutively Run in Beeline/JDBCServer Mode?.....	1459
26.8.2.21 Why Functions Cannot Be Used When Different JDBCServer Are Connected?.....	1459
26.8.2.22 Why Does an Exception Occur When I Drop Functions Created Using the Add Jar Statement?.....	1461
26.8.2.23 Why Does Spark2x Have No Access to DataSource Tables Created by Spark1.5?.....	1462
26.8.2.24 Why Does Spark-beeline Fail to Run and Error Message "Failed to create ThriftService instance" Is Displayed?.....	1463
26.8.2.25 Why Cannot I Query Newly Inserted Data in an ORC Hive Table Using Spark SQL?.....	1464
26.8.3 Spark Streaming.....	1465
26.8.3.1 Same DAG Log Is Recorded Twice for a Streaming Task.....	1465
26.8.3.2 What Can I Do If Spark Streaming Tasks Are Blocked?.....	1466
26.8.3.3 What Should I Pay Attention to When Optimizing Spark Streaming Task Parameters?.....	1467
26.8.3.4 Why Does the Spark Streaming Application Fail to Be Submitted After the Token Validity Period Expires?.....	1468
26.8.3.5 Why does Spark Streaming Application Fail to Restart from Checkpoint When It Creates an Input Stream Without Output Logic?.....	1469
26.8.3.6 Why Is the Input Size Corresponding to Batch Time on the Web UI Set to 0 Records When Kafka Is Restarted During Spark Streaming Running?.....	1470
26.8.4 Why the Job Information Obtained from the restful Interface of an Ended Spark Application Is Incorrect?.....	1471

26.8.5 Why Cannot I Switch from the Yarn Web UI to the Spark Web UI?.....	1472
26.8.6 What Can I Do If an Error Occurs when I Access the Application Page Because the Application Cached by HistoryServer Is Recycled?.....	1473
26.8.7 Why Is not an Application Displayed When I Run the Application with the Empty Part File?.....	1474
26.8.8 Why Does Spark2x Fail to Export a Table with the Same Field Name?.....	1475
26.8.9 Why JRE fatal error after running Spark application multiple times?.....	1475
26.8.10 Native Spark2x UI Fails to Be Accessed or Is Incorrectly Displayed when Internet Explorer Is Used for Access.....	1475
26.8.11 How Does Spark2x Access External Cluster Components?.....	1476
26.8.12 Why Does the Foreign Table Query Fail When Multiple Foreign Tables Are Created in the Same Directory?.....	1478
26.8.13 Why Is the Native Page of an Application in Spark2x JobHistory Displayed Incorrectly?.....	1478
26.8.14 Why Do I Fail to Create a Table in the Specified Location on OBS After Logging to spark-beeline?	1479
26.8.15 Spark Shuffle Exception Handling.....	1480
27 Using Sqoop.....	1481
27.1 Using Sqoop from Scratch.....	1481
27.2 Adapting Sqoop 1.4.7 to MRS 3.x Clusters.....	1485
27.3 Common Sqoop Commands and Parameters.....	1487
27.4 Common Issues About Sqoop.....	1490
27.4.1 What Should I Do If Class QueryProvider Is Unavailable?.....	1490
27.4.2 How Do I Do If PostgreSQL or GaussDB Fails to Connect?.....	1491
27.4.3 What Should I Do If Data Failed to Be Synchronized to a Hive Table on the OBS Using hive-table?	1493
27.4.4 What Should I Do If Data Failed to Be Synchronized to an ORC or Parquet Table Using hive-table?	1493
27.4.5 What Should I Do If Data Failed to Be Synchronized Using hive-table?.....	1494
27.4.6 What Should I Do If Data Failed to Be Synchronized to a Hive Parquet Table Using HCatalog?.	1494
27.4.7 What Should I Do If the Data Type of Fields timestamp and data Is Incorrect During Data Synchronization Between Hive and MySQL?.....	1495
28 Using Storm.....	1496
28.1 Using Storm from Scratch.....	1496
28.2 Using the Storm Client.....	1497
28.3 Submitting Storm Topologies on the Client.....	1498
28.4 Accessing the Storm Web UI.....	1499
28.5 Managing Storm Topologies.....	1501
28.6 Querying Storm Topology Logs.....	1502
28.7 Storm Common Parameters.....	1502
28.8 Configuring a Storm Service User Password Policy.....	1504
28.9 Migrating Storm Services to Flink.....	1506
28.9.1 Overview.....	1506
28.9.2 Completely Migrating Storm Services.....	1506
28.9.3 Performing Embedded Service Migration.....	1508

28.9.4 Migrating Services of External Security Components Interconnected with Storm.....	1508
28.10 Storm Log Introduction.....	1509
28.11 Performance Tuning.....	1514
28.11.1 Storm Performance Tuning.....	1514
29 Using Tez.....	1517
29.1 Precautions.....	1517
29.2 Common Tez Parameters.....	1517
29.3 Accessing TezUI.....	1517
29.4 Log Overview.....	1518
29.5 Common Issues.....	1520
29.5.1 TezUI Cannot Display Tez Task Execution Details.....	1520
29.5.2 Error Occurs When a User Switches to the Tez Web UI.....	1520
29.5.3 Yarn Logs Cannot Be Viewed on the TezUI Page.....	1521
29.5.4 Table Data Is Empty on the TezUI HiveQueries Page.....	1522
30 Using YARN.....	1523
30.1 Common YARN Parameters.....	1523
30.2 Creating Yarn Roles.....	1527
30.3 Using the YARN Client.....	1529
30.4 Configuring Resources for a NodeManager Role Instance.....	1531
30.5 Changing NodeManager Storage Directories.....	1532
30.6 Configuring Strict Permission Control for Yarn.....	1536
30.7 Configuring Container Log Aggregation.....	1538
30.8 Using CGroups with YARN.....	1545
30.9 Configuring the Number of ApplicationMaster Retries.....	1547
30.10 Configure the ApplicationMaster to Automatically Adjust the Allocated Memory.....	1548
30.11 Configuring the Access Channel Protocol.....	1549
30.12 Configuring Memory Usage Detection.....	1550
30.13 Configuring the Additional Scheduler WebUI.....	1551
30.14 Configuring Yarn Restart.....	1552
30.15 Configuring ApplicationMaster Work Preserving.....	1554
30.16 Configuring the Localized Log Levels.....	1555
30.17 Configuring Users That Run Tasks.....	1556
30.18 Yarn Log Overview.....	1557
30.19 Yarn Performance Tuning.....	1560
30.19.1 Preempting a Task.....	1560
30.19.2 Setting the Task Priority.....	1563
30.19.3 Optimizing Node Configuration.....	1564
30.20 Common Issues About Yarn.....	1570
30.20.1 Why Mounted Directory for Container is Not Cleared After the Completion of the Job While Using CGroups?.....	1570
30.20.2 Why the Job Fails with HDFS_DELEGATION_TOKEN Expired Exception?.....	1571
30.20.3 Why Are Local Logs Not Deleted After YARN Is Restarted?.....	1571

30.20.4 Why the Task Does Not Fail Even Though AppAttempts Restarts for More Than Two Times?...	1572
30.20.5 Why Is an Application Moved Back to the Original Queue After ResourceManager Restarts?...	1572
30.20.6 Why Does Yarn Not Release the Blacklist Even All Nodes Are Added to the Blacklist?	1572
30.20.7 Why Does the Switchover of ResourceManager Occur Continuously?.....	1573
30.20.8 Why Does a New Application Fail If a NodeManager Has Been in Unhealthy Status for 10 Minutes?.....	1574
30.20.9 Why Does an Error Occur When I Query the ApplicationID of a Completed or Non-existing Application Using the RESTful APIs?.....	1574
30.20.10 Why May A Single NodeManager Fault Cause MapReduce Task Failures in the Superior Scheduling Mode?.....	1575
30.20.11 Why Are Applications Suspended After They Are Moved From Lost_and_Found Queue to Another Queue?.....	1575
30.20.12 How Do I Limit the Size of Application Diagnostic Messages Stored in the ZKstore?.....	1576
30.20.13 Why Does a MapReduce Job Fail to Run When a Non-ViewFS File System Is Configured as ViewFS?.....	1577
30.20.14 Why Do Reduce Tasks Fail to Run in Some OSs After the Native Task Feature is Enabled?.....	1578
31 Using ZooKeeper.....	1579
31.1 Using ZooKeeper from Scratch.....	1579
31.2 Common ZooKeeper Parameters.....	1583
31.3 Using a ZooKeeper Client.....	1584
31.4 Configuring the ZooKeeper Permissions.....	1584
31.5 ZooKeeper Log Overview.....	1589
31.6 Common Issues About ZooKeeper.....	1591
31.6.1 Why Do ZooKeeper Servers Fail to Start After Many znodes Are Created?.....	1591
31.6.2 Why Does the ZooKeeper Server Display the java.io.IOException: Len Error Log?.....	1593
31.6.3 Why Four Letter Commands Don't Work With Linux netcat Command When Secure Netty Configurations Are Enabled at Zookeeper Server?.....	1594
31.6.4 How Do I Check Which ZooKeeper Instance Is a Leader?.....	1595
31.6.5 Why Cannot the Client Connect to ZooKeeper using the IBM JDK?.....	1595
31.6.6 What Should I Do When the ZooKeeper Client Fails to Refresh a TGT?.....	1596
31.6.7 Why Is Message "Node does not exist" Displayed when A Large Number of Znodes Are Deleted Using the deleteall Command.....	1596
32 Appendix.....	1597
32.1 Modifying Cluster Service Configuration Parameters.....	1597
32.2 Accessing Manager.....	1598
32.2.1 Accessing MRS Manager (Versions Earlier Than MRS 3.x).....	1599
32.2.2 Accessing FusionInsight Manager (MRS 3.x or Later).....	1601
32.3 Using an MRS Client.....	1603
32.3.1 Installing a Client (MRS 3.x or Later).....	1603
32.3.2 Installing a Client (Versions Earlier Than 3.x).....	1608
32.3.3 Updating a Client (Version 3.x or Later).....	1613
32.3.4 Updating a Client (Versions Earlier Than 3.x).....	1615

1 Using Alluxio

1.1 Configuring an Underlying Storage System

If you want to use a unified client API and a global namespace to access persistent storage systems including HDFS and OBS to separate computing from storage, you can configure the underlying storage system of Alluxio on MRS Manager. After a cluster is created, the default underlying storage address is **hdfs://hacluster/**, that is, the HDFS root directory is mapped to Alluxio.

Prerequisites

- Alluxio has been installed in a cluster.
- The password of user **admin** has been obtained. The password of user **admin** is specified by the user during MRS cluster creation.

Configuring HDFS as the Underlying File System of Alluxio

NOTE


Security clusters with Kerberos authentication enabled do not support this function.

Step 1 Go to the **All Configurations** page of Alluxio. See [Modifying Cluster Service Configuration Parameters](#).

Step 2 In the left pane, choose **Alluxio > Under Stores**, and modify the value of **alluxio.master.mount.table.root.ufs** to **hdfs://hacluster/XXX/**.

For example, if you want to use *HDFS root directory/alluxio/* as the root directory of Alluxio, modify the value of **alluxio.master.mount.table.root.ufs** to **hdfs://hacluster/alluxio/**.

Figure 1-1 Configuring HDFS as the underlying file system of Alluxio



Step 3 Click **Save Configuration**. In the displayed dialog box, select **Restart the affected services or instances**.

Step 4 Click **OK** to restart Alluxio.

----End

1.2 Accessing Alluxio Using a Data Application

The port number used for accessing the Alluxio file system is 19998, and the access address is **alluxio://<Master node IP address of Alluxio>:19998/<PATH>**. This section uses examples to describe how to access the Alluxio file system using data applications (Spark, Hive, Hadoop MapReduce, and Presto).

Using Alluxio as the Input and Output of a Spark Application

Step 1 Log in to the Master node in a cluster as user **root** using the password set during cluster creation.

Step 2 Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

Step 3 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step:

```
kinit MRS cluster user
```

Example: **kinit admin**

Step 4 Prepare an input file and copy local data to the Alluxio file system.

For example, prepare the input file **test_input.txt** in the local **/home** directory, and run the following command to save the **test_input.txt** file to Alluxio:

```
alluxio fs copyFromLocal /home/test_input.txt /input
```

Step 5 Run the following commands to start **spark-shell**:

```
spark-shell
```

Step 6 Run the following commands in **spark-shell**:

```
val s = sc.textFile("alluxio://<Name of the Alluxio node>:19998/input")
```

```
val double = s.map(line => line + line)
```

```
double.saveAsTextFile("alluxio://<Name of the Alluxio node>:19998/output")
```

NOTE

Replace ***Name of the Alluxio node*:19998** with the actual node name and port numbers of all nodes where the AlluxioMaster instance is deployed. Use commas (,) to separate the node name and port number, for example, **node-ana-coremspb.mrs-m0va.com:19998,node-master2kiww.mrs-m0va.com:19998,node-master1cqww.mrs-m0va.com:19998**.

Step 7 Press **Ctrl+C** to exit **spark-shell**.

Step 8 Run the `alluxio fs ls /` command to check whether the output directory `/output` containing double content of the input file exists in the root directory of Alluxio.

----End

Creating a Hive Table on Alluxio

Step 1 Log in to the Master node in a cluster as user `root` using the password set during cluster creation.

Step 2 Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

Step 3 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step:

```
kinit MRS cluster user
```

Example: `kinit admin`

Step 4 Prepare an input file. For example, prepare the `hive_load.txt` input file in the local `/home` directory. The file content is as follows:

```
1, Alice, company A  
2, Bob, company B
```

Step 5 Run the following command to import the `hive_load.txt` file to Alluxio:

```
alluxio fs copyFromLocal /home/hive_load.txt /hive_input
```

Step 6 Run the following command to start the Hive beeline:

```
beeline
```

Step 7 Run the following commands in beeline to create a table based on the input file in Alluxio:

```
CREATE TABLE u_user(id INT, name STRING, company STRING) ROW FORMAT  
DELIMITED FIELDS TERMINATED BY ',' STORED AS TEXTFILE;
```

```
LOAD DATA INPATH 'alluxio://<Name of the Alluxio node>:19998/hive_input'  
INTO TABLE u_user;
```

NOTE

Replace `<Name of the Alluxio node>:19998` with the actual node name and port numbers of all nodes where the AlluxioMaster instance is deployed. Use commas (,) to separate the node name and port number, for example, `node-ana-coremspb.mrs-m0va.com:19998,node-master2kiww.mrs-m0va.com:19998,node-master1cqww.mrs-m0va.com:19998`.

Step 8 Run the following command to view the created table:

```
select * from u_user;
```

----End

Running Hadoop Wordcount in Alluxio

Step 1 Log in to the Master node in a cluster as user **root** using the password set during cluster creation.

Step 2 Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

Step 3 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step:

```
kinit MRS cluster user
```

Example: **kinit admin**

Step 4 Prepare an input file and copy local data to the Alluxio file system.

For example, prepare the input file **test_input.txt** in the local **/home** directory, and run the following command to save the **test_input.txt** file to Alluxio:

```
alluxio fs copyFromLocal /home/test_input.txt /input
```

Step 5 Run the following command to execute the wordcount job:

```
yarn jar /opt/share/hadoop-mapreduce-examples-<Hadoop version>-mrs-  
<MRS cluster version>/hadoop-mapreduce-examples-<Hadoop version>-mrs-  
<MRS cluster version>.jar wordcount alluxio://<Name of the Alluxio  
node>:19998/input alluxio://<Name of the Alluxio node>:19998/output
```

NOTE

- Replace **<Hadoop version>** with the actual one.
- Replace **<MRS cluster version>** with the major version of MRS. For example, for a cluster of MRS 1.9.2, **mrs-1.9.0** is used.
- Replace **Name of the Alluxio node>:19998** with the actual node name and port numbers of all nodes where the AlluxioMaster instance is deployed. Use commas (,) to separate the node name and port number, for example, **node-ana-coremspb.mrs-m0va.com:19998,node-master2kiww.mrs-m0va.com:19998,node-master1cqvw.mrs-m0va.com:19998**.

Step 6 Run the **alluxio fs ls /** command to check whether the output directory **/output** containing the wordcount result exists in the root directory of Alluxio.

----End

Using Presto to Query Tables in Alluxio

Step 1 Log in to the Master node in a cluster as user **root** using the password set during cluster creation.

Step 2 Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

Step 3 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step:

kinit *MRS cluster user*

Example: **kinit admin**

Step 4 Run the following commands to start Hive Beeline to create a table on Alluxio.

beeline

```
CREATE TABLE u_user (id int, name string, company string) ROW FORMAT  
DELIMITED FIELDS TERMINATED BY ',' LOCATION 'alluxio://<Name of the  
Alluxio node>:19998/u_user';
```

```
insert into u_user values(1,'Alice','Company A'),(2, 'Bob', 'Company B');
```

 **NOTE**

Replace *Name of the Alluxio node*:19998 with the actual node name and port numbers of all nodes where the AlluxioMaster instance is deployed. Use commas (,) to separate the node name and port number, for example, **node-ana-coremspb.mrs-m0va.com:19998,node-master2kiww.mrs-m0va.com:19998,node-master1cqww.mrs-m0va.com:19998**.

Step 5 Start the Presto client. For details, see [Step 2](#) to [Step 8](#) in [Using a Client to Execute Query Statements](#).

Step 6 On the Presto client, run the **select * from hive.default.u_user;** statement to query the table created in Alluxio:

Figure 1-2 Using Presto to query the table created in Alluxio

```
presto> select * from hive.default.u_user;  
id | name | company  
----+-----+-----  
 1 | Alice | Company A  
 2 | Bob   | Company B  
(2 rows)
```

----End

1.3 Common Operations of Alluxio

Preparations

1. Create a cluster with Alluxio installed.
2. Log in to the active Master node in a cluster as user **root** using the password set during cluster creation.
3. Run the following command to configure environment variables:
source /opt/client/bigdata_env

Using the Alluxio Shell

The **Alluxio shell** contains multiple command line operations that interact with Alluxio.

- View a file system operation command list:
alluxio fs

- Run the **ls** command to list the files in Alluxio. For example, list all files in the root directory:
alluxio fs ls /
- Run the **copyFromLocal** command to copy local files to Alluxio:
alluxio fs copyFromLocal /home/test_input.txt /test_input.txt
Command output:
Copied file:///home/test_input.txt to /test_input.txt
- Run the **ls** command again to list the files in Alluxio. The copied **test_input.txt** file is listed:
alluxio fs ls /
Command output:
12 PERSISTED 11-28-2019 17:10:17:449 100% /test_input.txt
The **test_input.txt** file is displayed in Alluxio. The parameters in the file indicate the file size, whether the file is persistent, creation date, cache ratio of the file in Alluxio, and file name.
- Run the **cat** command to print file content:
alluxio fs cat /test_input.txt
Command output:
Test Alluxio

Mounting Function of Alluxio

Alluxio uses a unified namespace feature to unify the access to storage systems. For details, see <https://docs.alluxio.io/os/user/2.0/en/advanced/Namespcae-Management.html>.

This feature allows users to mount different storage systems to an Alluxio namespace and seamlessly access files across storage systems through the Alluxio namespace.

1. Create a directory as a mount point in Alluxio.

alluxio fs mkdir /mnt

Successfully created directory /mnt

2. Mount an existing OBS file system to Alluxio. (Prerequisite: An agency with the **OBS OperateAccess** permission has been configured for the cluster.) The **obs-mrstest** file system is used as an example. Replace the file system name with the actual one.

alluxio fs mount /mnt/obs obs://obs-mrstest/data

Mounted obs://obs-mrstest/data at /mnt/obs

3. List files in the OBS file system using the Alluxio namespace. Run the **ls** command to list the files in the OBS mount directory.

alluxio fs ls /mnt/obs

38 PERSISTED 11-28-2019 17:42:54:554 0% /mnt/obs/hive_load.txt

12 PERSISTED 11-28-2019 17:43:07:743 0% /mnt/obs/test_input.txt

You can also view the newly mounted files and directories on the Alluxio web UI.

4. After the mounting is complete, you can seamlessly exchange data between different storage systems through the unified namespace of Alluxio. For example, run the **ls -R** command to list all files in a directory recursively:

alluxio fs ls -R /

```

0    PERSISTED 11-28-2019 11:15:19:719 DIR /app-logs
1    PERSISTED 11-28-2019 11:18:36:885 DIR /apps
1    PERSISTED 11-28-2019 11:18:40:209 DIR /apps/templeton
239440292  PERSISTED 11-28-2019 11:18:40:209 0% /apps/templeton/hive.tar.gz
.....
1    PERSISTED 11-28-2019 19:00:23:879 DIR /mnt
2    PERSISTED 11-28-2019 19:00:23:879 DIR /mnt/obs
38   PERSISTED 11-28-2019 17:42:54:554 0% /mnt/obs/hive_load.txt
12   PERSISTED 11-28-2019 17:43:07:743 0% /mnt/obs/test_input.txt
.....

```

The command output shows all files that are from the mounted storage system in the root directory of the Alluxio file system (the default directory is the HDFS root directory, that is, **hdfs://hacluster/**). The **/app-logs** and **/apps** directories are in HDFS, and the **/mnt/obs/** directory is in OBS.

Using Alluxio to Accelerate Data Access

Alluxio can accelerate data access, because it uses memory to store data. Example commands are provided as follows:

1. Upload the **test_data.csv** file (a sample that records recipes) to the **/data** directory of the **obs-mrtest** file system. Run the **ls** command to display the file status.

```
alluxio fs ls /mnt/obs/test_data.csv
```

```
294520189  PERSISTED 11-28-2019 19:38:55:000 0% /mnt/obs/test_data.csv
```

The output indicates that the cache percentage of the file in Alluxio is 0%, that is, the file is not in Alluxio memory.

2. Count the occurrence times of the word "milk" in the file, and calculate the time consumed.

```
time alluxio fs cat /mnt/obs/test_data.csv | grep -c milk
```

```
52180
```

```
real 0m10.765s
user 0m5.540s
sys 0m0.696s
```

3. Data is stored in memory after being read for the first time. When Alluxio reads data again, the data access speed is increased. For example, after running the **cat** command to obtain a file, run the **ls** command to check the file status.

```
alluxio fs ls /mnt/obs/test_data.csv
```

```
294520189  PERSISTED 11-28-2019 19:38:55:000 100% /mnt/obs/test_data.csv
```

The output shows that the file has been fully loaded to Alluxio.

4. Access the file again, count the occurrence times of the word "eggs", and calculate the time consumed.

```
time alluxio fs cat /mnt/obs/test_data.csv | grep -c eggs
```

```
59510
```

```
real 0m5.777s
user 0m5.992s
sys 0m0.592s
```

According to the comparison of the two time consumption records, the time consumed for accessing data stored in Alluxio memory is significantly reduced.

2 Using CarbonData (for Versions Earlier Than MRS 3.x)

2.1 Using CarbonData from Scratch

This section is for MRS 3.x or earlier. For MRS 3.x or later, see [Using CarbonData \(for MRS 3.x or Later\)](#).

This section describes the procedure of using Spark CarbonData. All tasks are based on the Spark-beeline environment. The tasks include:

1. Connecting to Spark
Before performing any operation on CarbonData, users must connect CarbonData to Spark.
2. Creating a CarbonData table
After connecting to Spark, users must create a CarbonData table to load and query data.
3. Loading data to the CarbonData table
Users load data from CSV files in HDFS to the CarbonData table.
4. Querying data from the CarbonData table
After data is loaded to the CarbonData table, users can run query commands such as **groupby** and **where**.

Prerequisites

A client has been installed. For details, see [Using an MRS Client](#).

Procedure

Step 1 Connect CarbonData to Spark.

1. Prepare a client based on service requirements and use user **root** to log in to the node where the client is installed.
For example, if you have updated the client on the Master2 node, log in to the Master2 node to use the client. For details, see [Using an MRS Client](#).

- Run the following commands to switch the user and configure environment variables:

```
sudo su - omm  
source /opt/client/bigdata_env
```

- For clusters with Kerberos authentication enabled, run the following command to authenticate the user. For clusters with Kerberos authentication disabled, skip this step.

```
kinit Spark username
```

 **NOTE**

The user needs to be added to user groups **hadoop** (primary group) and **hive**.

- Run the following command to connect to the Spark environment.

```
spark-beeline
```

Step 2 Create a CarbonData table.

Run the following command to create a CarbonData table, which is used to load and query data.

```
CREATE TABLE x1 (imei string, deviceInformationId int, mac string,  
productdate timestamp, updatetime timestamp, gamePointId double,  
contractNumber double)  
  
STORED BY 'org.apache.carbondata.format'  
  
TBLPROPERTIES  
('DICTIONARY_EXCLUDE'='mac','DICTIONARY_INCLUDE'='deviceInformationId'  
);
```

The command output is as follows:

```
+-----+  
| result |  
+-----+  
+-----+  
No rows selected (1.551 seconds)
```

Step 3 Load data from CSV files to the CarbonData table.

Run the command to load data from CSV files based on the required parameters. Only CSV files are supported. The CSV column name and sequence configured in the **LOAD** command must be consistent with those in the CarbonData table. The data formats and number of data columns in the CSV files must also be the same as those in the CarbonData table.

The CSV files must be stored on HDFS. You can upload the files to OBS and import them from OBS to HDFS on the **Files** page of the MRS console.

If Kerberos authentication is enabled, prepare the CSV files in the work environment and import them to HDFS using open-source HDFS commands. In addition, assign the Spark user with the read and execute permissions of the files on HDFS by referring to [5](#).

For example, the **data.csv** file is saved in the **tmp** directory of HDFS with the following contents:

```
x123,111,dd,2017-04-20 08:51:27,2017-04-20 07:56:51,2222,33333
```

The command for loading data from that file is as follows:

```
LOAD DATA inpath 'hdfs://hacluster/tmp/data.csv' into table x1
options('DELIMITER',';', 'QUOTECHAR','"', 'FILEHEADER='imei,
deviceinformationid,mac,productdate,updatetime,gamepointid,contractnumb
er');
```

The command output is as follows:

```
+-----+
| Result |
+-----+
+-----+
No rows selected (3.039 seconds)
```

Step 4 Query data from the CarbonData.

- **Obtaining the number of records**

Run the following command to obtain the number of records in the CarbonData table:

```
select count(*) from x1;
```

- **Querying with the groupby condition**

Run the following command to obtain the **deviceinformationid** records without repetition in the CarbonData table:

```
select deviceinformationid,count (distinct deviceinformationid) from x1
group by deviceinformationid;
```

- **Querying with the where condition**

Run the following command to obtain specific **deviceinformationid** records:

```
select * from x1 where deviceinformationid='111';
```

Step 5 Run the following command to exit the Spark environment.

```
!quit
```

```
----End
```

2.2 About CarbonData Table

Description

CarbonData tables are similar to tables in the relational database management system (RDBMS). RDBMS tables consist of rows and columns to store data. CarbonData tables have fixed columns and also store structured data. In CarbonData, data is saved in entity files.

Supported Data Types

CarbonData tables support the following data types:

- Int
- String
- BigInt
- Decimal

- Double
- TimeStamp

Table 2-1 describes the details about each data type.

Table 2-1 CarbonData data types

Data Type	Description
Int	4-byte signed integer ranging from -2,147,483,648 to 2,147,483,647 NOTE If a non-dictionary column is of the int data type, it is internally stored as the BigInt type.
String	The maximum character string length is 100000.
BigInt	Data is saved using the 64-bit technology. The value ranges from -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807.
Decimal	The default value is (10,0) and maximum value is (38,38). NOTE When query with filters, append BD to the number to achieve accurate results. For example, select * from carbon_table where num = 1234567890123456.22BD .
Double	Data is saved using the 64-bit technology. The value ranges from 4.9E-324 to 1.7976931348623157E308.
TimeStamp	yyyy-MM-dd HH:mm:ss format is used by default.

 **NOTE**

Measurement of all Integer data is processed and displayed using the **BigInt** data type.

2.3 Creating a CarbonData Table

Scenario

A CarbonData table must be created to load and query data.

Creating a Table with Self-Defined Columns

Users can create a table by specifying its columns and data types. For analysis clusters with Kerberos authentication enabled, if a user wants to create a CarbonData table in a database other than the **default** database, the **Create** permission of the database must be added to the role to which the user is bound in Hive role management.

Sample command:

```
CREATE TABLE IF NOT EXISTS productdb.productSalesTable (
productNumber Int,
```

```

productName String,
storeCity String,
storeProvince String,
revenue Int)
STORED BY 'org.apache.carbondata.format'
TBLPROPERTIES (
'table_blocksize'='128',
'DICTIONARY_EXCLUDE'='productName',
'DICTIONARY_INCLUDE'='productNumber');
    
```

The following table describes parameters of preceding commands.

Table 2-2 Parameter description

Parameter	Description
productSalesTable	Table name. The table is used to load data for analysis. The table name consists of letters, digits, and underscores (_).
productdb	Database name. The database maintains logical connections with tables stored in it to identify and manage the tables. The database name consists of letters, digits, and underscores (_).
productNumber productName storeCity storeProvince revenue	Columns in the table. The columns are service entities for data analysis. The column name (field name) consists of letters, digits, and underscores (_). NOTE In CarbonData, you cannot configure a column's NOT NULL or default value, or the primary key of the table.

Parameter	Description
table_blocksize	<p>Block size of data files used by the CarbonData table. The value ranges from 1 MB to 2048 MB. The default is 1024 MB.</p> <ul style="list-style-type: none"> If the value of table_blocksize is too small, a large number of small files will be generated when data is loaded. This may affect the performance in using HDFS. If the value of table_blocksize is too large, a large volume of data must be read from a block and the read concurrency is low when data is queried. As a result, the query performance deteriorates. <p>You are advised to set the block size based on the data volume. For example, set the block size to 256 MB for GB-level data, 512 MB for TB-level data, and 1024 MB for PB-level data.</p>
DICTIONARY_EXCLUDE	<p>Specifies the columns that do not generate dictionaries. This function is optional and applicable to columns of high complexity. By default, the system generates dictionaries for columns of the String type. However, as the number of values in the dictionaries increases, conversion operations by the dictionaries increase and the system performance deteriorates.</p> <p>Generally, if a column has over 50,000 unique data records, it is considered as a highly complex column and dictionary generation must be disabled.</p> <p>NOTE Non-dictionary columns support only the String and Timestamp data types.</p>
DICTIONARY_INCLUDE	<p>Specifies the columns that generate dictionaries. This function is optional and applicable to columns of low complexity. It improves the performance of queries with the groupby condition. Generally, the complexity of a dictionary column cannot exceed 50,000.</p>

2.4 Deleting a CarbonData Table

Scenario

Unused CarbonData tables can be deleted. After a CarbonData table is deleted, its metadata and loaded data are deleted together.

Procedure

Step 1 Run the following command to delete a CarbonData table:

```
DROP TABLE [IF EXISTS] [db_name.]table_name;
```

db_name is optional. If **db_name** is not specified, the table named **table_name** in the current database is deleted.

For example, run the following command to delete the **productSalesTable** table in the **productdb** database:

```
DROP TABLE productdb.productSalesTable;
```

Step 2 Run the following command to confirm that the table is deleted:

```
SHOW TABLES;
```

```
----End
```

3 Using CarbonData (for MRS 3.x or Later)

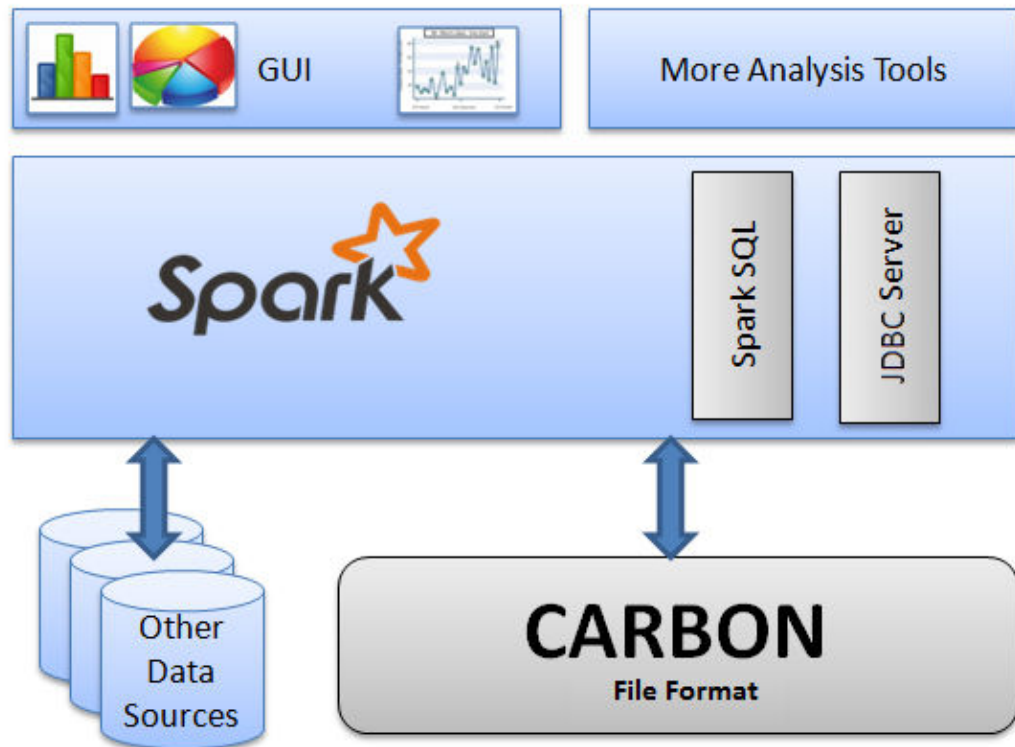
3.1 Overview

This section is for MRS 3.x or later. For MRS 3.x or earlier, see [Using CarbonData \(for Versions Earlier Than MRS 3.x\)](#).

3.1.1 CarbonData Overview

CarbonData is a new Apache Hadoop native data-store format. CarbonData allows faster interactive queries over PetaBytes of data using advanced columnar storage, index, compression, and encoding techniques to improve computing efficiency. In addition, CarbonData is also a high-performance analysis engine that integrates data sources with Spark.

Figure 3-1 Basic architecture of CarbonData



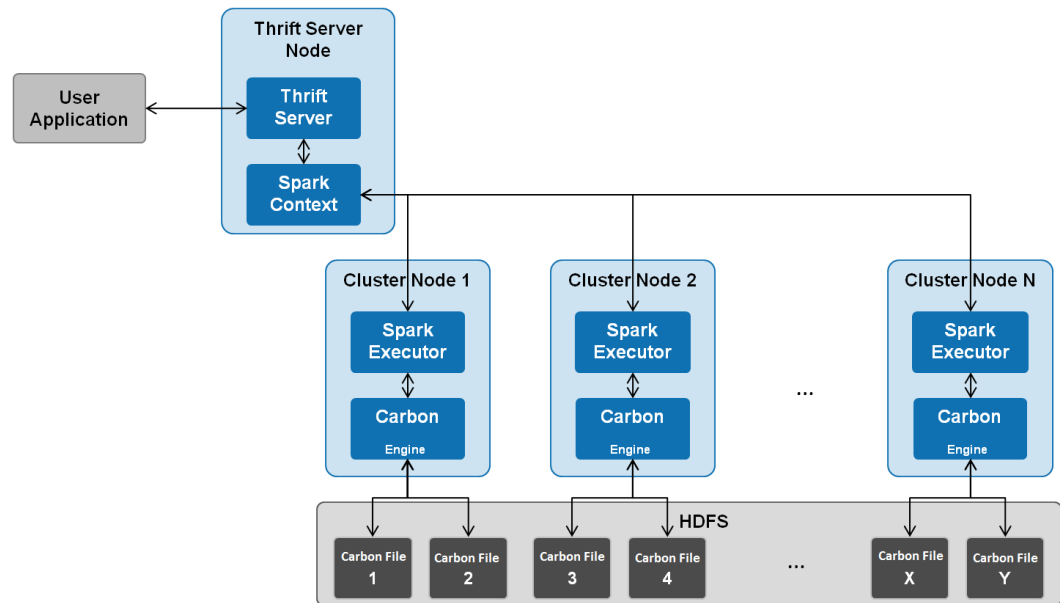
The purpose of using CarbonData is to provide quick response to ad hoc queries of big data. Essentially, CarbonData is an Online Analytical Processing (OLAP) engine, which stores data by using tables similar to those in Relational Database Management System (RDBMS). You can import more than 10 TB data to tables created in CarbonData format, and CarbonData automatically organizes and stores data using the compressed multi-dimensional indexes. After data is loaded to CarbonData, CarbonData responds to ad hoc queries in seconds.

CarbonData integrates data sources into the Spark ecosystem and you can query and analyze the data using Spark SQL. You can also use the third-party tool JDBCServer provided by Spark to connect to SparkSQL.

Topology of CarbonData

CarbonData runs as a data source inside Spark. Therefore, CarbonData does not start any additional processes on nodes in clusters. CarbonData engine runs inside the Spark executor.

Figure 3-2 Topology of CarbonData



Data stored in CarbonData Table is divided into several CarbonData data files. Each time when data is queried, CarbonData Engine reads and filters data sets. CarbonData Engine runs as a part of the Spark Executor process and is responsible for handling a subset of data file blocks.

Table data is stored in HDFS. Nodes in the same Spark cluster can be used as HDFS data nodes.

CarbonData Features

- SQL: CarbonData is compatible with Spark SQL and supports SQL query operations performed on Spark SQL.
- Simple Table dataset definition: CarbonData allows you to define and create datasets by using user-friendly Data Definition Language (DDL) statements. CarbonData DDL is flexible and easy to use, and can define complex tables.
- Easy data management: CarbonData provides various data management functions for data loading and maintenance. CarbonData supports bulk loading of historical data and incremental loading of new data. Loaded data can be deleted based on load time and a specific loading operation can be undone.
- CarbonData file format is a columnar store in HDFS. This format has many new column-based file storage features, such as table splitting and data compression. CarbonData has the following characteristics:
 - Stores data along with index: Significantly accelerates query performance and reduces the I/O scans and CPU resources, when there are filters in the query. CarbonData index consists of multiple levels of indices. A processing framework can leverage this index to reduce the task that needs to be scheduled and processed, and it can also perform skip scan in more finer grain unit (called blocklet) in task side scanning instead of scanning the whole file.
 - Operable encoded data: Through supporting efficient compression, CarbonData can query on compressed/encoded data. The data can be

converted just before returning the results to the users, which is called late materialized.

- Support for various use cases with one single data format: like interactive OLAP-style query, sequential access (big scan), and random access (narrow scan).

Key Technologies and Advantages of CarbonData

- Quick query response: CarbonData features high-performance query. The query speed of CarbonData is 10 times of that of Spark SQL. It uses dedicated data formats and applies multiple index technologies and multiple push-down optimizations, providing quick response to TB-level data queries.
- Efficient data compression: CarbonData compresses data by combining the lightweight and heavyweight compression algorithms. This significantly saves 60% to 80% data storage space and the hardware storage cost.

3.1.2 Main Specifications of CarbonData

Main Specifications of CarbonData

Table 3-1 Main Specifications of CarbonData

Entity	Tested Value	Test Environment
Number of tables	10000	3 nodes. 4 vCPUs and 20 GB memory for each executor. Driver memory: 5 GB, 3 executors. Total columns: 107 String: 75 Int: 13 BigInt: 7 Timestamp: 6 Double: 6
Number of table columns	2000	3 nodes. 4 vCPUs and 20 GB memory for each executor. Driver memory: 5 GB, 3 executors.
Maximum size of a raw CSV file	200 GB	17 cluster nodes. 150 GB memory and 25 vCPUs for each executor. Driver memory: 10 GB, 17 executors.

Entity	Tested Value	Test Environment
Number of CSV files in each folder	100 folders. Each folder has 10 files. The size of each file is 50 MB.	3 nodes. 4 vCPUs and 20 GB memory for each executor. Driver memory: 5 GB, 3 executors.
Number of load folders	10000	3 nodes. 4 vCPUs and 20 GB memory for each executor. Driver memory: 5 GB, 3 executors.

The memory required for data loading depends on the following factors:

- Number of columns
- Column values
- Concurrency (configured using **carbon.number.of.cores.while.loading**)
- Sort size in memory (configured using **carbon.sort.size**)
- Intermediate cache (configured using **carbon.graph.rowset.size**)

Data loading of an 8 GB CSV file that contains 10 million records and 300 columns with each row size being about 0.8 KB requires about 10 GB executor memory. That is, set **carbon.sort.size** to **100000** and retain the default values for other parameters.

Table Specifications

Table 3-2 Table specifications

Entity	Tested Value
Number of secondary index tables	10
Number of composite columns in a secondary index table	5
Length of column name in a secondary index table (unit: character)	120
Length of a secondary index table name (unit: character)	120
Cumulative length of all secondary index table names + column names in an index table* (unit: character)	3800**

 NOTE

- * Characters of column names in an index table refer to the upper limit allowed by Hive or the upper limit of available resources.
- ** Secondary index tables are registered using Hive and stored in HiveSERDEPROPERTIES in JSON format. The value of SERDEPROPERTIES supported by Hive can contain a maximum of 4,000 characters and cannot be changed.

3.2 Configuration Reference

This section provides the details of all the configurations required for the CarbonData System.

Table 3-3 System configurations in **carbon.properties**

Parameter	Default Value	Description
carbon.ddl.base.hdfs.url	hdfs://hacluster/opt/data	<p>HDFS relative path from the HDFS base path, which is configured in fs.defaultFS. The path configured in carbon.ddl.base.hdfs.url will be appended to the HDFS path configured in fs.defaultFS. If this path is configured, you do not need to pass the complete path while dataload.</p> <p>For example, if the absolute path of the CSV file is hdfs://10.18.101.155:54310/data/cnbc/2016/xyz.csv, the path hdfs://10.18.101.155:54310 will come from property fs.defaultFS and you can configure /data/cnbc/ as carbon.ddl.base.hdfs.url.</p> <p>During data loading, you can specify the CSV path as /2016/xyz.csv.</p>
carbon.badRecords.location	-	Storage path of bad records. This path is an HDFS path. The default value is Null . If bad records logging or bad records operation redirection is enabled, the path must be configured by the user.
carbon.badRecords.action	fail	<p>The following are four types of actions for bad records:</p> <p>FORCE: Data is automatically corrected by storing the bad records as NULL.</p> <p>REDIRECT: Bad records are written to the CSV file in carbon.badRecords.location instead of being loaded.</p> <p>IGNORE: Bad records are neither loaded nor written to the CSV file.</p> <p>FAIL: Data loading fails if any bad records are found.</p>

Parameter	Default Value	Description
carbon.update.sync.folder	/tmp/ carbodata	Specifies the modifiedTime.mdt file path. You can set it to an existing path or a new path. NOTE If you set this parameter to an existing path, ensure that all users can access the path and the path has the 777 permission.
carbon.enable.badrecord.action.redirect	false	Specifies whether to enable the REDIRECT mode to handle bad records during data loading. When it is enabled, bad records in source files will be recorded in a CSV file generated in a specified storage location each time data is loaded. CSV injection may occur when such CSV files are opened in Windows.

Table 3-4 Performance configurations in **carbon.properties**

Parameter	Default Value	Description
Data Loading Configuration		
carbon.sort.file.write.buffer.size	16384	CarbonData sorts data and writes it to a temporary file to limit memory usage. This parameter controls the size of the buffer used for reading and writing temporary files. The unit is bytes. The value ranges from 10240 to 10485760.
carbon.graph.rowset.size	100,000	Rowset size exchanged in data loading graph steps. The value ranges from 500 to 1,000,000.
carbon.number.of.cores.while.loading	6	Number of cores used during data loading. The greater the number of cores, the better the compaction performance. If the CPU resources are sufficient, you can increase the value of this parameter.
carbon.sort.size	500000	Number of records to be sorted
carbon.enableXXHash	true	Hashmap algorithm used for hashkey calculation
carbon.number.of.cores.block.sort	7	Number of cores used for sorting blocks during data loading

Parameter	Default Value	Description
carbon.max.driver.lru.cache.size	-1	Maximum size of LRU caching for data loading at the driver side. The unit is MB. The default value is -1 , indicating that there is no memory limit for the caching. Only integer values greater than 0 are accepted.
carbon.max.executor.lru.cache.size	-1	Maximum size of LRU caching for data loading at the executor side. The unit is MB. The default value is -1 , indicating that there is no memory limit for the caching. Only integer values greater than 0 are accepted. If this parameter is not configured, the value of carbon.max.driver.lru.cache.size is used.
carbon.merge.sort.prefetch	true	Whether to enable prefetch of data during merge sort while reading data from sorted temp files in the process of data loading
carbon.update.persist.enable	true	Configuration to enable the dataset of RDD/dataframe to persist data. Enabling this will reduce the execution time of UPDATE operation.
enable.unsafe.sort	true	Whether to use unsafe sort during data loading. Unsafe sort reduces the garbage collection during data load operation, resulting in better performance. The default value is true , indicating that unsafe sort is enabled.
enable.offheap.sort	true	Whether to use off-heap memory for sorting of data during data loading
offheap.sort.chunk.size.in.mb	64	Size of data chunks to be sorted, in MB. The value ranges from 1 to 1024.

Parameter	Default Value	Description
carbon.unsafe.working.memory.in.mb	512	<p>Size of the unsafe working memory. This will be used for sorting data and storing column pages. The unit is MB.</p> <p>Memory required for data loading: carbon.number.of.cores.while.loading [default value is 6] x Number of tables to load in parallel x offheap.sort.chunk.size.inmb [default value is 64 MB] + carbon.blockletgroup.size.in.mb [default value is 64 MB] + Current compaction ratio [64 MB/ 3.5] = Around 900 MB per table</p> <p>Memory required for data query: (SPARK_EXECUTOR_INSTANCES. [default value is 2] x (carbon.blockletgroup.size.in.mb [default value: 64 MB] + carbon.blockletgroup.size.in.mb [default value = 64 MB x 3.5) x Number of cores per executor [default value: 1]) = ~ 600 MB</p>
carbon.sort.inmemory.storage.size.in.mb	512	<p>Size of the intermediate sort data to be kept in the memory. Once the specified value is reached, the system writes data to the disk. The unit is MB.</p>
sort.inmemory.size.inmb	1024	<p>Size of the intermediate sort data to be kept in the memory. Once the specified value is reached, the system writes data to the disk. The unit is MB.</p> <p>If carbon.unsafe.working.memory.in.mb and carbon.sort.inmemory.storage.size.in.mb are configured, you do not need to set this parameter. If this parameter has been configured, 20% of the memory is used for working memory carbon.unsafe.working.memory.in.mb, and 80% is used for sort storage memory carbon.sort.inmemory.storage.size.in.mb.</p> <p>NOTE The value of spark.yarn.executor.memoryOverhead configured for Spark must be greater than the value of sort.inmemory.size.inmb configured for CarbonData. Otherwise, Yarn might stop the executor if off-heap access exceeds the configured executor memory.</p>

Parameter	Default Value	Description
carbon.blockletgroup.size.in.mb	64	The data is read as a group of blocklets which are called blocklet groups. This parameter specifies the size of each blocklet group. Higher value results in better sequential I/O access. The minimum value is 16 MB. Any value less than 16 MB will be reset to the default value (64 MB). The unit is MB.
enable.inmemory.merge.sort	false	Whether to enable inmemorymerge sort .
use.offheap.in.query.processing	true	Whether to enable offheap in query processing.
carbon.load.sort.scope	local_sort	Sort scope for the load operation. There are two types of sort: batch_sort and local_sort . If batch_sort is selected, the loading performance is improved but the query performance is reduced.
carbon.batch.sort.size.inmb	-	Size of data to be considered for batch sorting during data loading. The recommended value is less than 45% of the total sort data. The unit is MB. NOTE If this parameter is not set, its value is about 45% of the value of sort.inmemory.size.inmb by default.
enable.unsafe.columnpage	true	Whether to keep page data in heap memory during data loading or query to prevent garbage collection bottleneck.
carbon.use.local.dir	false	Whether to use Yarn local directories for multi-disk data loading. If this parameter is set to true , Yarn local directories are used to load multi-disk data to improve data loading performance.
carbon.use.multiple.temp.dir	false	Whether to use multiple temporary directories for storing temporary files to improve data loading performance.
carbon.load.datamaps.parallel.db_name.table_name	N/A	The value can be true or false . You can set the database name and table name to improve the first query performance of the table.
Compaction Configuration		

Parameter	Default Value	Description
carbon.number.of.cores.while.compacting	2	Number of cores to be used while compacting data. The greater the number of cores, the better the compaction performance. If the CPU resources are sufficient, you can increase the value of this parameter.
carbon.compaction.level.threshold	4,3	This configuration is for minor compaction which decides how many segments to be merged. For example, if this parameter is set to 2,3 , minor compaction is triggered every two segments. 3 is the number of level 1 compacted segments which is further compacted to new segment. The value ranges from 0 to 100.
carbon.major.compaction.size	1024	Major compaction size. Sum of the segments which is below this threshold will be merged. The unit is MB.
carbon.horizontal.compaction.enable	true	Whether to enable/disable horizontal compaction. After every DELETE and UPDATE statement, horizontal compaction may occur in case the incremental (DELETE/ UPDATE) files becomes more than specified threshold. By default, this parameter is set to true . You can set this parameter to false to disable horizontal compaction.
carbon.horizontal.update.compaction.threshold	1	Threshold limit on number of UPDATE delta files within a segment. In case the number of delta files goes beyond the threshold, the UPDATE delta files within the segment becomes eligible for horizontal compaction and are compacted into single UPDATE delta file. By default, this parameter is set to 1 . The value ranges from 1 to 10000 .
carbon.horizontal.delete.compaction.threshold	1	Threshold limit on number of DELETE incremental files within a block of a segment. In case the number of incremental files goes beyond the threshold, the DELETE incremental files for the particular block of the segment becomes eligible for horizontal compaction and are compacted into single DELETE incremental file. By default, this parameter is set to 1 . The value ranges from 1 to 10000 .
Query Configuration		
carbon.number.of.cores	4	Number of cores to be used during query

Parameter	Default Value	Description
carbon.limit.block.distribution.enable	false	Whether to enable the CarbonData distribution for limit query. The default value is false , indicating that block distribution is disabled for query statements that contain the keyword limit. For details about how to optimize this parameter, see Configurations for Performance Tuning .
carbon.custom.block.distribution	false	Whether to enable Spark or CarbonData block distribution. By default, the value is false , indicating that Spark block distribution is enabled. To enable CarbonData block distribution, change the value to true .
carbon.infilter.subquery.pushdown.enable	false	If this is set to true and a Select query is triggered in the filter with subquery, the subquery is executed and the output is broadcast as IN filter to the left table. Otherwise, SortMergeSemiJoin is executed. You are advised to set this to true when IN filter subquery does not return too many records. For example, when the IN sub-sentence query returns 10,000 or fewer records, enabling this parameter will give the query results faster. Example: <i>select * from flow_carbon_256b where cus_no in (select cus_no from flow_carbon_256b where dt>='20260101' and dt<='20260701' and txn_bk='tk_1' and txn_br='tr_1') limit 1000;</i>
carbon.scheduler.minRegisteredResourcesRatio	0.8	Minimum resource (executor) ratio needed for starting the block distribution. The default value is 0.8 , indicating that 80% of the requested resources are allocated for starting block distribution.
carbon.dynamicAllocation.schedulerTimeout	5	Maximum time that the scheduler waits for executors to be active. The default value is 5 seconds, and the maximum value is 15 seconds.
enable.unsafe.in.query.processing	true	Whether to use unsafe sort during query. Unsafe sort reduces the garbage collection during query, resulting in better performance. The default value is true , indicating that unsafe sort is enabled.
carbon.enable.vector.reader	true	Whether to enable vector processing for result collection to improve query performance
carbon.query.show.datamaps	true	SHOW TABLES lists all tables including the primary table and datamaps. To filter out the datamaps, set this parameter to false .
Secondary Index Configuration		

Parameter	Default Value	Description
carbon.secondary.index.creation.threads	1	Number of threads to concurrently process segments during secondary index creation. This property helps fine-tuning the system when there are a lot of segments in a table. The value ranges from 1 to 50.
carbon.si.lookup.partialstring	true	<ul style="list-style-type: none"> When the parameter value is true, it includes indexes started with, ended with, and contained. When the parameter value is false, it includes only secondary indexes started with.
carbon.si.segment.merge	true	<p>Enabling this property merges .carbondata files inside the secondary index segment. The merging will happen after the load operation. That is, at the end of the secondary index table load, small files are checked and merged.</p> <p>NOTE Table Block Size is used as the size threshold for merging small files.</p>

Table 3-5 Other configurations in **carbon.properties**

Parameter	Default Value	Description
Data Loading Configuration		
carbon.lock.type	HDFSLOCK	<p>Type of lock to be acquired during concurrent operations on a table.</p> <p>There are following types of lock implementation:</p> <ul style="list-style-type: none"> LOCALLOCK: Lock is created on local file system as a file. This lock is useful when only one Spark driver (or JDBCServer) runs on a machine. HDFSLOCK: Lock is created on HDFS file system as a file. This lock is useful when multiple Spark applications are running and no ZooKeeper is running on a cluster.
carbon.sort.intermediate.files.limit	20	Minimum number of intermediate files. After intermediate files are generated, sort and merge the files. For details about how to optimize this parameter, see Configurations for Performance Tuning .
carbon.csv.read.buffer.size.byte	1048576	Size of CSV reading buffer

Parameter	Default Value	Description
carbon.merge.sort.reader.thread	3	Maximum number of threads used for reading intermediate files for final merging.
carbon.concurrent.lock.retries	100	Maximum number of retries used to obtain the concurrent operation lock. This parameter is used for concurrent loading.
carbon.concurrent.lock.retry.timeout.sec	1	Interval between the retries to obtain the lock for concurrent operations.
carbon.lock.retries	3	Maximum number of retries to obtain the lock for any operations other than import.
carbon.lock.retry.timeout.sec	5	Interval between the retries to obtain the lock for any operation other than import.
carbon.tempstore.location	/opt/Carbon/TempStoreLoc	Temporary storage location. By default, the System.getProperty("java.io.tmpdir") method is used to obtain the value. For details about how to optimize this parameter, see the description of carbon.use.local.dir in Configurations for Performance Tuning .
carbon.load.log.counter	500000	Data loading records count in logs
SERIALIZATION_NULL_FORMAT	\N	Value to be replaced with NULL
carbon.skip.empty.line	false	Setting this property will ignore the empty lines in the CSV file during data loading.
carbon.load.datamaps.parallel	false	Whether to enable parallel datamap loading for all tables in all sessions. This property will improve the time to load datamaps into memory by distributing the job among executors, thus improving query performance.
Merging Configuration		
carbon.numberof.preserve.segments	0	If you want to preserve some number of segments from being compacted, then you can set this configuration. For example, if carbon.numberof.preserve.segments is set to 2 , the latest two segments will always be excluded from the compaction. No segments will be preserved by default.

Parameter	Default Value	Description
carbon.allowed.compaction.days	0	This configuration is used to control on the number of recent segments that needs to be merged. For example, if this parameter is set to 2 , the segments which are loaded in the time frame of past 2 days only will get merged. Segments which are loaded earlier than 2 days will not be merged. This configuration is disabled by default.
carbon.enable.auto.load.merge	false	Whether to enable compaction along with data loading.
carbon.merge.index.in.segment	true	This configuration enables to merge all the CarbonIndex files (.carbonindex) into a single MergeIndex file (.carbonindexmerge) upon data loading completion. This significantly reduces the delay in serving the first query.
Query Configuration		
max.query.execution.time	60	Maximum time allowed for one query to be executed. The unit is minute.
carbon.enableMinMax	true	MinMax is used to improve query performance. You can set this to false to disable this function.
carbon.lease.recovery.retry.count	5	Maximum number of attempts that need to be made for recovering a lease on a file. Minimum value: 1 Maximum value: 50
carbon.lease.recovery.retry.interval	1000 (ms)	Interval or pause time after a lease recovery attempt is made on a file. Minimum value: 1000 (ms) Maximum value: 10000 (ms)

Table 3-6 Spark configuration reference in **spark-defaults.conf**

Parameter	Default Value	Description
spark.driver.memory	4G	Memory to be used for the driver process. SparkContext has been initialized. NOTE In client mode, do not use SparkConf to set this parameter in the application because the driver JVM has been started. To configure this parameter, configure it in the --driver-memory command-line option or in the default property file.
spark.executor.memory	4 GB	Memory to be used for each executor process.
spark.sql.crossJoin.enabled	true	If the query contains a cross join, enable this property so that no error is thrown. In this case, you can use a cross join instead of a join for better performance.

Configure the following parameters in the **spark-defaults.conf** file on the Spark driver.

- In spark-sql mode:

Table 3-7 Parameter description

Parameter	Value	Description
spark.driver.extraJavaOptions	-Dlog4j.configuration=file:/opt/client/Spark2x/spark/conf/log4j.properties - Djetty.version=x.y.z - Dzookeeper.server.principal=zookeeper/hadoop.<System domain name> - Djava.security.krb5.conf=/opt/client/KrbClient/kerberos/var/krb5kdc/krb5.conf - Djava.security.auth.login.config=/opt/client/Spark2x/spark/conf/jaas.conf - Dorg.xerial.snappy.tmpdir=/opt/client/Spark2x/tmp - Dcarbon.properties.filepath=/opt/client/Spark2x/spark/conf/carbon.properties - Djava.io.tmpdir=/opt/client/Spark2x/tmp	The default value /opt/client/Spark2x/spark indicates CLIENT_HOME of the client and is added to the end of the value of spark.driver.extraJavaOptions . This parameter is used to specify the path of the carbon.properties file in Driver. NOTE Spaces next to equal marks (=) are not allowed.

Parameter	Value	Description
spark.sql.session.state.builder	org.apache.spark.sql.hive.FIHiveACLSessionStateBuilder	Session state constructor.
spark.carbon.sqlastbuilder.classname	org.apache.spark.sql.hive.CarbonInternalSqlAstBuilder	AST constructor.
spark.sql.catalog.class	org.apache.spark.sql.hive.HiveACLExternalCatalog	Hive External catalog to be used. This parameter is mandatory if Spark ACL is enabled.
spark.sql.hive.implementation	org.apache.spark.sql.hive.HiveACLClientImpl	How to call the Hive client. This parameter is mandatory if Spark ACL is enabled.
spark.sql.hiveClient.isolation.enabled	false	This parameter is mandatory if Spark ACL is enabled.

- In JDBCServer mode:

Table 3-8 Parameter description

Parameter	Value	Description
spark.driver.extraJavaOptions	-Xloggc:\${SPARK_LOG_DIR}/indexserver-omm-%p-gc.log - XX:+PrintGCDetails -XX:-OmitStackTracenFastThrow - XX:+PrintGCTimeStamps - XX:+PrintGCDateStamps - XX:MaxDirectMemorySize=512M - XX:MaxMetaspaceSize=512M - XX:+UseGCLogFileRotation - XX:NumberOfGCLogFiles=20 - XX:GCLogFileSize=10M - XX:OnOutOfMemoryError='kill -9 %p' - Djetty.version=x.y.z - Dorg.xerial.snappy.tmpdir=\${BIGDATA_HOME}/tmp/spark2x/JDBCServer/snappy_tmp - Djava.io.tmpdir=\${BIGDATA_HOME}/tmp/spark2x/JDBCServer/io_tmp - Dcarbon.properties.filepath=\${SPARK_CONF_DIR}/carbon.properties - Djdk.tls.ephemeralDHKeySize=20	The default value <code>\${SPARK_CONF_DIR}</code> depends on a specific cluster and is added to the end of the value of the <code>spark.driver.extraJavaOptions</code> parameter. This parameter is used to specify the path of the <code>carbon.properties</code> file in Driver. NOTE Spaces next to equal marks (=) are not allowed.

Parameter	Value	Description
	48 - Dspark.ssl.keyStore=\${SPARK_CONF_DIR}/child.keystore #{java_stack_prefer}	
spark.sql.session.state.builder	org.apache.spark.sql.hive.FIHiveACLSessionStateBuilder	Session state constructor.
spark.carbon.sqlastbuilder.classname	org.apache.spark.sql.hive.CarbonInternalSqlAstBuilder	AST constructor.
spark.sql.catalog.class	org.apache.spark.sql.hive.HiveACLExternalCatalog	Hive External catalog to be used. This parameter is mandatory if Spark ACL is enabled.
spark.sql.hive.implementation	org.apache.spark.sql.hive.HiveACLClientImpl	How to call the Hive client. This parameter is mandatory if Spark ACL is enabled.
spark.sql.hiveClient.isolation.enabled	false	This parameter is mandatory if Spark ACL is enabled.

3.3 CarbonData Operation Guide

3.3.1 CarbonData Quick Start

This section describes how to create CarbonData tables, load data, and query data. This quick start provides operations based on the Spark Beeline client. If you want to use Spark shell, wrap the queries with **spark.sql()**.

The following describes how to load data from a CSV file to a CarbonData table.

Table 3-9 CarbonData Quick Start

Operation	Description
Preparing a CSV File	Prepare the CSV file to be loaded to the CarbonData Table.

Operation	Description
Connecting to CarbonData	Connect to CarbonData before performing any operations on CarbonData.
Creating a CarbonData Table	Create a CarbonData table to load data and perform query operations.
Loading Data to a CarbonData Table	Load data from CSV to the created table.
Querying Data from a CarbonData Table	Perform query operations such as filters and groupby.

Preparing a CSV File

- Prepare a CSV file named **test.csv** on the local PC. An example is as follows:

```
13418592122,1001, MAC address, 2017-10-23 15:32:30,2017-10-24 15:32:30,62.50,74.56
13418592123 1002, MAC address, 2017-10-23 16:32:30,2017-10-24 16:32:30,17.80,76.28
13418592124,1003, MAC address, 2017-10-23 17:32:30,2017-10-24 17:32:30,20.40,92.94
13418592125 1004, MAC address, 2017-10-23 18:32:30,2017-10-24 18:32:30,73.84,8.58
13418592126,1005, MAC address, 2017-10-23 19:32:30,2017-10-24 19:32:30,80.50,88.02
13418592127 1006, MAC address, 2017-10-23 20:32:30,2017-10-24 20:32:30,65.77,71.24
13418592128,1007, MAC address, 2017-10-23 21:32:30,2017-10-24 21:32:30,75.21,76.04
13418592129,1008, MAC address, 2017-10-23 22:32:30,2017-10-24 22:32:30,63.30,94.40
13418592130, 1009, MAC address, 2017-10-23 23:32:30,2017-10-24 23:32:30,95.51,50.17
13418592131,1010, MAC address, 2017-10-24 00:32:30,2017-10-25 00:32:30,39.62,99.13
```
- Use WinSCP to import the CSV file to the directory of the node where the client is installed, for example, **/opt**.
- Log in to FusionInsight Manager and choose **System**. In the navigation pane on the left, choose **Permission > User**, click **Create** to create human-machine user **sparkuser**, and add the user to user groups hadoop (primary group) and hive.
- Run the following commands to go to the client installation directory, load environment variables, and authenticate the user.

```
cd /Client installation directory
source ./bigdata_env
source ./Spark2x/component_env
kinit sparkuser
```
- Run the following command to upload the CSV file to the **/data** directory of the HDFS.

```
hdfs dfs -put /opt/test.csv /data/
```

Connecting to CarbonData

- Use Spark SQL or Spark shell to connect to Spark and run Spark SQL commands.
- Run the following commands to start the JDBCServer and use a JDBC client (for example, Spark Beeline) to connect to the JDBCServer.

```
cd ./Spark2x/spark/bin
./spark-beeline
```

Creating a CarbonData Table

After connecting Spark Beeline with the JDBCServer, create a CarbonData table to load data and perform query operations. Run the following commands to create a simple table:

```
create table x1 (imei string, deviceInformationId int, mac string, productdate
timestamp, updatetime timestamp, gamePointId double, contractNumber
double) STORED AS carbondata TBLPROPERTIES
('SORT_COLUMNS'='imei,mac');
```

The command output is as follows:

```
+-----+
| Result |
+-----+
+-----+
No rows selected (1.093 seconds)
```

Loading Data to a CarbonData Table

After you have created a CarbonData table, you can load the data from CSV to the created table.

Run the following command with required parameters to load data from CSV. The column names of the CarbonData table must match the column names of the CSV file.

```
LOAD DATA inpath 'hdfs://hacluster/data/test.csv' into table x1
options('DELIMITER'=',', 'QUOTECHAR'='', 'FILEHEADER'='imei,
deviceinformationid,mac, productdate,updatetime,
gamepointid,contractnumber');
```

test.csv is the CSV file prepared in [Preparing a CSV File](#) and **x1** is the table name.

The CSV example file is as follows:

```
13418592122,1001, MAC address, 2017-10-23 15:32:30,2017-10-24 15:32:30,62.50,74.56
13418592123,1002, MAC address, 2017-10-23 16:32:30,2017-10-24 16:32:30,17.80,76.28
13418592124,1003, MAC address, 2017-10-23 17:32:30,2017-10-24 17:32:30,20.40,92.94
13418592125,1004, MAC address, 2017-10-23 18:32:30,2017-10-24 18:32:30,73.84,8.58
13418592126,1005, MAC address, 2017-10-23 19:32:30,2017-10-24 19:32:30,80.50,88.02
13418592127,1006, MAC address, 2017-10-23 20:32:30,2017-10-24 20:32:30,65.77,71.24
13418592128,1007, MAC address, 2017-10-23 21:32:30,2017-10-24 21:32:30,75.21,76.04
13418592129,1008, MAC address, 2017-10-23 22:32:30,2017-10-24 22:32:30,63.30,94.40
13418592130,1009, MAC address, 2017-10-23 23:32:30,2017-10-24 23:32:30,95.51,50.17
13418592131,1010, MAC address, 2017-10-24 00:32:30,2017-10-25 00:32:30,39.62,99.13
```

The command output is as follows:

```
+-----+
|Segment ID |
+-----+
|0          |
+-----+
No rows selected (3.039 seconds)
```

Querying Data from a CarbonData Table

After a CarbonData table is created and the data is loaded, you can perform query operations as required. Some query operations are provided as examples.

- **Obtaining the number of records**
Run the following command to obtain the number of records in the CarbonData table:
select count(*) from x1;
- **Querying with the groupby condition**
Run the following command to obtain the **deviceinformationid** records without repetition in the CarbonData table:
select deviceinformationid,count (distinct deviceinformationid) from x1 group by deviceinformationid;
- **Querying with Filter**
Run the following command to obtain specific **deviceinformationid** records:
select * from x1 where deviceinformationid='1010';

Using CarbonData on Spark-shell

If you need to use CarbonData on a Spark-shell, you need to create a CarbonData table, load data to the CarbonData table, and query data in CarbonData as follows:

```
spark.sql("CREATE TABLE x2(imei string, deviceInformationId int, mac string, productdate timestamp, updatetime timestamp, gamePointId double, contractNumber double) STORED AS carbondata")
spark.sql("LOAD DATA inpath 'hdfs://hacluster/data/x1_without_header.csv' into table x2
options('DELIMITER=', 'QUOTECHAR='\",'FILEHEADER'=imei, deviceinformationid,mac, productdate,updatetime, gamepointid,contractnumber)")
spark.sql("SELECT * FROM x2").show()
```

3.3.2 CarbonData Table Management

3.3.2.1 About CarbonData Table

Overview

In CarbonData, data is stored in entities called tables. CarbonData tables are similar to RDBMS tables. RDBMS data is stored in a table consisting of rows and columns. CarbonData tables store structured data, and have fixed columns and data types.

Supported Data Types

CarbonData tables support the following data types:

- Int
- String
- BigInt
- Smallint
- Char
- Varchar
- Boolean
- Decimal

- Double
- TimeStamp
- Date
- Array
- Struct
- Map

The following table describes supported data types and their respective values range.

Table 3-10 CarbonData data types

Data Type	Value Range
Int	4-byte signed integer ranging from -2,147,483,648 to 2,147,483,647. NOTE If a non-dictionary column is of the int data type, it is internally stored as the BigInt type.
String	100,000 characters NOTE If the CHAR or VARCHAR data type is used in CREATE TABLE , the two data types are automatically converted to the String data type. If a column contains more than 32,000 characters, add the column to the LONG_STRING_COLUMNS attribute of the tblproperties table during table creation.
BigInt	64-bit value ranging from -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
SmallInt	-32,768 to 32,767
Char	A to Z and a to z
Varchar	A to Z, a to z, and 0 to 9
Boolean	true or false
Decimal	The default value is (10,0) and maximum value is (38,38). NOTE When query with filters, append BD to the number to achieve accurate results. For example, select * from carbon_table where num = 1234567890123456.22BD .
Double	64-bit value ranging from 4.9E-324 to 1.7976931348623157E308
TimeStamp	The default format is yyyy-MM-dd HH:mm:ss .
Date	The DATE data type is used to store calendar dates. The default format is yyyy-MM-DD .

Data Type	Value Range
Array<data_type>	N/A NOTE Currently, only two layers of complex types can be nested.
Struct<col_name: data_type COMMENT col_comment, ...>	
Map<primitive_type, data_type>	

3.3.2.2 Creating a CarbonData Table

Scenario

A CarbonData table must be created to load and query data. You can run the **Create Table** command to create a table. This command is used to create a table using custom columns.

Creating a Table with Self-Defined Columns

Users can create a table by specifying its columns and data types.

Sample command:

```

CREATE TABLE IF NOT EXISTS productdb.productSalesTable (
    productNumber Int,
    productName String,
    storeCity String,
    storeProvince String,
    productCategory String,
    productBatch String,
    saleQuantity Int,
    revenue Int)
STORED AS carbondata
TBLPROPERTIES (
    'table_blocksize'='128');
    
```

The following table describes parameters of preceding commands.

Table 3-11 Parameter description

Parameter	Description
productSalesTable	Table name. The table is used to load data for analysis. The table name consists of letters, digits, and underscores (_).
productdb	Database name. The database maintains logical connections with tables stored in it to identify and manage the tables. The database name consists of letters, digits, and underscores (_).
productName storeCity storeProvince productCategory productBatch saleQuantity revenue	Columns in the table. The columns are service entities for data analysis. The column name (field name) consists of letters, digits, and underscores (_).
table_blocksize	Indicates the block size of data files used by the CarbonData table, in MB. The value ranges from 1 to 2048 . The default value is 1024 . If table_blocksize is too small, a large number of small files will be generated when data is loaded. This may affect the performance of HDFS. If table_blocksize is too large, during data query, the amount of block data that matches the index is large, and some blocks contain a large number of blocklets, affecting read concurrency and lowering query performance. You are advised to set the block size based on the data volume. For example, set the block size to 256 MB for GB-level data, 512 MB for TB-level data, and 1024 MB for PB-level data.

 **NOTE**

- Measurement of all Integer data is processed and displayed using the **BigInt** data type.
- CarbonData parses data strictly. Any data that cannot be parsed is saved as **null** in the table. For example, if the user loads the **double** value (3.14) to the **BigInt** column, the data is saved as **null**.
- The Short and Long data types used in the **Create Table** command are shown as **Smallint** and **BigInt** in the **DESCRIBE** command, respectively.
- You can run the **DESCRIBE** command to view the table data size and table index size.

Operation Result

Run the command to create a table.

3.3.2.3 Deleting a CarbonData Table

Scenario

You can run the **DROP TABLE** command to delete a table. After a CarbonData table is deleted, its metadata and loaded data are deleted together.

Procedure

Run the following command to delete a CarbonData table:

Run the following command:

```
DROP TABLE [IF EXISTS] [db_name.]table_name;
```

Once this command is executed, the table is deleted from the system. In the command, **db_name** is an optional parameter. If **db_name** is not specified, the table named **table_name** in the current database is deleted.

Example:

```
DROP TABLE productdb.productSalesTable;
```

Run the preceding command to delete the **productSalesTable** table from the **productdb** database.

Operation Result

Deletes the table specified in the command from the system. After the table is deleted, you can run the **SHOW TABLES** command to check whether the table is successfully deleted. For details, see [SHOW TABLES](#).

3.3.2.4 Modify the CarbonData Table

SET and UNSET

When the **SET** command is executed, the new properties overwrite the existing ones.

- SORT SCOPE

The following is an example of the **SET SORT SCOPE** command:

```
ALTER TABLE tablename SET TBLPROPERTIES('SORT_SCOPE'='no_sort')
```

After running the **UNSET SORT SCOPE** command, the default value **NO_SORT** is adopted.

The following is an example of the **UNSET SORT SCOPE** command:

```
ALTER TABLE tablename UNSET TBLPROPERTIES('SORT_SCOPE')
```

- SORT COLUMNS

The following is an example of the **SET SORT COLUMNS** command:

```
ALTER TABLE tablename SET TBLPROPERTIES('SORT_COLUMNS'='column1')
```


After this command is executed, the new value of **SORT_COLUMNS** is used. Users can adjust the **SORT_COLUMNS** based on the query results, but the original data is not affected. The operation does not affect the query performance of the original data segments which are not sorted by new **SORT_COLUMNS**.

The **UNSET** command is not supported, but the **SORT_COLUMNS** can be set to empty string instead of using the **UNSET** command.

```
ALTER TABLE tablename SET TBLPROPERTIES('SORT_COLUMNS'='')
```

NOTE

- The later version will enhance custom compaction to resort the old segments.
- The value of **SORT_COLUMNS** cannot be modified in the streaming table.
- If the **inverted index** column is removed from **SORT_COLUMNS**, **inverted index** will not be created in this column. However, the old configuration of **INVERTED_INDEX** will be kept.

3.3.3 CarbonData Table Data Management

3.3.3.1 Loading Data

Scenario

After a CarbonData table is created, you can run the **LOAD DATA** command to load data to the table for query. Once data loading is triggered, data is encoded in CarbonData format and files in multi-dimensional and column-based format are compressed and copied to the HDFS path of CarbonData files for quick analysis and queries. The HDFS path can be configured in the **carbon.properties** file. For details, see [Configuration Reference](#).

3.3.3.2 Deleting Segments

Scenario

If you want to modify and reload the data because you have loaded wrong data into a table, or there are too many bad records, you can delete specific segments by segment ID or data loading time.

NOTE

The segment deletion operation only deletes segments that are not compacted. You can run the **CLEAN FILES** command to clear compacted segments.

Deleting a Segment by Segment ID

Each segment has a unique ID. This segment ID can be used to delete the segment.

Step 1 Obtain the segment ID.

Command:

```
SHOW SEGMENTS FOR Table dbname.tablename LIMIT number_of_loads;
```

Example:

SHOW SEGMENTS FOR TABLE *carbonTable*;

Run the preceding command to show all the segments of the table named **carbonTable**.

SHOW SEGMENTS FOR TABLE *carbonTable LIMIT 2*;

Run the preceding command to show segments specified by *number_of_loads*.

The command output is as follows:

```
+-----+-----+-----+-----+-----+-----+-----+-----+
+
| ID | Status | Load Start Time | Load Time Taken | Partition | Data Size | Index Size | File Format |
+-----+-----+-----+-----+-----+-----+-----+-----+
+
| 3 | Success | 2020-09-28 22:53:26.336 | 3.726S | {} | 6.47KB | 3.30KB | columnar_v3 |
| 2 | Success | 2020-09-28 22:53:01.702 | 6.688S | {} | 6.47KB | 3.30KB | columnar_v3 |
+-----+-----+-----+-----+-----+-----+-----+-----+
+
```

 **NOTE**

The output of the **SHOW SEGMENTS** command includes ID, Status, Load Start Time, Load Time Taken, Partition, Data Size, Index Size, and File Format. The latest loading information is displayed in the first line of the command output.

Step 2 Run the following command to delete the segment after you have found the Segment ID:

Command:

DELETE FROM TABLE *tableName* WHERE SEGMENT.ID IN (*load_sequence_id1*, *load_sequence_id2*, ...);

Example:

DELETE FROM TABLE *carbonTable* WHERE SEGMENT.ID IN (1,2,3);

For details, see [DELETE SEGMENT by ID](#).

----End

Deleting a Segment by Data Loading Time

You can delete a segment based on the loading time.

Command:

DELETE FROM TABLE *db_name.table_name* WHERE SEGMENT.STARTTIME BEFORE *date_value*;

Example:

DELETE FROM TABLE *carbonTable* WHERE SEGMENT.STARTTIME BEFORE '2017-07-01 12:07:20';

The preceding command can be used to delete all segments before 2017-07-01 12:07:20.

For details, see [DELETE SEGMENT by DATE](#).

Result

Data of corresponding segments is deleted and is unavailable for query. You can run the **SHOW SEGMENTS** command to display the segment status and check whether the segment has been deleted.

NOTE

- Segments are not physically deleted after the execution of the **DELETE SEGMENT** command. Therefore, if you run the **SHOW SEGMENTS** command to check the status of a deleted segment, it will be marked as **Marked for Delete**. If you run the **SELECT * FROM tablename** command, the deleted segment will be excluded.
- The deleted segment will be deleted physically only when the next data loading reaches the maximum query execution duration, which is configured by the **max.query.execution.time** parameter. The default value of the parameter is 60 minutes.
- If you want to forcibly delete a physical segment file, run the **CLEAN FILES** command.

Example:

```
CLEAN FILES FOR TABLE table1;
```

This command will physically delete the segment file in the **Marked for delete** state.

If this command is executed before the time specified by **max.query.execution.time** arrives, the query may fail. **max.query.execution.time** indicates the maximum time allowed for a query, which is set in the **carbon.properties** file.

3.3.3.3 Combining Segments

Scenario

Frequent data access results in a large number of fragmented CarbonData files in the storage directory. In each data loading, data is sorted and indexing is performed. This means that an index is generated for each load. With the increase of data loading times, the number of indexes also increases. As each index works only on one loading, the performance of index is reduced. CarbonData provides loading and compression functions. In a compression process, data in each segment is combined and sorted, and multiple segments are combined into one large segment.

Prerequisites

Multiple data loadings have been performed.

Operation Description

There are three types of compaction: Minor, Major, and Custom.

- **Minor compaction:**
In minor compaction, you can specify the number of loads to be merged. If **carbon.enable.auto.load.merge** is set, minor compaction is triggered for every data load. If any segments are available to be merged, then compaction will run parallel with data load.

There are two levels in minor compaction:

- Level 1: Merging of the segments which are not yet compacted

- Level 2: Merging of the compacted segments again to form a larger segment
- Major compaction:
Multiple segments can be merged into one large segment. You can specify the compaction size so that all segments below the size will be merged. Major compaction is usually done during the off-peak time.
- Custom compaction:
In Custom compaction, you can specify the IDs of multiple segments to merge them into a large segment. The IDs of all the specified segments must exist and be valid. Otherwise, the compaction fails. Custom compaction is usually done during the off-peak time.

For details, see [ALTER TABLE COMPACTION](#).

Table 3-12 Compaction parameters

Parameter	Default Value	Application Type	Description
carbon.enable.automerge	false	Minor	Whether to enable compaction along with data loading. true: Compaction is automatically triggered when data is loaded. false: Compaction is not triggered when data is loaded.
carbon.compaction.level.threshold	4,3	Minor	This configuration is for minor compaction which decides how many segments to be merged. For example, if this parameter is set to 2,3 , minor compaction is triggered every two segments and segments form a single level 1 compacted segment. When the number of compacted level 1 segments reach 3, compaction is triggered again to merge them to form a single level 2 segment. The compaction policy depends on the actual data size and available resources. The value ranges from 0 to 100.

Parameter	Default Value	Application Type	Description
carbon.major.compaction.size	1024 MB	Major	<p>The major compaction size can be configured using this parameter. Sum of the segments which is below this threshold will be merged.</p> <p>For example, if this parameter is set to 1024 MB, and there are five segments whose sizes are 300 MB, 400 MB, 500 MB, 200 MB, and 100 MB used for major compaction, only segments whose total size is less than this threshold are compacted. In this example, only the segments whose sizes are 300 MB, 400 MB, 200 MB, and 100 MB are compacted.</p>
carbon.numberof.preserve.segments	0	Minor/Major	<p>If you want to preserve some number of segments from being compacted, then you can set this configuration.</p> <p>For example, if carbon.numberof.preserve.segments is set to 2, the latest two segments will always be excluded from the compaction.</p> <p>By default, no segments are reserved.</p>
carbon.allowed.compaction.days	0	Minor/Major	<p>This configuration is used to control on the number of recent segments that needs to be compacted.</p> <p>For example, if this parameter is set to 2, the segments which are loaded in the time frame of past 2 days only will get merged. Segments which are loaded earlier than 2 days will not be merged.</p> <p>This configuration is disabled by default.</p>
carbon.numberof.cores.while.compacting	2	Minor/Major	<p>Number of cores to be used while compacting data. The greater the number of cores, the better the compaction performance. If the CPU resources are sufficient, you can increase the value of this parameter.</p>

Parameter	Default Value	Application Type	Description
carbon.merge.index.in.segment	true	SEGMENT_INDEX	If this parameter is set to true , all the Carbon index (.carbonindex) files in a segment will be merged into a single Index (.carbonindexmerge) file. This enhances the first query performance.

Reference

You are advised not to perform minor compaction on historical data. For details, see [How to Avoid Minor Compaction for Historical Data?](#).

3.3.4 CarbonData Data Migration

Scenario

If you want to rapidly migrate CarbonData data from a cluster to another one, you can use the CarbonData backup and restoration commands. This method does not require data import in the target cluster, reducing required migration time.

Prerequisites

The Spark2x client has been installed in a directory, for example, **/opt/client**, in two clusters. The source cluster is cluster A, and the target cluster is cluster B.

Procedure

Step 1 Log in to the node where the client is installed in cluster A as a client installation user.

Step 2 Run the following commands to configure environment variables:

```
source /opt/client/bigdata_env
```

```
source /opt/client/Spark2x/component_env
```

Step 3 If the cluster is in security mode, run the following command to authenticate the user. In normal mode, skip user authentication.

```
kinit carbondatauser
```

carbondatauser indicates the user of the original data. That is, the user has the read and write permissions for the tables.

NOTE

You must add the user to the **hadoop** (primary group) and **hive** groups, and associate it with the **System_administrator** role.

Step 4 Run the following command to connect to the database and check the location for storing table data on HDFS:

```
spark-beeline
```

```
desc formatted Name of the table containing the original data;
```

Location in the displayed information indicates the directory where the data file resides.

Step 5 Log in to the node where the client is installed in cluster B as a client installation user and configure the environment variables:

```
source /opt/client/bigdata_env
```

```
source /opt/client/Spark2x/component_env
```

Step 6 If the cluster is in security mode, run the following command to authenticate the user. In normal mode, skip user authentication.

```
kinit carbondatauser2
```

carbondatauser2 indicates the user that uploads data.

 **NOTE**

You must add the user to the **hadoop** (primary group) and **hive** groups, and associate it with the **System_administrator** role.

Step 7 Run the **spark-beeline** command to connect to the database.

Step 8 Does the database that maps to the original data exist?

- If yes, go to [Step 9](#).
- If no, run the **create database** *Database name* command to create a database with the same name as that maps to the original data and go to [Step 9](#).

Step 9 Copy the original data from the HDFS directory in cluster A to that in cluster B.

When uploading data in cluster B, ensure that the upload directory has the directories with the same names as the database and table in the original directory and the upload user has the permission to write data to the upload directory. After the data is uploaded, the user has the permission to read and write the data.

For example, if the original data is stored in **/user/carboncadauser/warehouse/db1/tb1**, the data can be stored in **/user/carbondatauser2/warehouse/db1/tb1** in the new cluster.

1. Run the following command to download the original data to the **/opt/backup** directory of cluster A:

```
hdfs dfs -get /user/carboncadauser/warehouse/db1/tb1 /opt/backup
```
2. Run the following command to copy the original data of cluster A to the **/opt/backup** directory on the client node of cluster B.

```
scp /opt/backup root@IP address of the client node of cluster B:/opt/backup
```
3. Run the following command to upload the data copied to cluster B to HDFS:

```
hdfs dfs -put /opt/backup /user/carbondatauser2/warehouse/db1/tb1
```

Step 10 In the client environment of cluster B, run the following command to generate the metadata associated with the table corresponding to the original data in Hive:

```
REFRESH TABLE $dbName.$tbName;
```

\$dbName indicates the database name, and *\$tbName* indicates the table name.

Step 11 If the original table contains an index table, perform [Step 9](#) and [Step 10](#) to migrate the index table directory from cluster A to cluster B.

Step 12 Run the following command to register an index table for the CarbonData table (skip this step if no index table is created for the original table):

```
REGISTER INDEX TABLE $tableName ON $maintable;
```

\$tableName indicates the index table name, and *\$maintable* indicates the table name.

----End

3.3.5 Migrating Data on CarbonData from Spark 1.5 to Spark2x

Migration Solution Overview

This migration guides you to migrate the CarbonData table data of Spark 1.5 to that of Spark2x.

NOTE

Before performing this operation, you need to stop the data import service of the CarbonData table in Spark 1.5 and migrate data to the CarbonData table of Spark2x at a time. After the migration is complete, use Spark2x to perform service operations.

Migration roadmap:

1. Use Spark 1.5 to migrate historical data to the intermediate table.
2. Use Spark2x to migrate data from the intermediate table to the target table and change the target table name to the original table name.
3. After the migration is complete, use Spark2x to operate data in the CarbonData table.

Migration Solution and Commands

Migrating Historical Data

Step 1 Stop the CarbonData data import service, use spark-beeline of Spark 1.5 to view the ID and time of the latest segment in the CarbonData table, and record the segment ID.

```
show segments for table dbname.tablename;
```

Step 2 Run spark-beeline of Spark 1.5 as the user who has created the original CarbonData table to create an intermediate table in ORC or Parquet format. Then import the data in the original CarbonData table to the intermediate table. After the import is complete, the services of the CarbonData table can be restored.

Create an ORC table.

```
CREATE TABLE dbname.mid_tablename_orc STORED AS ORC as select * from  
dbname.tablename;
```

Create a Parquet table.

```
CREATE TABLE dbname.mid_tablename_parq STORED AS PARQUET as select *  
from dbname.tablename;
```

In the preceding command, **dbname** indicates the database name and **tablename** indicates the name of the original CarbonData table.

- Step 3** Run spark-beeline of Spark2x as the user who has created the original CarbonData table. Run the table creation statement of the old table to create a CarbonData table.

 **NOTE**

In the statement for creating a new table, the field sequence and type must be the same as those of the old table. In this way, the index column structure of the old table can be retained, which helps avoid errors caused by the use of **select *** statement during data insertion.

Run the spark-beeline command of Spark 1.5 to view the table creation statement of the old table: **SHOW CREATE TABLE dbname.tablename;**

Create a CarbonData table named **dbname.new_tablename**.

- Step 4** Run spark-beeline of Spark2x as the user who has created the original CarbonData table to load the intermediate table data in ORC (or PARQUET) format created in [Step 2](#) to the new table created in [Step 3](#). This step may take a long time (about 2 hours for 200 GB data). The following uses the ORC intermediate table as an example to describe the command for loading data:

```
insert into dbname.new_tablename select *  
from dbname.mid_tablename_orc;
```

- Step 5** Run spark-beeline of Spark2x as the user who has created the original CarbonData table to query and verify the data in the new table. If the data is correct, change the name of the original CarbonData table and then change the name of the new CarbonData table to the name of the original one.

```
ALTER TABLE dbname.tablename RENAME TO dbname.old_tablename;
```

```
ALTER TABLE dbname.new_tablename RENAME TO dbname.tablename;
```

- Step 6** Complete the migration. In this case, you can use Spark2x to query the new table and rebuild the secondary index.

----End

3.4 CarbonData Performance Tuning

3.4.1 Tuning Guidelines

Query Performance Tuning

There are various parameters that can be tuned to improve the query performance in CarbonData. Most of the parameters focus on increasing the parallelism in processing and optimizing system resource usage.

- Spark executor count: Executors are basic entities of parallelism in Spark. Raising the number of executors can increase the amount of parallelism in the cluster. For details about how to configure the number of executors, see the Spark documentation.
- Executor core: The number of concurrent tasks that an executor can run are controlled in each executor. Increasing the number of executor cores will add more concurrent processing tasks to improve performance.
- HDFS block size: CarbonData assigns query tasks by allocating different blocks to different executors for processing. HDFS block is the partition unit. CarbonData maintains a global block level index in Spark driver, which helps to reduce the quantity of blocks that need to be scanned for a query. Higher block size means higher I/O efficiency and lower global index efficiency. Reversely, lower block size means lower I/O efficiency, higher global index efficiency, and greater memory consumption.
- Number of scanner threads: Scanner threads control the number of parallel data blocks that are processed by each task. By increasing the number of scanner threads, you can increase the number of data blocks that are processed in parallel to improve performance. The **carbon.number.of.cores** parameter in the **carbon.properties** file is used to configure the number of scanner threads. For example, **carbon.number.of.cores = 4**.
- B-Tree caching: The cache memory can be optimized using the B-Tree least recently used (LRU) caching. In the driver, the B-Tree LRU caching configuration helps free up the cache by releasing table segments which are not accessed or not used. Similarly, in the executor, the B-Tree LRU caching configuration will help release table blocks that are not accessed or used. For details, see the description of **carbon.max.driver.lru.cache.size** and **carbon.max.executor.lru.cache.size** in [Table 3-4](#).

CarbonData Query Process

When CarbonData receives a table query task, for example query for table A, the index data of table A will be loaded to the memory for the query process. When CarbonData receives a query task for table A again, the system does not need to load the index data of table A.

When a query is performed in CarbonData, the query task is divided into several scan tasks, namely, task splitting based on HDFS blocks. Scan tasks are executed by executors on the cluster. Tasks can run in parallel, partially parallel, or in sequence, depending on the number of executors and configured number of executor cores.

Some parts of a query task can be processed at the individual task level, such as **select** and **filter**. Some parts of a query task can be processed at the individual task level, such as **group-by**, **count**, and **distinct count**.

Some operations cannot be performed at the task level, such as **Having Clause** (filter after grouping) and **sort**. Operations which cannot be performed at the task level or can be only performed partially at the task level require data (partial results) transmission across executors on the cluster. The transmission operation is called shuffle.

The more the tasks are, the more data needs to be shuffled. This affects query performance.

The number of tasks is depending on the number of HDFS blocks and the number of blocks is depending on the size of each block. You are advised to configure proper HDFS block size to achieve a balance among increased parallelism, the amount of data to be shuffled, and the size of aggregate tables.

Relationship Between Splits and Executors

If the number of splits is less than or equal to the executor count multiplied by the executor core count, the tasks are run in parallel. Otherwise, some tasks can start only after other tasks are complete. Therefore, ensure that the executor count multiplied by executor cores is greater than or equal to the number of splits. In addition, make sure that there are sufficient splits so that a query task can be divided into sufficient subtasks to ensure concurrency.

Configuring Scanner Threads

The scanner threads property decides the number of data blocks to be processed. If there are too many data blocks, a large number of small data blocks will be generated, affecting performance. If there are few data blocks, the parallelism is poor and the performance is affected. Therefore, when determining the number of scanner threads, you are advised to consider the average data size within a partition and select a value that makes the data block not small. Based on experience, you are advised to divide a single block size (unit: MB) by 250 and use the result as the number of scanner threads.

The number of actual available vCPUs is an important factor to consider when you want to increase the parallelism. The number of vCPUs that conduct parallel computation must not exceed 75% to 80% of actual vCPUs.

The number of vCPUs is approximately equal to:

Number of parallel tasks x Number of scanner threads. Number of parallel tasks is the smaller value of number of splits or executor count x executor cores.

Data Loading Performance Tuning

Tuning of data loading performance is different from that of query performance. Similar to query performance, data loading performance depends on the amount of parallelism that can be achieved. In case of data loading, the number of worker threads decides the unit of parallelism. Therefore, more executors mean more executor cores and better data loading performance.

To achieve better performance, you can configure the following parameters in HDFS.

Table 3-13 HDFS configuration

Parameter	Recommended Value
dfs.datanode.drop.cache.behind.reads	false
dfs.datanode.drop.cache.behind.writes	false
dfs.datanode.sync.behind.writes	true

Compression Tuning

CarbonData uses a few lightweight compression and heavyweight compression algorithms to compress data. Although these algorithms can process any type of data, the compression performance is better if the data is ordered with similar values being together.

During data loading, data is sorted based on the order of columns in the table to achieve good compression performance.

Since CarbonData sorts data in the order of columns defined in the table, the order of columns plays an important role in the effectiveness of compression. If the low cardinality dimension is on the left, the range of data partitions after sorting is small and the compression efficiency is high. If a high cardinality dimension is on the left, a range of data partitions obtained after sorting is relatively large, and compression efficiency is relatively low.

Memory Tuning

CarbonData provides a mechanism for memory tuning where data loading depends on the columns needed in the query. Whenever a query command is received, columns required by the query are fetched and data is loaded for those columns in memory. During this operation, if the memory threshold is reached, the least used loaded files are deleted to release memory space for columns required by the query.

3.4.2 Suggestions for Creating CarbonData Tables

Scenario

This section provides suggestions based on more than 50 test cases to help you create CarbonData tables with higher query performance.

Table 3-14 Columns in the CarbonData table

Column name	Data type	Cardinality	Attribution
msname	String	30 million	dimension
BEGIN_TIME	bigint	10,000	dimension
host	String	1 million	dimension

Column name	Data type	Cardinality	Attribution
dime_1	String	1,000	dimension
dime_2	String	500	dimension
dime_3	String	800	dimension
counter_1	numeric(20,0)	NA	measure
...	...	NA	measure
counter_100	numeric(20,0)	NA	measure

Procedure

- If the to-be-created table contains a column that is frequently used for filtering, for example, this column is used in more than 80% of filtering scenarios,
implement optimization as follows:
Place this column in the first column of **sort_columns**.
For example, if **msname** is used most frequently as a filter criterion in a query, it is placed in the first column. Run the following command to create a table. The query performance is good if **msname** is used as the filter criterion.

```
create table carbondata_table(
  msname String,
  ...
)STORED AS carbondata TBLPROPERTIES ('SORT_COLUMNS'='msname');
```

- If the to-be-created table has multiple columns which are frequently used to filter the results,
implement optimization as follows:
Create an index for the columns.
For example, if **msname**, **host**, and **dime_1** are frequently used columns, the **sort_columns** column sequence is "dime_1-> host-> msname..." based on cardinality. Run the following command to create a table. The following command can improve the filtering performance of **dime_1**, **host**, and **msname**.

```
create table carbondata_table(
  dime_1 String,
  host String,
  msname String,
  dime_2 String,
  dime_3 String,
  ...
)STORED AS carbondata
TBLPROPERTIES ('SORT_COLUMNS'='dime_1,host,msname');
```

- If the frequency of each column used for filtering is similar,
implement optimization as follows:
sort_columns is sorted in ascending order of cardinality.
Run the following command to create a table:

```
create table carbondata_table(
  Dime_1 String,
```

```
BEGIN_TIME bigint,
HOST String,
msname String,
...
)STORED AS carbondata
TBLPROPERTIES ('SORT_COLUMNS'='dime_2,dime_3,dime_1, BEGIN_TIME,host,msname');
```

- Create tables in ascending order of cardinalities. Then create secondary indexes for columns with more cardinalities. The statement for creating an index is as follows:

```
create index carbondata_table_index_msidx on tablecarbondata_table (
msname String) as 'carbondata' PROPERTIES ('table_blocksize'='128');
create index carbondata_table_index_host on tablecarbondata_table (
host String) as 'carbondata' PROPERTIES ('table_blocksize'='128');
```

- For columns of measure type, not requiring high accuracy, the numeric (20,0) data type is not required. You are advised to use the double data type to replace the numeric (20,0) data type to enhance query performance.

The result of performance analysis of test-case shows reduction in query execution time from 15 to 3 seconds, thereby improving performance by nearly 5 times. The command for creating a table is as follows:

```
create table carbondata_table(
Dime_1 String,
BEGIN_TIME bigint,
HOST String,
msname String,
counter_1 double,
counter_2 double,
...
counter_100 double,
)STORED AS carbondata
;
```

- If values (**start_time** for example) of a column are incremental: For example, if data is loaded to CarbonData every day, **start_time** is incremental for each load. In this case, it is recommended that the **start_time** column be put at the end of **sort_columns**, because incremental values are efficient in using min/max index. The command for creating a table is as follows:

```
create table carbondata_table(
Dime_1 String,
HOST String,
msname String,
counter_1 double,
counter_2 double,
BEGIN_TIME bigint,
...
counter_100 double,
)STORED AS carbondata
TBLPROPERTIES ( 'SORT_COLUMNS'='dime_2,dime_3,dime_1..BEGIN_TIME');
```

3.4.3 Configurations for Performance Tuning

Scenario

This section describes the configurations that can improve CarbonData performance.

Procedure

[Table 3-15](#) and [Table 3-16](#) describe the configurations about query of CarbonData.

Table 3-15 Number of tasks started for the shuffle process

Parameter	spark.sql.shuffle.partitions
Configuration File	spark-defaults.conf
Function	Data query
Scenario Description	Number of tasks started for the shuffle process in Spark
Tuning	You are advised to set this parameter to one to two times as much as the executor cores. In an aggregation scenario, reducing the number from 200 to 32 can reduce the query time by two folds.

Table 3-16 Number of executors and vCPUs, and memory size used for CarbonData data query

Parameter	spark.executor.cores spark.executor.instances spark.executor.memory
Configuration File	spark-defaults.conf
Function	Data query
Scenario Description	Number of executors and vCPUs, and memory size used for CarbonData data query
Tuning	In the bank scenario, configuring 4 vCPUs and 15 GB memory for each executor will achieve good performance. The two values do not mean the more the better. Configure the two values properly in case of limited resources. If each node has 32 vCPUs and 64 GB memory in the bank scenario, the memory is not sufficient. If each executor has 4 vCPUs and 12 GB memory, Garbage Collection may occur during query, time spent on query from increases from 3s to more than 15s. In this case, you need to increase the memory or reduce the number of vCPUs.

[Table 3-17](#), [Table 3-18](#), and [Table 3-19](#) describe the configurations for CarbonData data loading.

Table 3-17 Number of vCPUs used for data loading

Parameter	carbon.number.of.cores.while.loading
------------------	--------------------------------------

Configuration File	carbon.properties
Function	Data loading
Scenario Description	Number of vCPUs used for data processing during data loading in CarbonData
Tuning	If there are sufficient CPUs, you can increase the number of vCPUs to improve performance. For example, if the value of this parameter is changed from 2 to 4, the CSV reading performance can be doubled.

Table 3-18 Whether to use Yarn local directories for multi-disk data loading

Parameter	carbon.use.local.dir
Configuration File	carbon.properties
Function	Data loading
Scenario Description	Whether to use Yarn local directories for multi-disk data loading
Tuning	If this parameter is set to true , CarbonData uses local Yarn directories for multi-table load disk load balance, improving data loading performance.

Table 3-19 Whether to use multiple directories during loading

Parameter	carbon.use.multiple.temp.dir
Configuration File	carbon.properties
Function	Data loading
Scenario Description	Whether to use multiple temporary directories to store temporary sort files
Tuning	If this parameter is set to true , multiple temporary directories are used to store temporary sort files during data loading. This configuration improves data loading performance and prevents single points of failure (SPOFs) on disks.

[Table 3-20](#) describes the configurations for CarbonData data loading and query.

Table 3-20 Number of vCPUs used for data loading and query

Parameter	carbon.compaction.level.threshold
Configuration File	carbon.properties
Function	Data loading and query
Scenario Description	For minor compaction, specifies the number of segments to be merged in stage 1 and number of compacted segments to be merged in stage 2.
Tuning	Each CarbonData load will create one segment, if every load is small in size, it will generate many small files over a period of time impacting the query performance. Configuring this parameter will merge the small segments to one big segment which will sort the data and improve the performance. The compaction policy depends on the actual data size and available resources. For example, a bank loads data once a day and at night when no query is performed. If resources are sufficient, the compaction policy can be 6 or 5.

Table 3-21 Whether to enable data pre-loading when the index cache server is used

Parameter	carbon.indexserver.enable.prepriming
Configuration File	carbon.properties
Function	Data loading
Scenario Description	Enabling data pre-loading during the use of the index cache server can improve the performance of the first query.
Tuning	You can set this parameter to true to enable the pre-loading function. The default value is false .

3.5 CarbonData Access Control

The following table provides details about Hive ACL permissions required for performing operations on CarbonData tables.

Prerequisites

Parameters listed in [Table 3-7](#) or [Table 3-8](#) have been configured.

Hive ACL permissions

Table 3-22 Hive ACL permissions required for CarbonData table-level operations

Scenario	Required Permission
DESCRIBE TABLE	SELECT (of table)
SELECT	SELECT (of table)
EXPLAIN	SELECT (of table)
CREATE TABLE	CREATE (of database)
CREATE TABLE As SELECT	CREATE (on database), INSERT (on table), RW on data file, and SELECT (on table)
LOAD	INSERT (of table) RW on data file
DROP TABLE	OWNER (of table)
DELETE SEGMENTS	DELETE (of table)
SHOW SEGMENTS	SELECT (of table)
CLEAN FILES	DELETE (of table)
INSERT OVERWRITE / INSERT INTO	INSERT (of table) RW on data file and SELECT (of table)
CREATE INDEX	OWNER (of table)
DROP INDEX	OWNER (of table)
SHOW INDEXES	SELECT (of table)
ALTER TABLE ADD COLUMN	OWNER (of table)
ALTER TABLE DROP COLUMN	OWNER (of table)
ALTER TABLE CHANGE DATATYPE	OWNER (of table)
ALTER TABLE RENAME	OWNER (of table)
ALTER TABLE COMPACTION	INSERT (on table)
FINISH STREAMING	OWNER (of table)
ALTER TABLE SET STREAMING PROPERTIES	OWNER (of table)
ALTER TABLE SET TABLE PROPERTIES	OWNER (of table)
UPDATE CARBON TABLE	UPDATE (of table)
DELETE RECORDS	DELETE (of table)
REFRESH TABLE	OWNER (of main table)

Scenario	Required Permission
REGISTER INDEX TABLE	OWNER (of table)
SHOW PARTITIONS	SELECT (on table)
ALTER TABLE ADD PARTITION	OWNER (of table)
ALTER TABLE DROP PARTITION	OWNER (of table)

 **NOTE**

- If tables in the database are created by multiple users, the **Drop database** command fails to be executed even if the user who runs the command is the owner of the database.
- In a secondary index, when the parent table is triggered, **insert** and **compaction** are triggered on the index table. If you select a query that has a filter condition that matches index table columns, you should provide selection permissions for the parent table and index table.
- The LockFiles folder and lock files created in the LockFiles folder will have full permissions, as the LockFiles folder does not contain any sensitive data.
- If you are using ACL, ensure you do not configure any path for DDL or DML which is being used by other process. You are advised to create new paths.
Configure the path for the following configuration items:
 - 1) carbon.badRecords.location
 - 2) Db_Path and other items during database creation
- For Carbon ACL in a non-security cluster, **hive.server2.enable.doAs** in the **hive-site.xml** file must be set to **false**. Then the query will run as the user who runs the hiveserver2 process.

3.6 CarbonData Syntax Reference

3.6.1 DDL

3.6.1.1 CREATE TABLE

Function

This command is used to create a CarbonData table by specifying the list of fields along with the table properties.

Syntax

```
CREATE TABLE [IF NOT EXISTS] [db_name.]table_name
[(col_name data_type, ...)]
STORED AS carbodata
[TBLPROPERTIES (property_name=property_value, ...)];
```

Additional attributes of all tables are defined in **TBLPROPERTIES**.

Parameter Description

Table 3-23 CREATE TABLE parameters

Parameter	Description
db_name	Database name that contains letters, digits, and underscores (_).
col_name data_type	List with data types separated by commas (,). The column name contains letters, digits, and underscores (_). NOTE When creating a CarbonData table, do not use tupleId, PositionId, and PositionReference as column names because columns with these names are internally used by secondary index commands.
table_name	Table name of a database that contains letters, digits, and underscores (_).
STORED AS	The carbondata parameter defines and creates a CarbonData table.
TBLPROPERTIES	List of CarbonData table properties.

Precautions

Table attributes are used as follows:

- Block size

The block size of a data file can be defined for a single table using **TBLPROPERTIES**. The larger one between the actual size of the data file and the defined block size is selected as the actual block size of the data file in HDFS. The unit is MB. The default value is 1024 MB. The value ranges from 1 MB to 2048 MB. If the value is beyond the range, the system reports an error.

Once the block size reaches the configured value, the write program starts a new block of CarbonData data. Data is written in multiples of the page size (32,000 records). Therefore, the boundary is not strict at the byte level. If the new page crosses the boundary of the configured block, the page is written to the new block instead of the current block.

```
TBLPROPERTIES('table_blocksize='128')
```

 NOTE

- If a small block size is configured in the CarbonData table while the size of the data file generated by the loaded data is large, the block size displayed in HDFS is different from the configured value. This is because when data is written to a local block file for the first time, even though the size of the to-be-written data is larger than the configured value of the block size, data will still be written into the block. Therefore, the actual value of block size in HDFS is the larger value between the size of the data to be written and the configured block size.
- If **block.num** is less than the parallelism, the blocks are split into new blocks so that new blocks.num is greater than parallelism and all cores can be used. This optimization is called block distribution.
- **SORT_SCOPE** specifies the sort scope during table creation. There are four types of sort scopes:
 - **GLOBAL_SORT**: It improves query performance, especially for point queries. *TBLPROPERTIES('SORT_SCOPE'='GLOBAL_SORT')*
 - **LOCAL_SORT**: Data is sorted locally (task-level sorting).
 - **NO_SORT**: The default sorting mode is used. Data is loaded in unsorted manner, which greatly improves loading performance.

- **SORT_COLUMNS**

This table property specifies the order of sort columns.

TBLPROPERTIES('SORT_COLUMNS'='column1, column3')

 NOTE

- If this attribute is not specified, no columns are sorted by default.
- If this property is specified but with empty argument, then the table will be loaded without sort. For example, *('SORT_COLUMNS='')*.
- **SORT_COLUMNS** supports the string, date, timestamp, short, int, long, byte, and boolean data types.
- **RANGE_COLUMN**

This property is used to specify a column to partition the input data by range. Only one column can be configured. During data import, you can use **global_sort_partitions** or **scale_factor** to avoid generating small files.

TBLPROPERTIES('RANGE_COLUMN'='column1')

- **LONG_STRING_COLUMNS**

The length of a common string cannot exceed 32,000 characters. To store a string of more than 32,000 characters, set **LONG_STRING_COLUMNS** to the target column.

TBLPROPERTIES('LONG_STRING_COLUMNS'='column1, column3')

 NOTE

LONG_STRING_COLUMNS can be set only for columns of the STRING, CHAR, or VARCHAR type.

Scenarios

Creating a Table by Specifying Columns

The **CREATE TABLE** command is the same as that of Hive DDL. The additional configurations of CarbonData are provided as table properties.

```
CREATE TABLE [IF NOT EXISTS] [db_name.]table_name  
[(col_name data_type , ...)]  
STORED AS carbodata  
[TBLPROPERTIES (property_name=property_value, ...)];
```

Examples

```
CREATE TABLE IF NOT EXISTS productdb.productSalesTable (  
productNumber Int,  
productName String,  
storeCity String,  
storeProvince String,  
productCategory String,  
productBatch String,  
saleQuantity Int,  
revenue Int)  
STORED AS carbodata  
TBLPROPERTIES (  
'table_blocksize'='128',  
'SORT_COLUMNS'='productBatch, productName')
```

System Response

A table will be created and the success message will be logged in system logs.

3.6.1.2 CREATE TABLE As SELECT

Function

This command is used to create a CarbonData table by specifying the list of fields along with the table properties.

Syntax

```
CREATE TABLE [IF NOT EXISTS] [db_name.]table_name STORED AS carbodata  
[TBLPROPERTIES (key1=val1, key2=val2, ...)] AS select_statement;
```

Parameter Description

Table 3-24 CREATE TABLE parameters

Parameter	Description
db_name	Database name that contains letters, digits, and underscores (_).
table_name	Table name of a database that contains letters, digits, and underscores (_).
STORED AS	Used to store data in CarbonData format.
TBLPROPERTIES	List of CarbonData table properties. For details, see Precautions .

Precautions

N/A

Examples

```
CREATE TABLE ctas_select_parquet STORED AS carbondata as select * from
parquet_ctas_test;
```

System Response

This example will create a Carbon table from any Parquet table and load all the records from the Parquet table.

3.6.1.3 DROP TABLE

Function

This command is used to delete an existing table.

Syntax

```
DROP TABLE [IF EXISTS] [db_name.]table_name;
```

Parameter Description

Table 3-25 DROP TABLE parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is selected.
table_name	Name of the table to be deleted

Precautions

In this command, **IF EXISTS** and **db_name** are optional.

Example

```
DROP TABLE IF EXISTS productDatabase.productSalesTable;
```

System Response

The table will be deleted.

3.6.1.4 SHOW TABLES

Function

SHOW TABLES command is used to list all tables in the current or a specific database.

Syntax

```
SHOW TABLES [IN db_name];
```

Parameter Description

Table 3-26 SHOW TABLE parameters

Parameter	Description
IN db_name	Name of the database. This parameter is required only when tables of this specific database are to be listed.

Usage Guidelines

IN db_Name is optional.

Examples

```
SHOW TABLES IN ProductDatabase;
```

System Response

All tables are listed.

3.6.1.5 ALTER TABLE COMPACTION

Function

The **ALTER TABLE COMPACTION** command is used to merge a specified number of segments into a single segment. This improves the query performance of a table.

Syntax

```
ALTER TABLE [db_name.]table_name COMPACT 'MINOR/MAJOR/  
SEGMENT_INDEX';
```

```
ALTER TABLE [db_name.]table_name COMPACT 'CUSTOM' WHERE SEGMENT.ID IN  
(id1, id2, ...);
```

Parameter Description

Table 3-27 ALTER TABLE COMPACTION parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is selected.
table_name	Table name.
MINOR	Minor compaction. For details, see Combining Segments .
MAJOR	Major compaction. For details, see Combining Segments .
SEGMENT_INDEX	This configuration enables you to merge all the CarbonData index files (.carbonindex) inside a segment to a single CarbonData index merge file (.carbonindexmerge). This enhances the first query performance. For more information, see Table 3-12 .
CUSTOM	Custom compaction. For details, see Combining Segments .

Precautions

N/A

Examples

```
ALTER TABLE ProductDatabase COMPACT 'MINOR';
```

```
ALTER TABLE ProductDatabase COMPACT 'MAJOR';
```

```
ALTER TABLE ProductDatabase COMPACT 'SEGMENT_INDEX';
```

```
ALTER TABLE ProductDatabase COMPACT 'CUSTOM' WHERE SEGMENT.ID IN  
(0, 1);
```

System Response

ALTER TABLE COMPACTION does not show the response of the compaction because it is run in the background.

If you want to view the response of minor and major compactions, you can check the logs or run the **SHOW SEGMENTS** command.

Example:

```

+-----+-----+-----+-----+-----+-----+-----+
+--+
| ID | Status | Load Start Time | Load Time Taken | Partition | Data Size | Index Size | File
Format |
+-----+-----+-----+-----+-----+-----+-----+
+--+
| 3 | Success | 2020-09-28 22:53:26.336 | 3.726S | {} | 6.47KB | 3.30KB | columnar_v3 |
| 2 | Success | 2020-09-28 22:53:01.702 | 6.688S | {} | 6.47KB | 3.30KB | columnar_v3 |
| 1 | Compacted | 2020-09-28 22:51:15.242 | 5.82S | {} | 6.50KB | 3.43KB |
columnar_v3 |
| 0.1 | Success | 2020-10-30 20:49:24.561 | 16.66S | {} | 12.87KB | 6.91KB | columnar_v3 |
| 0 | Compacted | 2020-09-28 22:51:02.6 | 6.819S | {} | 6.50KB | 3.43KB | columnar_v3 |
+-----+-----+-----+-----+-----+-----+-----+
+--+

```

In the preceding information:

- **Compacted** indicates that data has been compacted.
- **0.1** indicates the compacting result of segment 0 and segment 1.

The compact operation does not incur any change to other operations.

Compacted segments, such as segment 0 and segment 1, become useless. To save space, before you perform other operations, run the **CLEAN FILES** command to delete compacted segments. For more information about the **CLEAN FILES** command, see [CLEAN FILES](#).

3.6.1.6 TABLE RENAME

Function

This command is used to rename an existing table.

Syntax

```
ALTER TABLE [db_name.]table_name RENAME TO new_table_name;
```

Parameter Description

Table 3-28 RENAME parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is selected.

Parameter	Description
table_name	Current name of the existing table
new_table_name	New name of the existing table

Precautions

- Parallel queries (using table names to obtain paths for reading CarbonData storage files) may fail during this operation.
- The secondary index table cannot be renamed.

Example

```
ALTER TABLE carbon RENAME TO carbondata;
```

```
ALTER TABLE test_db.carbon RENAME TO test_db.carbondata;
```

System Response

The new table name will be displayed in the CarbonData folder. You can run **SHOW TABLES** to view the new table name.

3.6.1.7 ADD COLUMNS

Function

This command is used to add a column to an existing table.

Syntax

```
ALTER TABLE [db_name.]table_name ADD COLUMNS (col_name data_type,...)
TBLPROPERTIES ('COLUMNPROPERTIES.columnName.shared_column'='sharedFolder.sharedColumnName,...', 'DEFAULT.VALUE.COLUMN_NAME'='default_value');
```

Parameter Description

Table 3-29 ADD COLUMNS parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is selected.
table_name	Table name.

Parameter	Description
col_name data_type	Name of a comma-separated column with a data type. It consists of letters, digits, and underscores (_). NOTE When creating a CarbonData table, do not name columns as tupleId, PositionId, and PositionReference because they will be used in UPDATE, DELETE, and secondary index commands.

Precautions

- Only **shared_column** and **default_value** are read. If any other property name is specified, no error will be thrown and the property will be ignored.
- If no default value is specified, the default value of the new column is considered null.
- If filter is applied to the column, new columns will not be added during sort. New columns may affect query performance.

Examples

- **ALTER TABLE carbon ADD COLUMNS (a1 INT, b1 STRING);**
- **ALTER TABLE carbon ADD COLUMNS (a1 INT, b1 STRING)
TBLPROPERTIES ('COLUMNPROPERTIES.b1.shared_column'='sharedFolder.b1');**
- **ALTER TABLE carbon ADD COLUMNS (a1 INT, b1 STRING)
TBLPROPERTIES ('DEFAULT.VALUE.a1'='10');**

System Response

The newly added column can be displayed by running the **DESCRIBE** command.

3.6.1.8 DROP COLUMNS

Function

This command is used to delete one or more columns from a table.

Syntax

```
ALTER TABLE [db_name.]table_name DROP COLUMNS (col_name, ...);
```

Parameter Description

Table 3-30 DROP COLUMNS parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is selected.

Parameter	Description
table_name	Table name.
col_name	Name of a column in a table. Multiple columns are supported. It consists of letters, digits, and underscores (_).

Precautions

After a column is deleted, at least one key column must exist in the schema. Otherwise, an error message is displayed, and the column fails to be deleted.

Examples

Assume that the table contains four columns named a1, b1, c1, and d1.

- Delete a column:
ALTER TABLE *carbon* DROP COLUMNS (*b1*);
ALTER TABLE *test_db.carbon* DROP COLUMNS (*b1*);
- Delete multiple columns:
ALTER TABLE *carbon* DROP COLUMNS (*b1,c1*);
ALTER TABLE *test_db.carbon* DROP COLUMNS (*b1,c1*);

System Response

If you run the **DESCRIBE** command, the deleted columns will not be displayed.

3.6.1.9 CHANGE DATA TYPE

Function

This command is used to change the data type from INT to BIGINT or decimal precision from lower to higher.

Syntax

```
ALTER TABLE [db_name.]table_name CHANGE col_name col_name changed_column_type;
```

Parameter Description

Table 3-31 CHANGE DATA TYPE parameters

Parameter	Description
db_name	Name of the database. If this parameter is left unspecified, the current database is selected.
table_name	Name of the table.

Parameter	Description
col_name	Name of columns in a table. Column names contain letters, digits, and underscores (_).
changed_column_type	The change in the data type.

Usage Guidelines

- Change of decimal data type from lower precision to higher precision will only be supported for cases where there is no data loss.
Example:
 - **Invalid scenario** - Change of decimal precision from (10,2) to (10,5) is not valid as in this case only scale is increased but total number of digits remain the same.
 - **Valid scenario** - Change of decimal precision from (10,2) to (12,3) is valid as the total number of digits are increased by 2 but scale is increased only by 1 which will not lead to any data loss.
- The allowed range is 38,38 (precision, scale) and is a valid upper case scenario which is not resulting in data loss.

Examples

- Changing data type of column a1 from INT to BIGINT.
ALTER TABLE test_db.carbon CHANGE a1 a1 BIGINT;
- Changing decimal precision of column a1 from 10 to 18.
ALTER TABLE test_db.carbon CHANGE a1 a1 DECIMAL(18,2);

System Response

By running DESCRIBE command, the changed data type for the modified column is displayed.

3.6.1.10 REFRESH TABLE

Function

This command is used to register Carbon table to Hive meta store catalogue from existing Carbon table data.

Syntax

REFRESH TABLE db_name.table_name;

Parameter Description

Table 3-32 REFRESH TABLE parameters

Parameter	Description
db_name	Name of the database. If this parameter is left unspecified, the current database is selected.
table_name	Name of the table.

Usage Guidelines

- The new database name and the old database name should be same.
- Before executing this command the old table schema and data should be copied into the new database location.
- If the table is aggregate table, then all the aggregate tables should be copied to the new database location.
- For old store, the time zone of the source and destination cluster should be same.
- If old cluster used HIVE meta store to store schema, refresh will not work as schema file does not exist in file system.

Examples

```
REFRESH TABLE dbcarbon.productSalesTable;
```

System Response

By running this command, the Carbon table will be registered to Hive meta store catalogue from existing Carbon table data.

3.6.1.11 REGISTER INDEX TABLE

Function

This command is used to register an index table with the primary table.

Syntax

```
REGISTER INDEX TABLE indextable_name ON db_name.maintable_name;
```

Parameter Description

Table 3-33 REFRESH INDEX TABLE parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is selected.
indextable_name	Index table name.
maintable_name	Primary table name.

Precautions

Before running this command, run **REFRESH TABLE** to register the primary table and secondary index table with the Hive metastore.

Examples

```
create database productdb;
```

```
use productdb;
```

```
CREATE TABLE productSalesTable(a int,b string,c string) stored as carbondata;
```

```
create index productNameIndexTable on table productSalesTable(c) as  
'carbondata';
```

```
insert into table productSalesTable select 1,'a','aaa';
```

```
create database productdb2;
```

Run the **hdfs** command to copy **productSalesTable** and **productNameIndexTable** in the **productdb** database to the **productdb2** database.

```
refresh table productdb2.productSalesTable ;
```

```
refresh table productdb2.productNameIndexTable ;
```

```
explain select * from productdb2.productSalesTable where c = 'aaa'; / The  
query command does not use an index table.
```

```
REGISTER INDEX TABLE productNameIndexTable ON  
productdb2.productSalesTable;
```

```
explain select * from productdb2.productSalesTable where c = 'aaa'; // The  
query command uses an index table.
```

System Response

By running this command, the index table will be registered to the primary table.

3.6.2 DML

3.6.2.1 LOAD DATA

Function

This command is used to load user data of a particular type, so that CarbonData can provide good query performance.

 **NOTE**

Only the raw data on HDFS can be loaded.

Syntax

```
LOAD DATA INPATH 'folder_path' INTO TABLE [db_name.]table_name  
OPTIONS(property_name=property_value, ...);
```

Parameter Description

Table 3-34 LOAD DATA parameters

Parameter	Description
folder_path	Path of the file or folder used for storing the raw CSV data.
db_name	Database name. If this parameter is not specified, the current database is used.
table_name	Name of a table in a database.

Precautions

The following configuration items are involved during data loading:

- **DELIMITER:** Delimiters and quote characters provided in the load command. The default value is a comma (,).

```
OPTIONS('DELIMITER',';', 'QUOTECHAR'='')
```

You can use '**DELIMITER**'='\t' to separate CSV data using tabs.

```
OPTIONS('DELIMITER'='\t')
```

CarbonData also supports **\001** and **\017** as delimiters.

 **NOTE**

When the delimiter of CSV data is a single quotation mark ('), the single quotation mark must be enclosed in double quotation marks (" "). For example, '**DELIMITER**'="''".

- **QUOTECHAR:** Delimiters and quote characters provided in the load command. The default value is double quotation marks (").

```
OPTIONS('DELIMITER',';', 'QUOTECHAR'='')
```
- **COMMENTCHAR:** Comment characters provided in the load command. During data loading, if there is a comment character at the beginning of a line, the line is regarded as a comment line and data in the line will not be loaded. The default value is a pound key (#).

OPTIONS('COMMENTCHAR'='#')

- **FILEHEADER:** If the source file does not contain any header, add a header to the **LOAD DATA** command.

OPTIONS('FILEHEADER'='column1,column2')

- **ESCAPECHAR:** Is used to perform strict verification of the escape character on CSV files. The default value is backslash (\).

OPTIONS('ESCAPECHAR'='\')

 **NOTE**

Enter **ESCAPECHAR** in the CSV data. **ESCAPECHAR** must be enclosed in double quotation marks (" "). For example, "a\b".

- Bad records handling:

In order for the data processing application to provide benefits, certain data integration is required. In most cases, data quality problems are caused by data sources.

Methods of handling bad records are as follows:

- Load all of the data before dealing with the errors.
- Clean or delete bad records before loading data or stop the loading when bad records are found.

There are many options for clearing source data during CarbonData data loading, as listed in [Table 3-35](#).

Table 3-35 Bad Records Logger

Configuration Item	Default Value	Description
BAD_RECORDS_LOGGER_ENABLE	false	Whether to create logs with details about bad records

Configuration Item	Default Value	Description
BAD_RECORDS_ACTION	FAIL	<p>The four types of actions for bad records are as follows:</p> <ul style="list-style-type: none"> ● FORCE: Auto-corrects the data by storing the bad records as NULL. ● REDIRECT: Bad records cannot be loaded and written to the CSV file in BAD_RECORD_PATH. This mode is disabled by default. To enable it, set carbon.enable.badrecord.action.redirect to true. ● IGNORE: Bad records are neither loaded nor written to the CSV file. ● FAIL: Data loading fails if any bad records are found. <p>NOTE In loaded data, if all records are bad records, BAD_RECORDS_ACTION is invalid and the load operation fails.</p>
IS_EMPTY_DATA_BAD_RECORD	false	Whether empty data of a column to be considered as bad record or not. If this parameter is set to false , empty data ("",', or,) is not considered as bad records. If this parameter is set to true , empty data is considered as bad records.
BAD_RECORD_PATH	-	HDFS path where bad records are stored. The default value is Null . If bad records logging or bad records operation redirection is enabled, the path must be configured by the user.

Example:

```
LOAD DATA INPATH 'filepath.csv' INTO TABLE tablename
OPTIONS('BAD_RECORDS_LOGGER_ENABLE'='true',
'BAD_RECORD_PATH'='hdfs://hacluster/tmp/carbon',
'BAD_RECORDS_ACTION'='REDIRECT',
'IS_EMPTY_DATA_BAD_RECORD'='false');
```

 **NOTE**

If **REDIRECT** is used, CarbonData will add all bad records into a separate CSV file. However, this file must not be used for subsequent data loading because the content may not exactly match the source record. You must clean up the source record for further data ingestion. This option is used to remind you which records are bad.

- **MAXCOLUMNS:** (Optional) Specifies the maximum number of columns parsed by a CSV parser in a line.

OPTIONS('MAXCOLUMNS'='400')

Table 3-36 MAXCOLUMNS

Name of the Optional Parameter	Default Value	Maximum Value
MAXCOLUMNS	2000	20000

Table 3-37 Behavior chart of MAXCOLUMNS

MAXCOLUMNS Value	Number of Columns in the File Header	Final Value Considered
Not specified in Load options	5	2000
Not specified in Load options	6000	6000
40	7	Max (column count of file header, MAXCOLUMNS value)
22000	40	20000
60	Not specified in Load options	Max (Number of columns in the first line of the CSV file, MAXCOLUMNS value)

 **NOTE**

There must be sufficient executor memory for setting the maximum value of **MAXCOLUMNS Option**. Otherwise, data loading will fail.

- If **SORT_SCOPE** is set to **GLOBAL_SORT** during table creation, you can specify the number of partitions to be used when sorting data. If this parameter is not set or is set to a value less than **1**, the number of map tasks is used as the number of reduce tasks. It is recommended that each reduce task process 512 MB to 1 GB data.

OPTIONS('GLOBAL_SORT_PARTITIONS'='2')

 **NOTE**

To increase the number of partitions, you may need to increase the value of **spark.driver.maxResultSize**, as the sampling data collected in the driver increases with the number of partitions.

- **DATEFORMAT**: Specifies the date format of the table.

```
OPTIONS('DATEFORMAT'='dateFormat')
```

 **NOTE**

Date formats are specified by date pattern strings. The date pattern letters in Carbon are same as in JAVA.

- **TIMESTAMPFORMAT**: Specifies the timestamp of a table.
- *OPTIONS('TIMESTAMPFORMAT'='timestampFormat')*
- **SKIP_EMPTY_LINE**: Ignores empty rows in the CSV file during data loading.

```
OPTIONS('SKIP_EMPTY_LINE'='TRUE/FALSE')
```

- **Optional: SCALE_FACTOR**: Used to control the number of partitions for **RANGE_COLUMN, SCALE_FACTOR**. The formula is as follows:

```
splitSize = max(blocklet_size, (block_size - blocklet_size)) * scale_factor  
numPartitions = total size of input data / splitSize
```

The default value is **3**. The value ranges from **1** to **300**.

```
OPTIONS('SCALE_FACTOR'='10')
```

 **NOTE**

- If **GLOBAL_SORT_PARTITIONS** and **SCALE_FACTOR** are used at the same time, only **GLOBAL_SORT_PARTITIONS** is valid.
- The compaction on **RANGE_COLUMN** will use **LOCAL_SORT** by default.

Scenarios

To load a CSV file to a CarbonData table, run the following statement:

```
LOAD DATA INPATH 'folder path' INTO TABLE tablename  
OPTIONS(property_name=property_value, ...);
```

Examples

The data in the **data.csv** file is as follows:

```
ID,date,country,name,phonetype,serialname,salary  
4,2014-01-21 00:00:00,xxx,aaa4,phone2435,ASD66902,15003  
5,2014-01-22 00:00:00,xxx,aaa5,phone2441,ASD90633,15004  
6,2014-03-07 00:00:00,xxx,aaa6,phone294,ASD59961,15005
```

```
CREATE TABLE carbontable(ID int, date Timestamp, country String, name String,  
phonetype String, serialname String,salary int) STORED AS carbondata;
```

```
LOAD DATA inpath 'hdfs://hacluster/tmp/data.csv' INTO table carbontable  
options('DELIMITER'=',');
```

System Response

Success or failure will be recorded in the driver logs.

3.6.2.2 UPDATE CARBON TABLE

Function

This command is used to update the CarbonData table based on the column expression and optional filtering conditions.

Syntax

- Syntax 1:
UPDATE <CARBON TABLE> SET (column_name1, column_name2, ... column_name n) = (column1_expression , column2_expression , column3_expression ... column n_expression) [WHERE { <filter_condition> }];
- Syntax 2:
UPDATE <CARBON TABLE> SET (column_name1, column_name2,) = (select sourceColumn1, sourceColumn2 from sourceTable [WHERE { <filter_condition> }]) [WHERE { <filter_condition> }];

Parameter Description

Table 3-38 UPDATE parameters

Parameter	Description
CARBON TABLE	Name of the CarbonData table to be updated
column_name	Target column to be updated
sourceColumn	Column value of the source table that needs to be updated in the target table
sourceTable	Table from which the records are updated to the target table

Precautions

Note the following before running this command:

- The UPDATE command fails if multiple input rows in the source table are matched with a single row in the target table.
- If the source table generates empty records, the UPDATE operation completes without updating the table.
- If rows in the source table do not match any existing rows in the target table, the UPDATE operation completes without updating the table.
- UPDATE is not allowed in the table with secondary index.
- In a subquery, if the source table and target table are the same, the UPDATE operation fails.
- The UPDATE operation fails if the subquery used in the UPDATE command contains an aggregate function or a GROUP BY clause.

For example, **update t_carbn01 a set (a.item_type_code, a.profit) = (select b.item_type_cd, sum(b.profit) from t_carbn01b b where item_type_cd =2 group by item_type_code);**

In the preceding example, aggregate function **sum(b.profit)** and GROUP BY clause are used in the subquery. As a result, the UPDATE operation will fail.

- If the **carbon.input.segments** property has been set for the queried table, the UPDATE operation fails. To solve this problem, run the following statement before the query:

Syntax:

SET carbon.input.segments. <database_name>. <table_name>=*;

Examples

- Example 1:
update carbonTable1 d set (d.column3,d.column5) = (select s.c33 ,s.c55 from sourceTable1 s where d.column1 = s.c11) where d.column1 = 'country' exists(select * from table3 o where o.c2 > 1);
- Example 2:
update carbonTable1 d set (c3) = (select s.c33 from sourceTable1 s where d.column1 = s.c11) where exists(select * from iud.other o where o.c2 > 1);
- Example 3:
update carbonTable1 set (c2, c5) = (c2 + 1, concat(c5 , "y"));
- Example 4:
update carbonTable1 d set (c2, c5) = (c2 + 1, "xyx") where d.column1 = 'india';
- Example 5:
update carbonTable1 d set (c2, c5) = (c2 + 1, "xyx") where d.column1 = 'india' and exists(select * from table3 o where o.column2 > 1);

System Response

Success or failure will be recorded in the driver log and on the client.

3.6.2.3 DELETE RECORDS from CARBON TABLE

Function

This command is used to delete records from a CarbonData table.

Syntax

DELETE FROM CARBON_TABLE [WHERE expression];

Parameter Description

Table 3-39 DELETE RECORDS parameters

Parameter	Description
CARBON TABLE	Name of the CarbonData table in which the DELETE operation is performed

Precautions

- If a segment is deleted, all secondary indexes associated with the segment are deleted as well.
- If the **carbon.input.segments** property has been set for the queried table, the DELETE operation fails. To solve this problem, run the following statement before the query:

Syntax:

```
SET carbon.input.segments. <database_name>.<table_name>=*;
```

Examples

- Example 1:
delete from columncarbonTable1 d where d.column1 = 'country';
- Example 2:
delete from dest where column1 IN ('country1', 'country2');
- Example 3:
delete from columncarbonTable1 where column1 IN (select column11 from sourceTable2);
- Example 4:
delete from columncarbonTable1 where column1 IN (select column11 from sourceTable2 where column1 = 'xxx');
- Example 5:
delete from columncarbonTable1 where column2 >= 4;

System Response

Success or failure will be recorded in the driver log and on the client.

3.6.2.4 INSERT INTO CARBON TABLE

Function

This command is used to add the output of the SELECT command to a Carbon table.

Syntax

```
INSERT INTO [CARBON TABLE] [select query];
```

Parameter Description

Table 3-40 INSERT INTO parameters

Parameter	Description
CARBON TABLE	Name of the CarbonData table to be inserted
select query	SELECT query on the source table (CarbonData, Hive, and Parquet tables are supported)

Precautions

- A table has been created.
- You must belong to the data loading group in order to perform data loading operations. By default, the data loading group is named **ficommon**.
- CarbonData tables cannot be overwritten.
- The data type of the source table and the target table must be the same. Otherwise, data in the source table will be regarded as bad records.
- The **INSERT INTO** command does not support partial success. If bad records exist, the command fails.
- When you insert data of the source table to the target table, you cannot upload or update data of the source table.

To enable data loading or updating during the INSERT operation, set the following parameter to **true**.

carbon.insert.persist.enable=true

By default, the preceding parameters are set to **false**.

NOTE

Enabling this property will reduce the performance of the INSERT operation.

Example

```
create table carbon01(a int,b string,c string) stored as carbondata;  
insert into table carbon01 values(1,'a','aa'),(2,'b','bb'),(3,'c','cc');  
create table carbon02(a int,b string,c string) stored as carbondata;  
INSERT INTO carbon02 select * from carbon01 where a > 1;
```

System Response

Success or failure will be recorded in the driver logs.

3.6.2.5 DELETE SEGMENT by ID

Function

This command is used to delete segments by the ID.

Syntax

```
DELETE FROM TABLE db_name.table_name WHERE SEGMENT.ID IN  
(segment_id1,segment_id2);
```

Parameter Description

Table 3-41 DELETE SEGMENT parameters

Parameter	Description
segment_id	ID of the segment to be deleted.
db_name	Database name. If the parameter is not specified, the current database is used.
table_name	The name of the table in a specific database.

Usage Guidelines

Segments cannot be deleted from the stream table.

Examples

```
DELETE FROM TABLE CarbonDatabase.CarbonTable WHERE SEGMENT.ID IN  
(0);
```

```
DELETE FROM TABLE CarbonDatabase.CarbonTable WHERE SEGMENT.ID IN  
(0,5,8);
```

System Response

Success or failure will be recorded in the CarbonData log.

3.6.2.6 DELETE SEGMENT by DATE

Function

This command is used to delete segments by loading date. Segments created before a specific date will be deleted.

Syntax

```
DELETE FROM TABLE db_name.table_name WHERE SEGMENT.STARTTIME  
BEFORE date_value;
```

Parameter Description

Table 3-42 DELETE SEGMENT by DATE parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is used.
table_name	Name of a table in the specified database
date_value	Valid date when segments are started to be loaded. Segments before the date will be deleted.

Precautions

Segments cannot be deleted from the stream table.

Example

```
DELETE FROM TABLE db_name.table_name WHERE SEGMENT.STARTTIME
BEFORE '2017-07-01 12:07:20';
```

STARTTIME indicates the loading start time of different loads.

System Response

Success or failure will be recorded in CarbonData logs.

3.6.2.7 SHOW SEGMENTS

Function

This command is used to list the segments of a CarbonData table.

Syntax

```
SHOW SEGMENTS FOR TABLE [db_name.]table_name LIMIT number_of_loads;
```

Parameter Description

Table 3-43 SHOW SEGMENTS FOR TABLE parameters

Parameter	Description
db_name	Database name. If this parameter is not specified, the current database is used.
table_name	Name of a table in the specified database
number_of_loads	Threshold of records to be listed

Precautions

None

Examples

```
create table carbon01(a int,b string,c string) stored as carbondata;
insert into table carbon01 select 1,'a','aa';
insert into table carbon01 select 2,'b','bb';
insert into table carbon01 select 3,'c','cc';
SHOW SEGMENTS FOR TABLE carbon01 LIMIT 2;
```

System Response

ID	Status	Load Start Time	Load Time Taken	Partition	Data Size	Index Size	File Format
3	Success	2020-09-28 22:53:26.336	3.726S	{}	6.47KB	3.30KB	columnar_v3
2	Success	2020-09-28 22:53:01.702	6.688S	{}	6.47KB	3.30KB	columnar_v3

3.6.2.8 CREATE SECONDARY INDEX

Function

This command is used to create secondary indexes in the CarbonData tables.

Syntax

```
CREATE INDEX index_name
ON TABLE [db_name.]table_name (col_name1, col_name2)
AS 'carbondata'
PROPERTIES ('table_blocksize'='256');
```

Parameter Description

Table 3-44 CREATE SECONDARY INDEX parameters

Parameter	Description
index_name	Index table name. It consists of letters, digits, and special characters (_).
db_name	Database name. It consists of letters, digits, and special characters (_).

Parameter	Description
table_name	Name of the database table. It consists of letters, digits, and special characters (_).
col_name	Name of a column in a table. Multiple columns are supported. It consists of letters, digits, and special characters (_).
table_blocksize	Block size of a data file. For details, see Block Size .

Precautions

db_name is optional.

Examples

```
create table productdb.productSalesTable(id int,price int,productName string,city string) stored as carbondata;
```

```
CREATE INDEX productNameIndexTable on table productdb.productSalesTable (productName,city) as 'carbondata' ;
```

In this example, a secondary table named **productdb.productNameIndexTable** is created and index information of the provided column is loaded.

System Response

A secondary index table will be created. Index information related to the provided column will be loaded into the secondary index table. The success message will be recorded in system logs.

3.6.2.9 SHOW SECONDARY INDEXES

Function

This command is used to list all secondary index tables in the CarbonData table.

Syntax

```
SHOW INDEXES ON db_name.table_name;
```

Parameter Description

Table 3-45 SHOW SECONDARY INDEXES parameters

Parameter	Description
db_name	Database name. It consists of letters, digits, and special characters (_).

Parameter	Description
table_name	Name of the database table. It consists of letters, digits, and special characters (_).

Precautions

db_name is optional.

Examples

```
create table productdb.productSalesTable(id int,price int,productName
string,city string) stored as carbondata;
```

```
CREATE INDEX productNameIndexTable on table productdb.productSalesTable
(productName,city) as 'carbondata' ;
```

```
SHOW INDEXES ON productdb.productSalesTable;
```

System Response

All index tables and corresponding index columns in a given CarbonData table will be listed.

3.6.2.10 DROP SECONDARY INDEX

Function

This command is used to delete the existing secondary index table in a specific table.

Syntax

```
DROP INDEX [IF EXISTS] index_name ON [db_name.]table_name;
```

Parameter Description

Table 3-46 DROP SECONDARY INDEX parameters

Parameter	Description
index_name	Name of the index table. Table name contains letters, digits, and underscores (_).
db_Name	Name of the database. If the parameter is not specified, the current database is used.
table_name	Name of the table to be deleted.

Usage Guidelines

In this command, **IF EXISTS** and **db_name** are optional.

Examples

```
DROP INDEX if exists productNameIndexTable ON productdb.productSalesTable;
```

System Response

Secondary Index Table will be deleted. Index information will be cleared in CarbonData table and the success message will be recorded in system logs.

3.6.2.11 CLEAN FILES

Function

After the **DELETE SEGMENT** command is executed, the deleted segments are marked as the **delete** state. After the segments are merged, the status of the original segments changes to **compacted**. The data files of these segments are not physically deleted. If you want to forcibly delete these files, run the **CLEAN FILES** command.

However, running this command may result in a query command execution failure.

Syntax

```
CLEAN FILES FOR TABLE [db_name.]table_name ;
```

Parameter Description

Table 3-47 CLEAN FILES FOR TABLE parameters

Parameter	Description
db_name	Database name. It consists of letters, digits, and underscores (_).
table_name	Name of the database table. It consists of letters, digits, and underscores (_).

Precautions

None

Examples

Add Carbon configuration parameters.

```
carbon.clean.file.force.allowed = true
```

```
create table carbon01(a int,b string,c string) stored as carbondata;
```

```
insert into table carbon01 select 1,'a','aa';
insert into table carbon01 select 2,'b','bb';
delete from table carbon01 where segment.id in (0);
show segments for table carbon01;
CLEAN FILES FOR TABLE carbon01 options('force'='true');
show segments for table carbon01;
```

In this example, all the segments marked as **deleted** and **compacted** are physically deleted.

System Response

Success or failure will be recorded in the driver logs.

3.6.2.12 SET/RESET

Function

This command is used to dynamically add, update, display, or reset the CarbonData properties without restarting the driver.

Syntax

- Add or Update parameter value:
SET *parameter_name=parameter_value*
This command is used to add or update the value of **parameter_name**.
- Display property value:
SET *parameter_name*
This command is used to display the value of **parameter_name**.
- Display session parameter:
SET
This command is used to display all supported session parameters.
- Display session parameters along with usage details:
SET -v
This command is used to display all supported session parameters and their usage details.
- Reset parameter value:
RESET
This command is used to clear all session parameters.

Parameter Description

Table 3-48 SET parameters

Parameter	Description
parameter_name	Name of the parameter whose value needs to be dynamically added, updated, or displayed
parameter_value	New value of parameter_name to be set

Precautions

The following table lists the properties which you can set or clear using the SET or RESET command.

Table 3-49 Properties

Property	Description
carbon.options.bad.records.logger.enable	Whether to enable bad record logger.
carbon.options.bad.records.action	Operations on bad records, for example, force, redirect, fail, or ignore. For more information, see Bad record handling .
carbon.options.is.empty.data.bad.record	Whether the empty data is considered as a bad record. For more information, see Bad record handling .
carbon.options.sort.scope	Scope of the sort during data loading.
carbon.options.bad.record.path	HDFS path where bad records are stored.
carbon.custom.block.distribution	Whether to enable Spark or CarbonData block distribution.
enable.unsafe.sort	Whether to use unsafe sort during data loading. Unsafe sort reduces the garbage collection during data loading, thereby achieving better performance.

Property	Description
carbon.si.lookup.partialstring	<p>If this is set to TRUE, the secondary index uses the starts-with, ends-with, contains, and LIKE partition condition strings.</p> <p>If this is set to FALSE, the secondary index uses only the starts-with partition condition string.</p>
carbon.input.segments	<p>Segment ID to be queried. This property allows you to query a specified segment of a specified table. CarbonScan reads data only from the specified segment ID.</p> <p>Syntax: carbon.input.segments. <database_name>. <table_name> = <list of segment ids ></p> <p>If you want to query a specified segment in multi-thread mode, you can use CarbonSession.threadSet instead of the SET statement.</p> <p>Syntax: CarbonSession.threadSet ("carbon.input.segments. <database_name>. <table_name>","<list of segment ids >");</p> <p>NOTE You are advised not to set this property in the carbon.properties file because all sessions contain the segment list unless session-level or thread-level overwriting occurs.</p>

Examples

- Add or Update:
SET enable.unsafe.sort=true
- Display property value:
SET enable.unsafe.sort
- Show the segment ID list, segment status, and other required details, and specify the segment list to be read:
SHOW SEGMENTS FOR TABLE carbontable1;
SET carbon.input.segments.db.carbontable1 = 1, 3, 9;
- Query a specified segment in multi-thread mode:
CarbonSession.threadSet
("carbon.input.segments.default.carbon_table_MULTI_Thread", "1,3");

- Use **CarbonSession.threadSet** to query segments in a multi-thread environment (Scala code is used as an example):

```
def main(args: Array[String]) {
  Future
  {
    CarbonSession.threadSet("carbon.input.segments.default.carbon_table_Multi_Thread", "1")
    spark.sql("select count(empno) from carbon_table_Multi_Thread").show()
  }
}
```

- Reset:
RESET

System Response

- Success will be recorded in the driver log.
- Failure will be displayed on the UI.

3.6.3 Operation Concurrent Execution

Before performing **DDL** and **DML** operations, you need to obtain the corresponding locks. See **Table 3-50** for details about the locks that need to be obtained for each operation. The check mark (√) indicates that the lock is required. An operation can be performed only after all required locks are obtained.

You can check whether any two operations can be executed concurrently by using the following method: The first two lines in **Table 3-50** indicate two operations. If no column in the two lines is marked with the check mark (√), the two operations can be executed concurrently. That is, if the columns with check marks (√) in the two lines do not exist, the two operations can be executed concurrently.

Table 3-50 List of obtaining locks for operations

Operation	METADATA_LOCK	COMPACTION_LOCK	DROP_TABLE_LOCK	DELETE_SEGMENT_LOCK	CLEANFILES_LOCK	ALTER_PARTITION_LOCK	UPDATE_LOCK	STREAMING_LOCK	CURRENT_LOAD_LOCK	SEGMENT_LOCK
CREATE TABLE	-	-	-	-	-	-	-	-	-	-
CREATE TABLE As SELECT	-	-	-	-	-	-	-	-	-	-
DROP TABLE	√	-	√	-	-	-	-	√	-	-

Operation	METADATA_LOCK	COMPACTION_LOCK	DROP_TABLE_LOCK	DELETE_SEGMENT_LOCK	CLEANFILES_LOCK	ALTER_PARTITION_LOCK	UPDATE_LOCK	STREAMING_LOCK	CURRENT_LOAD_LOCK	SEGMENT_LOCK
ALTER TABLE COMPACTION	-	√	-	-	-	-	√	-	-	-
TABLE RENAME	-	-	-	-	-	-	-	-	-	-
ADD COLUMNS	√	√	-	-	-	-	-	-	-	-
DROP COLUMNS	√	√	-	-	-	-	-	-	-	-
CHANGE DATA TYPE	√	√	-	-	-	-	-	-	-	-
REFRESH TABLE	-	-	-	-	-	-	-	-	-	-
REGISTER INDEX TABLE	√	-	-	-	-	-	-	-	-	-
REFRESH INDEX	-	√	-	-	-	-	-	-	-	-

Operation	METADATA_LOCK	COMPACTION_LOCK	DROP_TABLE_LOCK	DELETE_SEGMENT_LOCK	CLEANFILES_LOCK	ALTER_PARTITION_LOCK	UPDATE_LOCK	STREAMING_LOCK	CURRENT_LOAD_LOCK	SEGMENT_LOCK
LOAD DATA/INSERT INTO	-	-	-	-	-	-	-	-	√	√
UPDATE CARBON TABLE	√	√	-	-	-	-	√	-	-	-
DELETE RECORDS from CARBON TABLE	√	√	-	-	-	-	√	-	-	-
DELETE SEGMENT by ID	-	-	-	√	√	-	-	-	-	-
DELETE SEGMENT by DATE	-	-	-	√	√	-	-	-	-	-
SHOW SEGMENTS	-	-	-	-	-	-	-	-	-	-

Operation	METADATA_LOCK	COMPACTION_LOCK	DROP_TABLE_LOCK	DELETE_SEGMENT_LOCK	CLEANFILES_LOCK	ALTER_PARTITION_LOCK	UPDATE_LOCK	STREAMING_LOCK	CURRENT_LOAD_LOCK	SEGMENT_LOCK
CREATE SECONDARY INDEX	√	√	-	√	-	-	-	-	-	-
SHOW SECONDARY INDEXES	-	-	-	-	-	-	-	-	-	-
DROP SECONDARY INDEX	√	-	√	-	-	-	-	-	-	-
CLEAN FILES	-	-	-	-	-	-	-	-	-	-
SET/RESET	-	-	-	-	-	-	-	-	-	-
Add Hive Partition	-	-	-	-	-	-	-	-	-	-
Drop Hive Partition	√	√	√	√	√	√	-	-	-	-
Drop Partition	√	√	√	√	√	√	-	-	-	-

Operation	METADATA_LOCK	COMPACTION_LOCK	DROP_TABLE_LOCK	DELETE_SEGMENT_LOCK	CLEANFILES_LOCK	ALTER_PARTITION_LOCK	UPDATE_LOCK	STREAMING_LOCK	CURRENT_LOAD_LOCK	SEGMENT_LOCK
Alter table set	√	√	-	-	-	-	-	-	-	-

3.6.4 API

This section describes the APIs and usage methods of Segment. All methods are in the `org.apache.spark.util.CarbonSegmentUtil` class.

The following methods have been abandoned:

```
/**
 * Returns the valid segments for the query based on the filter condition
 * present in carbonScanRdd.
 *
 * @param carbonScanRdd
 * @return Array of valid segments
 */
@deprecated def getFilteredSegments(carbonScanRdd: CarbonScanRDD[InternalRow]): Array[String];
```

Usage Method

Use the following methods to obtain CarbonScanRDD from the query statement:

```
val df=carbon.sql("select * from table where age='12'")
val myscan=df.queryExecution.sparkPlan.collect {
case scan: CarbonDataSourceScan if scan.rdd.isInstanceOf[CarbonScanRDD[InternalRow]] => scan.rdd
case scan: RowDataSourceScanExec if scan.rdd.isInstanceOf[CarbonScanRDD[InternalRow]] => scan.rdd
}.head
val carbonrdd=myscan.asInstanceOf[CarbonScanRDD[InternalRow]]
```

Example:

```
CarbonSegmentUtil.getFilteredSegments(carbonrdd)
```

The filtered segment can be obtained by importing SQL statements.

```
/**
 * Returns an array of valid segment numbers based on the filter condition provided in the sql
 * NOTE: This API is supported only for SELECT Sql (insert into,ctas,... is not supported)
 *
 * @param sql
 * @param sparkSession
 * @return Array of valid segments
 * @throws UnsupportedOperationException because Get Filter Segments API supports if and only
 * if only one carbon main table is present in query.
 */
def getFilteredSegments(sql: String, sparkSession: SparkSession): Array[String];
```

Example:

```
CarbonSegmentUtil.getFilteredSegments("select * from table where age='12'", sparkSession)
```

Import the database name and table name to obtain the list of segments to be merged. The obtained segments can be used as parameters of the `getMergedLoadName` function.

```
/**
 * Identifies all segments which can be merged with MAJOR compaction type.
 * NOTE: This result can be passed to getMergedLoadName API to get the merged load name.
 *
 * @param sparkSession
 * @param tableName
 * @param dbName
 * @return list of LoadMetadataDetails
 */
def identifySegmentsToBeMerged(sparkSession: SparkSession,
    tableName: String,
    dbName: String) : util.List[LoadMetadataDetails];
```

Example:

```
CarbonSegmentUtil.identifySegmentsToBeMerged(sparkSession, "table_test", "default")
```

Import the database name, table name, and obtain all segments which can be merged with CUSTOM compaction type. The obtained segments can be transferred as the parameter of the getMergedLoadName function.

```
/**
 * Identifies all segments which can be merged with CUSTOM compaction type.
 * NOTE: This result can be passed to getMergedLoadName API to get the merged load name.
 *
 * @param sparkSession
 * @param tableName
 * @param dbName
 * @param customSegments
 * @return list of LoadMetadataDetails
 * @throws UnsupportedOperationException if customSegments is null or empty.
 * @throws MalformedCarbonCommandException if segment does not exist or is not valid
 */
def identifySegmentsToBeMergedCustom(sparkSession: SparkSession,
    tableName: String,
    dbName: String,
    customSegments: util.List[String]): util.List[LoadMetadataDetails];
```

Example:

```
val customSegments = new util.ArrayList[String]()
customSegments.add("1")
customSegments.add("2")
CarbonSegmentUtil.identifySegmentsToBeMergedCustom(sparkSession, "table_test", "default",
    customSegments)
```

If a segment list is specified, the merged load name is returned.

```
/**
 * Returns the Merged Load Name for given list of segments
 *
 * @param list of segments
 * @return Merged Load Name
 * @throws UnsupportedOperationException if list of segments is less than 1
 */
def getMergedLoadName(list: util.List[LoadMetadataDetails]): String;
```

Example:

```
val carbonTable = CarbonEnv.getCarbonTable(Option(databaseName), tableName)(sparkSession)
val loadMetadataDetails = SegmentStatusManager.readLoadMetadata(carbonTable.getMetadataPath)
CarbonSegmentUtil.getMergedLoadName(loadMetadataDetails.toList.asJava)
```

3.6.5 Spatial Indexes

Quick Example

```
create table IF NOT EXISTS carbonTable
(
```



```
COLUMN1 BIGINT,  
LONGITUDE BIGINT,  
LATITUDE BIGINT,  
COLUMN2 BIGINT,  
COLUMN3 BIGINT  
)  
STORED AS carbondata  
TBLPROPERTIES  
(  
'SPATIAL_INDEX.mygeohash.type'='geohash',  
'SPATIAL_INDEX.mygeohash.sourcecolumns'='longitude,  
latitude',  
'SPATIAL_INDEX.mygeohash.originLatitude'='39.850713',  
'SPATIAL_INDEX.mygeohash.gridSize'='50',  
'SPATIAL_INDEX.mygeohash.minLongitude'='115.828503',  
'SPATIAL_INDEX.mygeohash.maxLongitude'='720.000  
000',  
'SPATIAL_INDEX.mygeohash.minLatitude'='39.850713',  
'SPATIAL_INDEX.mygeohash.maxLatitude'='720.0  
00000',  
'SPATIAL_INDEX'='mygeohash',  
'SPATIAL_INDEX.mygeohash.conversionRatio'='1000000',  
'SORT_COLUMNS'='column1,column2,column3,latitude,longitude');
```

Introduction to Spatial Indexes


Spatial data includes multidimensional points, lines, rectangles, cubes, polygons, and other geometric objects. A spatial data object occupies a certain region of space, called spatial scope, characterized by its location and boundary. The spatial data can be either point data or region data.

- Point data: A point has a spatial extent characterized completely by its location. It does not occupy space and has no associated boundary. Point data consists of a collection of points in a two-dimensional space. Points can be stored as a pair of longitude and latitude.
- Region data: A region has a spatial extent with a location, and boundary. The location can be considered as the position of a fixed point in the region, such as its centroid. In two dimensions, the boundary can be visualized as a line (for finite regions, a closed loop). Region data contains a collection of regions.

Currently, only point data is supported, and it can be stored.

Longitude and latitude can be encoded as a unique GeoID. Geohash is a public-domain geocoding system invented by Gustavo Niemeyer. It encodes geographical locations into a short string of letters and digits. It is a hierarchical spatial data structure which subdivides the space into buckets of grid shape, which is one of the many applications of what is known as the Z-order curve, and generally the space-filling curve.

The Z value of a point in multiple dimensions is calculated by interleaving the binary representation of its coordinate value, as shown in the following figure. When Geohash is used to create a GeoID, data is sorted by GeoID instead of longitude and latitude. Data is stored by spatial proximity.



Creating a Table

GeoHash encoding:

```
create table IF NOT EXISTS carbonTable
(
...
`LONGITUDE` BIGINT,
`LATITUDE` BIGINT,
...
)
STORED AS carbondata
TBLPROPERTIES
('SPATIAL_INDEX.mygeohash.type='geohash','SPATIAL_INDEX.mygeohash.sourcecolumns='longitude,
latitude','SPATIAL_INDEX.mygeohash.originLatitude='xx.xxxxxx','SPATIAL_INDEX.mygeohash.gridSize='xx','SP
ATIAL_INDEX.mygeohash.minLongitude='xxx.xxxxxx','SPATIAL_INDEX.mygeohash.maxLongitude='xxx.xxxxxx',
'SPATIAL_INDEX.mygeohash.minLatitude='xx.xxxxxx','SPATIAL_INDEX.mygeohash.maxLatitude='xxx.xxxxxx',
'SPATIAL_INDEX='mygeohash','SPATIAL_INDEX.mygeohash.conversionRatio='1000000','SORT_COLUMNS'=co
lumn1,column2,column3,latitude,longitude');
```

SPATIAL_INDEX is a user-defined index handler. This handler allows users to create new columns from the table-structure column set. The new column name is the same as that of the handler name. The **type** and **sourcecolumns** properties of the handler are mandatory. Currently, the value of **type** supports only **geohash**. Carbon provides a default implementation class that can be easily used. You can extend the default implementation class to mount the customized implementation class of **geohash**. The default handler also needs to provide the following table properties:

- **SPATIAL_INDEX.xxx.originLatitude**: specifies the origin latitude. (**Double** type.)
- **SPATIAL_INDEX.xxx.gridSize**: specifies the grid length in meters. (**Int** type.)

- **SPATIAL_INDEX.xxx.minLongitude**: specifies the minimum longitude. (**Double** type.)
- **SPATIAL_INDEX.xxx.maxLongitude**: specifies the maximum longitude. (**Double** type.)
- **SPATIAL_INDEX.xxx.minLatitude**: specifies the minimum latitude. (**Double** type.)
- **SPATIAL_INDEX.xxx.maxLatitude**: specifies the maximum latitude. (**Double** type.)
- **SPATIAL_INDEX.xxx.conversionRatio**: used to convert the small value of the longitude and latitude to an integer. (**Int** type.)

You can add your own table properties to the handlers in the above format and access them in your custom implementation class. **originLatitude**, **gridSize**, and **conversionRatio** are mandatory. Other parameters are optional in Carbon. You can use the **SPATIAL_INDEX.xxx.class** property to specify their implementation classes.

The default implementation class can generate handler column values for **sourcecolumns** in each row and support query based on the **sourcecolumns** filter criteria. The generated handler column is invisible to users. Except the **SORT_COLUMNS** table properties, no DDL commands or properties are allowed to contain the handler column.

 **NOTE**

- By default, the generated handler column is regarded as the sorting column. If **SORT_COLUMNS** does not contain any **sourcecolumns**, add the handler column to the end of the existing **SORT_COLUMNS**. If the handler column has been specified in **SORT_COLUMNS**, its order in **SORT_COLUMNS** remains unchanged.
- If **SORT_COLUMNS** contains any **sourcecolumns** but does not contain the handler column, the handler column is automatically inserted before **sourcecolumns** in **SORT_COLUMNS**.
- If **SORT_COLUMNS** needs to contain any **sourcecolumns**, ensure that the handler column is listed before the **sourcecolumns** so that the handler column can take effect during sorting.

GeoSOT encoding:

```
CREATE TABLE carbontable(
...
longitude DOUBLE,
latitude DOUBLE,
...)
STORED AS carbondata
TBLPROPERTIES ('SPATIAL_INDEX'='xxx',
'SPATIAL_INDEX.xxx.type'='geosot',
'SPATIAL_INDEX.xxx.sourcecolumns'='longitude, latitude',
'SPATIAL_INDEX.xxx.level'='21',
'SPATIAL_INDEX.xxx.class'='org.apache.carbondata.geo.GeoSOTIndex')
```

Table 3-51 Parameter description

Parameter	Description
SPATIAL_INDEX	Specifies the spatial index. Its value is the same as the column name.

Parameter	Description
SPATIAL_INDEX.xxx.type	(Mandatory) The value is set to geosot .
SPATIAL_INDEX.xxx.source columns	(Mandatory) Specifies the source columns for calculating the spatial index. The value must be two existing columns of the double type.
SPATIAL_INDEX.xxx.level	(Optional) Specifies the columns for calculating the spatial index. The default value is 17 , through which you can obtain an accurate result and improve the computing performance.
SPATIAL_INDEX.xxx.class	(Optional) Specifies the implementation class of GeoSOT. The default value is org.apache.carbondata.geo.GeoSOTIndex .

Example:

```
create table geosot(
timevalue bigint,
longitude double,
latitude double)
stored as carbondata
TBLPROPERTIES ('SPATIAL_INDEX'='mygeosot',
'SPATIAL_INDEX.mygeosot.type'='geosot',
'SPATIAL_INDEX.mygeosot.level'='21', 'SPATIAL_INDEX.mygeosot.sourcecolumns'='longitude, latitude');
```

Preparing Data

- Data file 1: **geosotdata.csv**

```
timevalue,longitude,latitude
1575428400000,116.285807,40.084087
1575428400000,116.372142,40.129503
1575428400000,116.187332,39.979316
1575428400000,116.337069,39.951887
1575428400000,116.359102,40.154684
1575428400000,116.736367,39.970323
1575428400000,116.720179,40.009893
1575428400000,116.346961,40.13355
1575428400000,116.302895,39.930753
1575428400000,116.288955,39.999101
1575428400000,116.17609,40.129953
1575428400000,116.725575,39.981115
1575428400000,116.266922,40.179415
1575428400000,116.353706,40.156483
1575428400000,116.362699,39.942444
1575428400000,116.325378,39.963129
```

- Data file 2: **geosotdata2.csv**

```
timevalue,longitude,latitude
1575428400000,120.17708,30.326882
1575428400000,120.180685,30.326327
1575428400000,120.184976,30.327105
1575428400000,120.189311,30.327549
1575428400000,120.19446,30.329698
1575428400000,120.186965,30.329133
1575428400000,120.177481,30.328911
1575428400000,120.169713,30.325614
1575428400000,120.164563,30.322243
1575428400000,120.171558,30.319613
1575428400000,120.176365,30.320687
```

```
1575428400000,120.179669,30.323688
1575428400000,120.181001,30.320761
1575428400000,120.187094,30.32354
1575428400000,120.193574,30.323651
1575428400000,120.186192,30.320132
1575428400000,120.190055,30.317464
1575428400000,120.195376,30.318094
1575428400000,120.160786,30.317094
1575428400000,120.168211,30.318057
1575428400000,120.173618,30.316612
1575428400000,120.181001,30.317316
1575428400000,120.185162,30.315908
1575428400000,120.192415,30.315871
1575428400000,120.161902,30.325614
1575428400000,120.164306,30.328096
1575428400000,120.197093,30.325985
1575428400000,120.19602,30.321651
1575428400000,120.198638,30.32354
1575428400000,120.165421,30.314834
```

Importing Data

The GeoHash default implementation class extends the customized index abstract class. If the handler property is not set to a customized implementation class, the default implementation class is used. You can extend the default implementation class to mount the customized implementation class of **geohash**. The methods of the customized index abstract class are as follows:

- **Init** method: Used to extract, verify, and store the handler property. If the operation fails, the system throws an exception and displays the error information.
- **Generate** method: Used to generate indexes. It generates an index for each row of data.
- **Query** method: Used to generate an index value range list for given input.

The commands for importing data are the same as those for importing common Carbon tables.

```
LOAD DATA inpath '/tmp/geosotdata.csv' INTO TABLE geosot OPTIONS
('DELIMITER'=',');
```

```
LOAD DATA inpath '/tmp/geosotdata2.csv' INTO TABLE geosot OPTIONS
('DELIMITER'=',');
```

NOTE

For details about **geosotdata.csv** and **geosotdata2.csv**, see [Preparing Data](#).

Aggregate Query of Irregular Spatial Sets

Query statements and filter UDFs

- Filtering data based on polygon
IN_POLYGON(pointList)
UDF input parameter

Parameter	Type	Description
pointList	String	Enter multiple points as a string. Each point is presented as longitude latitude . Longitude and latitude are separated by a space. Each pair of longitude and latitude is separated by a comma (,). The longitude and latitude values at the start and end of the string must be the same.

UDF output parameter

Parameter	Type	Description
inOrNot	Boolean	Checks whether data is in the specified polygon_list .

Example:

```
select longitude, latitude from geosot where IN_POLYGON('116.321011 40.123503, 116.137676 39.947911, 116.560993 39.935276, 116.321011 40.123503');
```

- Filtering data based on the polygon list

IN_POLYGON_LIST(polygonList, opType)

UDF input parameters

Parameter	Type	Description
polygonList	String	Inputs multiple polygons as a string. Each polygon is presented as POLYGON ((longitude1 latitude1, longitude2 latitude2, ...)) . Note that there is a space after POLYGON . Longitudes and latitudes are separated by spaces. Each pair of longitude and latitude is separated by a comma (,). The longitudes and latitudes at the start and end of a polygon must be the same. IN_POLYGON_LIST requires at least two polygons. Example: POLYGON ((116.137676 40.163503, 116.137676 39.935276, 116.560993 39.935276, 116.137676 40.163503))

Parameter	Type	Description
opType	String	Performs union, intersection, and subtraction on multiple polygons. Currently, the following operation types are supported: <ul style="list-style-type: none"> • OR: $A \cup B \cup C$ (Assume that three polygons A, B, and C are input.) • AND: $A \cap B \cap C$

UDF output parameter

Parameter	Type	Description
inOrNot	Boolean	Checks whether data is in the specified polygon_list .

Example:

```
select longitude, latitude from geosot where IN_POLYGON_LIST('POLYGON ((120.176433
30.327431,120.171283 30.322245,120.181411 30.314540, 120.190509 30.321653,120.185188
30.329358,120.176433 30.327431)), POLYGON ((120.191603 30.328946,120.184179
30.327465,120.181819 30.321464, 120.190359 30.315388,120.199242 30.324464,120.191603
30.328946))', 'OR');
```

- Filtering data based on the polyline list

IN_POLYLINE_LIST(polylineList, bufferInMeter)

UDF input parameters

Parameter	Type	Description
polylineList	String	Inputs multiple polylines as a string. Each polyline is presented as LINSTRING (longitude1 latitude1, longitude2 latitude2, ...) . Note that there is a space after LINSTRING . Longitudes and latitudes are separated by spaces. Each pair of longitude and latitude is separated by a comma (,). A union will be output based on the data in multiple polylines. Example: LINSTRING (116.137676 40.163503, 116.137676 39.935276, 116.260993 39.935276)
bufferInMeter	Float	Polyline buffer distance, in meters. Right angles are used at the end to create a buffer.

UDF output parameter

Parameter	Type	Description
inOrNot	Boolean	Checks whether data is in the specified polyline_list .

Example:

```
select longitude, latitude from geosot where IN_POLYLINE_LIST('LINESTRING (120.184179 30.327465, 120.191603 30.328946, 120.199242 30.324464, 120.190359 30.315388)', 65);
```

- Filtering data based on the GeoID range list

IN_POLYGON_RANGE_LIST(polygonRangeList, opType)

UDF input parameters

Parameter	Type	Description
polygonRangeList	String	Inputs multiple rangeLists as a string. Each rangeList is presented as RANGELIST (startGeold1 endGeold1, startGeold2 endGeold2, ...) . Note that there is a space after RANGELIST . Start GeolDs and end GeolDs are separated by spaces. Each group of GeolD ranges is separated by a comma (,). Example: RANGELIST (855279368848 855279368850, 855280799610 855280799612, 855282156300 855282157400)
opType	String	Performs union, intersection, and subtraction on multiple rangeLists. Currently, the following operation types are supported: <ul style="list-style-type: none"> • OR: A U B U C (Assume that three rangeLists A, B, and C are input.) • AND: A ∩ B ∩ C

UDF output parameter

Parameter	Type	Description
inOrNot	Boolean	Checks whether data is in the specified polyRange_list .

Example:


```
select mygeosot, longitude, latitude from geosot where IN_POLYGON_RANGE_LIST('RANGELIST
(526549722865860608 526549722865860618, 532555655580483584 532555655580483594)', 'OR');
```

- Performing polygon query

IN_POLYGON_JOIN(GEO_HASH_INDEX_COLUMN, POLYGON_COLUMN)

Perform join query on two tables. One is a spatial data table containing the longitude, latitude, and GeoHashIndex columns, and the other is a dimension table that saves polygon data.

During query, **IN_POLYGON_JOIN UDF**, **GEO_HASH_INDEX_COLUMN**, and **POLYGON_COLUMN** of the polygon table are used. **Polygon_column** specifies the column containing multiple points (longitude and latitude pairs). The first and last points in each row of the Polygon table must be the same. All points in each row form a closed geometric shape.

UDF input parameters

Parameter	Type	Description
GEO_HASH_INDEX_COLUMN	Long	GeoHashIndex column of the spatial data table.
POLYGON_COLUMN	String	Polygon column of the polygon table, the value of which is represented by the string of polygon, for example, POLYGON ((longitude1 latitude1, longitude2 latitude2, ...)) .

Example:

```
CREATE TABLE polygonTable(
polygon string,
poiType string,
poId String)
STORED AS carbondata;

insert into polygonTable select 'POLYGON ((120.176433 30.327431,120.171283 30.322245,
120.181411 30.314540,120.190509 30.321653,120.185188 30.329358,120.176433 30.327431))','abc','1';

insert into polygonTable select 'POLYGON ((120.191603 30.328946,120.184179 30.327465,
120.181819 30.321464,120.190359 30.315388,120.199242 30.324464,120.191603 30.328946))','abc','2';

select t1.longitude,t1.latitude from geosot t1
inner join
(select polygon,poId from polygonTable where poiType='abc') t2
on in_polygon_join(t1.mygeosot,t2.polygon) group by t1.longitude,t1.latitude;
```

- Performing range_list query

IN_POLYGON_JOIN_RANGE_LIST(GEO_HASH_INDEX_COLUMN, POLYGON_COLUMN)

Use the **IN_POLYGON_JOIN_RANGE_LIST UDF** to associate the spatial data table with the polygon dimension table based on **Polygon_RangeList**. By using a range list, you can skip the conversion between a polygon and a range list.

UDF input parameters

Parameter	Type	Description
GEO_HASH_INDEX_COLUMN	Long	GeoHashIndex column of the spatial data table.
POLYGON_COLUMN	String	Rangelist column of the Polygon table, the value of which is represented by the string of rangeList, for example, RANGELIST (startGeoid1 endGeoid1, startGeoid2 endGeoid2, ...) .

Example:

```
CREATE TABLE polygonTable(
polygon string,
poiType string,
poild String)
STORED AS carbondata;

insert into polygonTable select 'RANGELIST (526546455897309184 526546455897309284,
526549831217315840 526549831217315850, 532555655580483534 532555655580483584)', 'xyz', '2';

select t1.*
from geosot t1
inner join
(select polygon, poild from polygonTable where poiType='xyz') t2
on in_polygon_join_range_list(t1.mygeosot, t2.polygon);
```

UDFs of spatial index tools

- Obtaining row number and column number of a grid converted from Geoid
GeoidToGridXy(geoid)

UDF input parameter

Parameter	Type	Description
geoid	Long	Calculates the row number and column number of the grid based on Geoid.

UDF output parameter

Parameter	Type	Description
gridArray	Array[Int]	Returns the grid row and column numbers contained in Geoid in array. The first digit indicates the row number, and the second digit indicates the column number.

Example:

```
select longitude, latitude, mygeohash, GeoldToGridXy(mygeohash) as GridXY from geoTable;
```

- Converting longitude and latitude to Geold

LatLngToGeold(latitude, longitude oriLatitude, gridSize)

UDF input parameters

Parameter	Type	Description
longitude	Long	Longitude. Note: The value is an integer after conversion.
latitude	Long	Latitude. Note: The value is an integer after conversion.
oriLatitude	Double	Origin latitude, required for calculating Geold.
gridSize	Int	Grid size, required for calculating Geold.

UDF output parameter

Parameter	Type	Description
geold	Long	Returns a number that indicates the longitude and latitude after coding.

Example:

```
select longitude, latitude, mygeohash, LatLngToGeold(latitude, longitude, 39.832277, 50) as geold from geoTable;
```

- Converting Geold to longitude and latitude

GeoldToLatLng(geold, oriLatitude, gridSize)

UDF input parameters

Parameter	Type	Description
geold	Long	Calculates the longitude and latitude based on Geold.
oriLatitude	Double	Origin latitude, required for calculating the longitude and latitude.
gridSize	Int	Grid size, required for calculating the longitude and latitude.

 **NOTE**

GeoID is generated based on the grid coordinates, which are the grid center. Therefore, the calculated longitude and latitude are the longitude and latitude of the grid center. There may be an error ranging from 0 degree to half of the grid size between the calculated longitude and latitude and the longitude and latitude of the generated GeoID.

UDF output parameter

Parameter	Type	Description
latitudeAndLongitude	Array[Double]	Returns the longitude and latitude coordinates of the grid center that represent the GeoID in array. The first digit indicates the latitude, and the second digit indicates the longitude.

Example:

```
select longitude, latitude, mygeohash, GeoidToLatLng(mygeohash, 39.832277, 50) as LatitudeAndLongitude from geoTable;
```

- Calculating the upper-layer GeoID of the pyramid model

ToUpperLayerGeoid(geoid)

UDF input parameter

Parameter	Type	Description
geoid	Long	Calculates the upper-layer GeoID of the pyramid model based on the input GeoID.

UDF output parameter

Parameter	Type	Description
geoid	Long	Returns the upper-layer GeoID of the pyramid model.

Example:

```
select longitude, latitude, mygeohash, ToUpperLayerGeoid(mygeohash) as upperLayerGeoid from geoTable;
```

- Obtaining the GeoID range list using the input polygon

ToRangeList(polygon, oriLatitude, gridSize)

UDF input parameters

Parameter	Type	Description
polygon	String	Input polygon string, which is a pair of longitude and latitude. Longitude and latitude are separated by a space. Each pair of longitude and latitude is separated by a comma (,). The longitude and latitude at the start and end must be the same.
oriLatitude	Double	Origin latitude, required for calculating GeoID.
gridSize	Int	Grid size, required for calculating GeoID.

UDF output parameter

Parameter	Type	Description
geoldList	Buffer[Array[Long]]	Converts polygons into GeoID range lists.

Example:

```
select ToRangeList('116.321011 40.123503, 116.137676 39.947911, 116.560993 39.935276, 116.321011 40.123503', 39.832277, 50) as rangeList from geoTable;
```

- Calculating the upper-layer longitude of the pyramid model

ToUpperLongitude (longitude, gridSize, oriLat)

UDF input parameters

Parameter	Type	Description
longitude	Long	Input longitude, which is a long integer.
gridSize	Int	Grid size, required for calculating longitude.
oriLatitude	Double	Origin latitude, required for calculating longitude.

UDF output parameter

Parameter	Type	Description
longitude	Long	Returns the upper-layer longitude.

Example:

```
select ToUpperLongitude (-23575161504L, 50, 39.832277) as upperLongitude from geoTable;
```

- Calculating the upper-layer latitude of the pyramid model

ToUpperLatitude(Latitude, gridSize, oriLat)

UDF input parameters

Parameter	Type	Description
latitude	Long	Input latitude, which is a long integer.
gridSize	Int	Grid size, required for calculating latitude.
oriLatitude	Double	Origin latitude, required for calculating latitude.

UDF output parameter

Parameter	Type	Description
Latitude	Long	Returns the upper-layer latitude.

Example:

```
select ToUpperLatitude (-23575161504L, 50, 39.832277) as upperLatitude from geoTable;
```

- Converting longitude and latitude to GeoSOT

LatLngToGridCode(latitude, longitude, level)

UDF input parameters

Parameter	Type	Description
latitude	Double	Latitude.
longitude	Double	Longitude.
level	Int	Level. The value range is [0, 32].

UDF output parameter

Parameter	Type	Description
geold	Long	A number that indicates the longitude and latitude after GeoSOT encoding.

Example:

```
select LatLngToGridCode(39.930753, 116.302895, 21) as geold;
```

3.7 CarbonData Troubleshooting

3.7.1 Filter Result Is not Consistent with Hive when a Big Double Type Value Is Used in Filter

Symptom

When double data type values with higher precision are used in filters, incorrect values are returned by filtering results.

Possible Causes

When double data type values with higher precision are used in filters, values are rounded off before comparison. Therefore, values of double data type with different fraction part are considered same.

Troubleshooting Method

NA.

Procedure

To avoid this problem, use decimal data type when high precision data comparisons are required, such as financial applications, equality and inequality checks, and rounding operations.

Reference Information

NA.

3.7.2 Query Performance Deterioration

Symptom

The query performance fluctuates when the query is executed in different query periods.

Possible Causes

During data loading, the memory configured for each executor program instance may be insufficient, resulting in more Java GCs. When GC occurs, the query performance deteriorates.

Troubleshooting Method

On the Spark UI, the GC time of some executors is obviously higher than that of other executors, or all executors have high GC time.

Procedure

Log in to Manager and choose **Cluster > Services > Spark2x**. On the displayed page, click the **Configurations** tab and then **All Configurations**, search for **spark.executor.memory** in the search box, and set its value to a larger value.

spark.executor.memory	4G
-----------------------	----

Reference

None

3.8 CarbonData FAQ

3.8.1 Why Is Incorrect Output Displayed When I Perform Query with Filter on Decimal Data Type Values?

Question

Why is incorrect output displayed when I perform query with filter on decimal data type values?

For example:

```
select * from carbon_table where num = 1234567890123456.22;
```

Output:

```
+-----+-----+-----+
| name |    num    |
+-----+-----+-----+
| IAA | 1234567890123456.22 |
| IAA | 1234567890123456.21 |
+-----+-----+-----+
```

Answer

To obtain accurate output, append BD to the number.

For example:

```
select * from carbon_table where num = 1234567890123456.22BD;
```

Output:

```
+-----+-----+-----+
| name |    num    |
+-----+-----+-----+
| IAA | 1234567890123456.22 |
+-----+-----+-----+
```


3.8.2 How to Avoid Minor Compaction for Historical Data?

Question

How to avoid minor compaction for historical data?

Answer

If you want to load historical data first and then the incremental data, perform following steps to avoid minor compaction of historical data:

1. Load all historical data.
2. Configure the major compaction size to a value smaller than the segment size of historical data.
3. Run the major compaction once on historical data so that these segments will not be considered later for minor compaction.
4. Load the incremental data.
5. You can configure the minor compaction threshold as required.

For example:

1. Assume that you have loaded all historical data to CarbonData and the size of each segment is 500 GB.
2. Set the threshold of major compaction property to **carbon.major.compaction.size = 491520** (480 GB x 1024).
3. Run major compaction. All segments will be compacted because the size of each segment is more than configured size.
4. Perform incremental loading.
5. Configure the minor compaction threshold to **carbon.compaction.level.threshold = 6,6**.
6. Run minor compaction. As a result, only incremental data is compacted.

3.8.3 How to Change the Default Group Name for CarbonData Data Loading?

Question

How to change the default group name for CarbonData data loading?

Answer

By default, the group name for CarbonData data loading is **ficommon**. You can perform the following operation to change the default group name:

1. Edit the **carbon.properties** file.
2. Change the value of the key **carbon.dataload.group.name** as required. The default value is **ficommon**.

3.8.4 Why Does INSERT INTO CARBON TABLE Command Fail?

Question

Why does the *INSERT INTO CARBON TABLE* command fail and the following error message is displayed?

```
Data load failed due to bad record
```

Answer

The *INSERT INTO CARBON TABLE* command fails in the following scenarios:

- If the data type of source and target table columns are not the same, the data from the source table will be treated as bad records and the *INSERT INTO* command fails.
- If the result of aggregation function on a source column exceeds the maximum range of the target column, then the *INSERT INTO* command fails.

Solution:

You can use the cast function on corresponding columns when inserting records.

For example:

- a. Run the *DESCRIBE* command to query the target and source table.

```
DESCRIBE newcarbontable;
```

Result:

```
col1 int  
col2 bigint
```

```
DESCRIBE sourcetable;
```

Result:

```
col1 int  
col2 int
```

- b. Add the cast function to convert bigint value to integer.

```
INSERT INTO newcarbontable select col1, cast(col2 as integer) from  
sourcetable;
```

3.8.5 Why Is the Data Logged in Bad Records Different from the Original Input Data with Escape Characters?

Question

Why is the data logged in bad records different from the original input data with escaped characters?

Answer

An escape character is a backslash (\) followed by one or more characters. If the input records contain escape characters such as \t, \b, \n, \r, \f, \', \", \\, java will process the escape character '\' and the following characters together to obtain the escaped meaning.

For example, if the CSV data type `2010\\10,test` is inserted to String,int type, the value is treated as bad records, because `test` cannot be converted to int. The value logged in the bad records is `2010\10` because java processes `\\` as `\`.

3.8.6 Why Data Load Performance Decreases due to Bad Records?

Question

Why data load performance decreases due to bad records?

Answer

If bad records are present in the data and `BAD_RECORDS_LOGGER_ENABLE` is `true` or `BAD_RECORDS_ACTION` is `redirect` then load performance will decrease due to extra I/O for writing failure reason in log file or redirecting the records to raw CSV.

3.8.7 Why INSERT INTO/LOAD DATA Task Distribution Is Incorrect and the Opened Tasks Are Less Than the Available Executors when the Number of Initial Executors Is Zero?

Question

Why `INSERT INTO` or `LOAD DATA` task distribution is incorrect, and the opened tasks are less than the available executors when the number of initial executors is zero?

Answer

In case of `INSERT INTO` or `LOAD DATA`, CarbonData distributes one task per node. If the executors are not allocated from the distinct nodes then CarbonData will launch fewer tasks.

Solution:

Configure higher value for the executor memory and core so that the yarn can launch only one executor per node.

1. Configure the number of the Executor cores.
 - Configure the `spark.executor.cores` in `spark-defaults.conf` or the `SPARK_EXECUTOR_CORES` in `spark-env.sh` appropriately.
 - Add `--executor-cores NUM` parameter to configure the cores during use the `spark-submit` command.
2. Configure the Executor memory.
 - Configure the `spark.executor.memory` in `spark-defaults.conf` or the `SPARK_EXECUTOR_MEMORY` in `spark-env.sh` appropriately.
 - Add `--executor-memory MEM` parameter to configure the memory during use the `spark-submit` command.

3.8.8 Why Does CarbonData Require Additional Executors Even Though the Parallelism Is Greater Than the Number of Blocks to Be Processed?

Question

Why does CarbonData require additional executors even though the parallelism is greater than the number of blocks to be processed?

Answer

CarbonData block distribution optimizes data processing as follows:

1. Optimize data processing parallelism.
2. Optimize parallel reading of block data.

To optimize parallel processing and parallel read, CarbonData requests executors based on the locality of blocks so that it can obtain executors on all nodes.

If you are using dynamic allocation, you need to configure the following properties:

1. Set **spark.dynamicAllocation.executorIdleTimeout** to 15 minutes (or the average query time).
2. Set **spark.dynamicAllocation.maxExecutors** correctly. The default value **2048** is not recommended. Otherwise, CarbonData will request the maximum number of executors.
3. For a bigger cluster, set **carbon.dynamicAllocation.schedulerTimeout** to a value ranging from 10 to 15 seconds. The default value is 5 seconds.
4. Set **carbon.scheduler.minRegisteredResourcesRatio** to a value ranging from 0.1 to 1.0. The default value is **0.8**. Block distribution can be started as long as the value of **carbon.scheduler.minRegisteredResourcesRatio** is within the range.

3.8.9 Why Data loading Fails During off heap?

Question

Why Data Loading fails during off heap?

Answer

YARN Resource Manager will consider (Java heap memory + **spark.yarn.am.memoryOverhead**) as memory limit, so during the off heap, the memory can exceed this limit. So you need to increase the memory by increasing the value of the parameter **spark.yarn.am.memoryOverhead**.

3.8.10 Why Do I Fail to Create a Hive Table?

Question

Why do I fail to create a hive table?

Answer

Creating a Hive table fails, when source table or sub query has more number of partitions. The implementation of the query requires a lot of tasks, then the number of files will be output a lot, resulting OOM in Driver.

It can be solved by using ***distribute by*** on suitable cardinality(distinct values) column in the statement of Hive table creation.

distribute by clause limits number of hive table partitions. It considers cardinality of given column or **`spark.sql.shuffle.partitions`** which ever is minimal. For example, if **`spark.sql.shuffle.partitions`** is 200, if cardinality of column is 100, out files is 200, but the other 100 files are empty. So using very low cardinality column like 1 will cause data skew and will effect later query distribution.

So we suggest using the column with cardinality greater than **`spark.sql.shuffle.partitions`**. It can be greater than 2 to 3 times.

Example:

```
create table hivetable1 as select * from sourcetable1 distribute by col_age;
```

3.8.11 Why CarbonData tables created in V100R002C50RC1 not reflecting the privileges provided in Hive Privileges for non-owner?

Question

Why CarbonData tables created in V100R002C50RC1 not reflecting the privileges provided in Hive Privileges for non-owner?

Answer

The Hive ACL is implemented after the version V100R002C50RC1, hence the Hive ACL Privileges are not reflecting.

To support HIVE ACL Privileges for CarbonData tables created in V100R002C50RC1, following two ALTER TABLE commands must be executed by owner of the table.

```
ALTER TABLE $dbname.$tablename SET LOCATION '$carbon.store/$dbname/$tablename';
```

```
ALTER TABLE $dbname.$tablename SET SERDEPROPERTIES ('path'='$carbon.store/$dbname/$tablename');
```

Example:

Assume database name is 'carbondb', table name is 'carbontable', and CarbonData store location is 'hdfs://hacluster/user/hive/warehouse/carbon.store', then the commands should be executed is as follows:

```
ALTER TABLE carbondb.carbontable SET LOCATION 'hdfs://hacluster/user/hive/warehouse/carbon.store/carbondb/carbontable';
```

```
ALTER TABLE carbondb.carbontable SET SERDEPROPERTIES ('path'='hdfs://hacluster/user/hive/warehouse/carbon.store/carbondb/carbontable');
```

3.8.12 How Do I Logically Split Data Across Different Namespaces?

Question

How do I logically split data across different namespaces?

Answer

- Configuration:
To logically split data across different namespaces, you must update the following configuration in the **core-site.xml** file of HDFS, Hive, and Spark.

NOTE

Changing the Hive component will change the locations of carbonstore and warehouse.

- Configuration in HDFS
 - **fs.defaultFS**: Name of the default file system. The URI mode must be set to **viewfs**. When **viewfs** is used, the permission part must be **ClusterX**.
 - **fs.viewfs.mountable.ClusterX.homedir**: Home directory base path. You can use the `getHomeDirectory()` method defined in **FileSystem/FileContext** to access the home directory.
 - **fs.viewfs.mountable.default.link.<dir_name>**: ViewFS mount table.

Example:

```
<property>
<name>fs.defaultFS</name>
<value>viewfs://ClusterX</value>
</property>
<property>
<name>fs.viewfs.mounttable.ClusterX.link./folder1</name>
<value>hdfs://NS1/folder1</value>
</property>
<property>
<name>fs.viewfs.mounttable.ClusterX.link./folder2</name>
<value>hdfs://NS2/folder2</value>
</property>
```

- Configurations in Hive and Spark
 - fs.defaultFS**: Name of the default file system. The URI mode must be set to **viewfs**. When **viewfs** is used, the permission part must be **ClusterX**.
- Syntax:

```
LOAD DATA INPATH 'path to data' INTO TABLE table_name OPTIONS ('...');
```

 NOTE

When Spark is configured with the viewFS file system and attempts to load data from HDFS, users must specify a path such as **viewfs://** or a relative path as the file path in the **LOAD** statement.

- Example:

- Sample viewFS path:

```
LOAD DATA INPATH 'viewfs://ClusterX/dir/data.csv' INTO TABLE  
table_name OPTIONS ('...');
```

- Sample relative path:

```
LOAD DATA INPATH '/apps/input_data1.txt' INTO TABLE table_name;
```

3.8.13 Why Missing Privileges Exception is Reported When I Perform Drop Operation on Databases?

Question

Why drop database cascade is throwing the following exception?

```
Error: org.apache.spark.sql.AnalysisException: Missing Privileges;(State=,code=0)
```

Answer

This error is thrown when the owner of the database performs **drop database <database_name> cascade** which contains tables created by other users.

3.8.14 Why the UPDATE Command Cannot Be Executed in Spark Shell?

Question

Why the UPDATE command cannot be executed in Spark Shell?

Answer

The syntax and examples provided in this document are about Beeline commands instead of Spark Shell commands.

To run the UPDATE command in Spark Shell, use the following syntax:

- Syntax 1

```
<carbon_context>.sql("UPDATE <CARBON TABLE> SET (column_name1,  
column_name2, ... column_name n) = (column1_expression ,  
column2_expression , column3_expression ... column n_expression)  
[ WHERE { <filter_condition> } ]").show
```

- Syntax 2

```
<carbon_context>.sql("UPDATE <CARBON TABLE> SET (column_name1,  
column_name2,) = (select sourceColumn1, sourceColumn2 from  
sourceTable [ WHERE { <filter_condition> } ] ) [ WHERE  
{ <filter_condition> } ]").show
```

Example:

If the context of CarbonData is **carbon**, run the following command:

```
carbon.sql("update carbonTable1 d set (d.column3,d.column5) = (select s.c33 ,s.c55 from sourceTable1 s where d.column1 = s.c11) where d.column1 = 'country' exists( select * from table3 o where o.c2 > 1);").show
```

3.8.15 How Do I Configure Unsafe Memory in CarbonData?

Question

How do I configure unsafe memory in CarbonData?

Answer

In the Spark configuration, the value of **spark.yarn.executor.memoryOverhead** must be greater than the sum of (**sort.inmemory.size.inmb + Netty offheapmemory required**), or the sum of (**carbon.unsafe.working.memory.in.mb + carbon.sort.inmemory.storage.size.in.mb + Netty offheapmemory required**). Otherwise, if off-heap access exceeds the configured executor memory, Yarn may stop the executor.

If **spark.shuffle.io.preferDirectBufs** is set to **true**, the netty transfer service in Spark takes off some heap memory (around 384 MB or 0.1 x executor memory) from **spark.yarn.executor.memoryOverhead**.

For details, see [Configuring Executor Off-Heap Memory](#).

3.8.16 Why Exception Occurs in CarbonData When Disk Space Quota is Set for Storage Directory in HDFS?

Question

Why exception occurs in CarbonData when Disk Space Quota is set for the storage directory in HDFS?

Answer

The data will be written to HDFS when you during create table, load table, update table, and so on. If the configured HDFS directory does not have sufficient disk space quota, then the operation will fail and throw following exception.

```
org.apache.hadoop.hdfs.protocol.DSQuotaExceededException:  
The DiskSpace quota of /user/tenant is exceeded:  
quota = 314572800 B = 300 MB but disk space consumed = 402653184 B = 384 MB at  
org.apache.hadoop.hdfs.server.namenode.DirectoryWithQuotaFeature.verifyStoragespaceQuota(DirectoryWith  
hQuotaFeature.java:211) at  
org.apache.hadoop.hdfs.server.namenode.DirectoryWithQuotaFeature.verifyQuota(DirectoryWithQuotaFeatu  
re.java:239) at  
org.apache.hadoop.hdfs.server.namenode.FSDirectory.verifyQuota(FSDirectory.java:941) at  
org.apache.hadoop.hdfs.server.namenode.FSDirectory.updateCount(FSDirectory.java:745)
```

If such exception occurs, configure a sufficient disk space quota for the tenant.

For example:

If the HDFS replication factor is 3 and HDFS default block size is 128 MB, then at least 384 MB (no. of block x block_size x replication_factor of the schema file = 1 x 128 x 3 = 384 MB) disk space quota is required to write a table schema file to HDFS.

 NOTE

In case of fact files, as the default block size is 1024 MB, the minimum space required is 3072 MB per fact file for data load.

3.8.17 Why Does Data Query or Loading Fail and "org.apache.carbondata.core.memory.MemoryException: Not enough memory" Is Displayed?

Question

Why does data query or loading fail and "org.apache.carbondata.core.memory.MemoryException: Not enough memory" is displayed?

Answer

This exception is thrown when the out-of-heap memory required for data query and loading in the executor is insufficient.

In this case, increase the values of **carbon.unsafe.working.memory.in.mb** and **spark.yarn.executor.memoryOverhead**.

For details, see [How Do I Configure Unsafe Memory in CarbonData?](#)

The memory is shared by data query and loading. Therefore, if the loading and query operations need to be performed at the same time, you are advised to set **carbon.unsafe.working.memory.in.mb** and **spark.yarn.executor.memoryOverhead** to a value greater than 2,048 MB.

The following formula can be used for estimation:

Memory required for data loading:

$\text{carbon.number.of.cores.while.loading [default value is 6]} \times \text{Number of tables to load in parallel} \times \text{offheap.sort.chunk.size.inmb [default value is 64 MB]} + \text{carbon.blockletgroup.size.in.mb [default value is 64 MB]} + \text{Current compaction ratio [64 MB/3.5]}$

= Around 900 MB per table

Memory required for data query:

$(\text{SPARK_EXECUTOR_INSTANCES. [default value is 2]} \times (\text{carbon.blockletgroup.size.in.mb [default value: 64 MB]} + \text{carbon.blockletgroup.size.in.mb [default value = 64 MB} \times 3.5]) \times \text{Number of cores per executor [default value: 1]})$

= ~ 600 MB

3.8.18 Why Do Files of a Carbon Table Exist in the Recycle Bin Even If the drop table Command Is Not Executed When Mis-deletion Prevention Is Enabled?

Question

Why do files of a Carbon table exist in the recycle bin even if the **drop table** command is not executed when mis-deletion prevention is enabled?

Answer

After the mis-deletion prevention is enabled for a Carbon table, calling a file deletion command will move the deleted files to the recycle bin. The intermediate file **.carbonindex** is deleted duration the execution of the **insert** or **load** command. Therefore, the table files may exist in the recycle bin even through the **drop table** command is not executed. If you run the **drop table** command, a table directory with a timestamp is generated. The files in the directory are complete.

4 Using ClickHouse

4.1 Using ClickHouse from Scratch

ClickHouse is a column-based database oriented to online analysis and processing. It supports SQL query and provides good query performance. The aggregation analysis and query performance based on large and wide tables is excellent, which is one order of magnitude faster than other analytical databases.

Prerequisites

The client has been installed in a directory, for example, `/opt/client`. The client directory in the following operations is only an example. Change it to the actual installation directory. Before using the client, download and update the client configuration file, and ensure that the active management node of Manager is available.

Procedure

Step 1 Log in to the node where the client is installed as the client installation user.

Step 2 Run the following command to go to the client installation directory:

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the user. The current user must have the permission to create ClickHouse tables. If Kerberos authentication is disabled for the current cluster, skip this step.

For an MRS 3.1.0 cluster, run the **export CLICKHOUSE_SECURITY_ENABLED=true** command first.

```
kinit Component service user
```

Example: **kinit clickhouseuser**

Step 5 Run the client command of the ClickHouse component.

Run the **clickhouse -h** command to view the command help of ClickHouse.

The command output is as follows:

```
Use one of the following commands:
clickhouse local [args]
clickhouse client [args]
clickhouse benchmark [args]
clickhouse server [args]
clickhouse performance-test [args]
clickhouse extract-from-config [args]
clickhouse compressor [args]
clickhouse format [args]
clickhouse copier [args]
clickhouse obfuscator [args]
...
```

For MRS 3.1.0, run the **clickhouse client** command to connect to the ClickHouse server.

- Command for using a non-SSL mode to log in to a ClickHouse cluster with Kerberos authentication disabled
clickhouse client --host *IP address of the ClickHouse instance* --port 9000 --user *Username* --password
Enter the user password.
- Using SSL for login when Kerberos authentication is enabled for the current cluster:
There are no default users in clusters with Kerberos authentication enabled. You must create a user on FusionInsight Manager.
After the user authentication is successful, you do not need to carry the **--user** and **--password** parameters when logging in to the client as the authenticated user.
clickhouse client --host *IP address of the ClickHouse instance* --port 9440 --secure

For MRS 3.1.2 or later, run the **clickhouse client** command to connect to the ClickHouse server.

- Command for using a non-SSL mode to log in to a ClickHouse cluster with Kerberos authentication disabled
clickhouse client --host *IP address of the ClickHouse instance* --port 9000 --user *Username* --password
Enter the user password.
- Using SSL for login when Kerberos authentication is enabled for the current cluster:
There are no default users in clusters with Kerberos authentication enabled. You must create a user on FusionInsight Manager.
clickhouse client --host *IP address of the ClickHouse instance* --port 9440 --user *Username* --password --secure
Enter the user password.

The following table describes the parameters of the **clickhouse client** command.

Table 4-1 Parameters of the **clickhouse client** command

Parameter	Description
--host	<p>Host name of the server. The default value is localhost. You can use the host name or IP address of the node where the ClickHouse instance is located.</p> <p>NOTE You can log in to FusionInsight Manager and choose Cluster > Services > ClickHouse > Instance to obtain the service IP address of the ClickHouseServer instance.</p>
--port	<p>Port for connection.</p> <ul style="list-style-type: none"> • If the SSL security connection is used, the default port number is 9440, the parameter --secure must be carried. For details about the port number, search for the tcp_port_secure parameter in the ClickHouseServer instance configuration. • If non-SSL security connection is used, the default port number is 9000, the parameter --secure does not need to be carried. For details about the port number, search for the tcp_port parameter in the ClickHouseServer instance configuration.
--user	<p>Username.</p> <p>You can create a user on FusionInsight Manager and bind roles to it.</p> <ul style="list-style-type: none"> • If Kerberos authentication has been enabled for the current cluster (the cluster is in security mode) and the user authentication is successful, you do not need to carry the --user and --password parameters during your login to the client as the authenticated user. You must create a user with this name on Manager because there is no default user in the Kerberos cluster scenario. • If Kerberos authentication has not been enabled for the current cluster (the cluster is in normal mode), you cannot use the ClickHouse user created on FusionInsight Manager if you need to specify the username and password when you log in to the client. You need to execute the create user SQL statement on the client to create a ClickHouse user. If you do not need to specify the username and password during your login to the client, the default user is used by default.
--password	<p>Password. The default password is an empty string. This parameter is used together with the --user parameter. You can set a password when creating a user on Manager.</p>
--query	<p>Query to process when using non-interactive mode.</p>
--database	<p>Current default database. The default value is default, which is the default configuration on the server.</p>
--multiline	<p>If this parameter is specified, multiline queries are allowed. (Enter only indicates line feed and does not indicate that the query statement is complete.)</p>

Parameter	Description
-- multiquery	If this parameter is specified, multiple queries separated with semicolons (;) can be processed. This parameter is valid only in non-interactive mode.
--format	Specified default format used to output the result.
--vertical	If this parameter is specified, the result is output in vertical format by default. In this format, each value is printed on a separate line, which helps to display a wide table.
--time	If this parameter is specified, the query execution time is printed to stderr in non-interactive mode.
--stacktrace	If this parameter is specified, stack trace information will be printed when an exception occurs.
--config-file	Name of the configuration file.
--secure	If this parameter is specified, the server will be connected in SSL mode.
-- history_file	Path of files that record command history.
-- param_<name>	Query with parameters. Pass values from the client to the server.

----End

4.2 ClickHouse Table Engine Overview

Background

Table engines play a key role in ClickHouse to determine:

- Where to write and read data
- Supported query modes
- Whether concurrent data access is supported
- Whether indexes can be used
- Whether multi-thread requests can be executed
- Parameters used for data replication

This section describes MergeTree and Distributed engines, which are the most important and frequently used ClickHouse table engines.

MergeTree Family

Engines of the MergeTree family are the most universal and functional table engines for high-load tasks. They have the following key features:

- Data is stored by partition and block based on partitioning keys.
- Data index is sorted based on primary keys and the **ORDER BY** sorting keys.
- Data replication is supported by table engines prefixed with Replicated.
- Data sampling is supported.

When data is written, a table with this type of engine divides data into different folders based on the partitioning key. Each column of data in the folder is an independent file. A file that records serialized index sorting is created. This structure reduces the volume of data to be retrieved during data reading, greatly improving query efficiency.

- MergeTree

Syntax for creating a table:

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]
(
  name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1] [TTL expr1],
  name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2] [TTL expr2],
  ...
  INDEX index_name1 expr1 TYPE type1(...) GRANULARITY value1,
  INDEX index_name2 expr2 TYPE type2(...) GRANULARITY value2
) ENGINE = MergeTree()
ORDER BY expr
[PARTITION BY expr]
[PRIMARY KEY expr]
[SAMPLE BY expr]
[TTL expr [DELETE|TO DISK 'xxx'|TO VOLUME 'xxx'], ...]
[SETTINGS name=value, ...]
```

Example:

```
CREATE TABLE default.test (
  name1 DateTime,
  name2 String,
  name3 String,
  name4 String,
  name5 Date,
  ...
) ENGINE = MergeTree()
PARTITION BY toYYYYMM(name5)
ORDER BY (name1, name2)
SETTINGS index_granularity = 8192
```

Parameters in the example are described as follows:

- **ENGINE = MergeTree()**: specifies the MergeTree engine.
- **PARTITION BY toYYYYMM(name4)**: specifies the partition. The sample data is partitioned by month, and a folder is created for each month.
- **ORDER BY**: specifies the sorting field. A multi-field index can be sorted. If the first field is the same, the second field is used for sorting, and so on.
- **index_granularity = 8192**: specifies the index granularity. One index value is recorded for every 8,192 data records.

If the data to be queried exists in a partition or sorting field, the data query time can be greatly reduced.

- ReplacingMergeTree

Different from MergeTree, ReplacingMergeTree deletes duplicate entries with the same sorting key. ReplacingMergeTree is suitable for clearing duplicate data to save space, but it does not guarantee the absence of duplicate data. Generally, it is not recommended.

Syntax for creating a table:

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]
(
  name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1],
  name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2],
  ...
) ENGINE = ReplacingMergeTree([ver])
[PARTITION BY expr]
[ORDER BY expr]
[SAMPLE BY expr]
[SETTINGS name=value, ...]
```

- **SummingMergeTree**

When merging data parts in SummingMergeTree tables, ClickHouse merges all rows with the same primary key into one row that contains summed values for the columns with the numeric data type. If the primary key is composed in a way that a single key value corresponds to large number of rows, storage volume can be significantly reduced and the data query speed can be accelerated.

Syntax for creating a table:

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]
(
  name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1],
  name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2],
  ...
) ENGINE = SummingMergeTree([columns])
[PARTITION BY expr]
[ORDER BY expr]
[SAMPLE BY expr]
[SETTINGS name=value, ...]
```

Example:

Create a SummingMergeTree table named **testTable**.

```
CREATE TABLE testTable
(
  id UInt32,
  value UInt32
)
ENGINE = SummingMergeTree()
ORDER BY id
```

Insert data into the table.

```
INSERT INTO testTable Values(5,9),(5,3),(4,6),(1,2),(2,5),(1,4),(3,8);
INSERT INTO testTable Values(88,5),(5,5),(3,7),(3,5),(1,6),(2,6),(4,7),(4,6),(43,5),(5,9),(3,6);
```

Query all data in unmerged parts.

```
SELECT * FROM testTable
```

id	value
1	6
2	5
3	8
4	6
5	12

id	value
1	6
2	6
3	18
4	13
5	14
43	5
88	5

If ClickHouse has not summed up all rows and you need to aggregate data by ID, use the **sum** function and **GROUP BY** statement.


```
SELECT id, sum(value) FROM testTable GROUP BY id
```

id	sum(value)
4	19
3	26
88	5
2	11
5	26
1	12
43	5

Merge rows manually.

```
OPTIMIZE TABLE testTable
```

Query data in the **testTable** table again.

```
SELECT * FROM testTable
```

id	value
1	12
2	11
3	26
4	19
5	26
43	5
88	5

SummingMergeTree uses the **ORDER BY** sorting keys as the condition keys to aggregate data. That is, if sorting keys are the same, data records are merged into one and the specified merged fields are aggregated.

Data is pre-aggregated only when merging is executed in the background, and the merging execution time cannot be predicted. Therefore, it is possible that some data has been pre-aggregated and some data has not been aggregated. Therefore, the **GROUP BY** statement must be used during aggregation.

- AggregatingMergeTree

AggregatingMergeTree is a pre-aggregation engine used to improve aggregation performance. When merging partitions, the AggregatingMergeTree engine aggregates data based on predefined conditions, calculates data based on predefined aggregate functions, and saves the data in binary format to tables.

Syntax for creating a table:

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]
(
    name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1],
    name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2],
    ...
) ENGINE = AggregatingMergeTree()
[PARTITION BY expr]
[ORDER BY expr]
[SAMPLE BY expr]
[TTL expr]
[SETTINGS name=value, ...]
```

Example:

You do not need to set the AggregatingMergeTree parameter separately. When partitions are merged, data in each partition is aggregated based on the **ORDER BY** sorting key. You can set the aggregate functions to be used and column fields to be calculated by defining the AggregateFunction type, as shown in the following example:

```
create table test_table (
name1 String,
```

```
name2 String,
name3 AggregateFunction(uniq,String),
name4 AggregateFunction(sum,Int),
name5 DateTime
) ENGINE = AggregatingMergeTree()
PARTITION BY toYYYYMM(name5)
ORDER BY (name1,name2)
PRIMARY KEY name1;
```

When data of the AggregateFunction type is written or queried, the ***state** and ***merge** functions need to be called. The asterisk (*) indicates the aggregate functions used for defining the field type. For example, the **uniq** and **sum** functions are specified for the **name3** and **name4** fields defined in the **test_table**, respectively. Therefore, you need to call the **uniqState** and **sumState** functions and run the **INSERT** and **SELECT** statements when writing data into the table.

```
insert into test_table select '8','test1',uniqState('name1'),sumState(toInt32(100)), '2021-04-30
17:18:00';
insert into test_table select '8','test1',uniqState('name1'),sumState(toInt32(200)), '2021-04-30
17:18:00';
```

When querying data, you need to call the corresponding functions **uniqMerge** and **sumMerge**.

```
select name1,name2,uniqMerge(name3),sumMerge(name4) from test_table group by name1,name2;
```

name1	name2	uniqMerge(name3)	sumMerge(name4)
8	test1	1	300

AggregatingMergeTree is more commonly used with materialized views, which are query views of other data tables at the upper layer.

- CollapsingMergeTree

CollapsingMergeTree defines a **Sign** field to record status of data rows. If **Sign** is **1**, the data in this row is valid. If **Sign** is **-1**, the data in this row needs to be deleted.

Syntax for creating a table:

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]
(
  name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1],
  name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2],
  ...
) ENGINE = CollapsingMergeTree(sign)
[PARTITION BY expr]
[ORDER BY expr]
[SAMPLE BY expr]
[SETTINGS name=value, ...]
```

- VersionedCollapsingMergeTree

The VersionedCollapsingMergeTree engine adds **Version** to the table creation statement to record the mapping between a **state** row and a **cancel** row in case that rows are out of order. The rows with the same primary key, same **Version**, and opposite **Sign** will be deleted during compaction.

Syntax for creating a table:

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]
(
  name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1],
  name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2],
  ...
) ENGINE = VersionedCollapsingMergeTree(sign, version)
[PARTITION BY expr]
[ORDER BY expr]
[SAMPLE BY expr]
[SETTINGS name=value, ...]
```

- GraphiteMergeTree
The GraphiteMergeTree engine is used to store data in the time series database Graphite.

Syntax for creating a table:

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]
(
    Path String,
    Time DateTime,
    Value <Numeric_type>,
    Version <Numeric_type>
    ...
) ENGINE = GraphiteMergeTree(config_section)
[PARTITION BY expr]
[ORDER BY expr]
[SAMPLE BY expr]
[SETTINGS name=value, ...]
```

Replicated*MergeTree Engines

All engines of the MergeTree family in ClickHouse prefixed with Replicated become MergeTree engines that support replicas.



Replicated series engines use ZooKeeper to synchronize data. When a replicated table is created, all replicas of the same shard are synchronized based on the information registered with ZooKeeper.

Template for creating a Replicated engine:

```
ENGINE = Replicated*MergeTree('Storage path in ZooKeeper','Replica name', ...)
```

Two parameters need to be specified for a Replicated engine:


- *Storage path in ZooKeeper*: specifies the path for storing table data in ZooKeeper. The path format is ***/clickhouse/tables/{shard}/Database name/ Table name***.
- *Replica name*: Generally, **{replica}** is used.

For details about the example, see [Creating a ClickHouse Table](#).

Distributed Engine

The Distributed engine does not store any data. It serves as a transparent proxy for data shards and can automatically transmit data to each node in the cluster. Distributed tables need to work with other local data tables. Distributed tables distribute received read and write tasks to each local table where data is stored.

Figure 4-1 Working principle of the Distributed engine



Template for creating a Distributed engine:

```
ENGINE = Distributed(cluster_name, database_name, table_name, [sharding_key])
```

Parameters of a distributed table are described as follows:

- **cluster_name:** specifies the cluster name. When a distributed table is read or written, the cluster configuration information is used to search for the corresponding ClickHouse instance node.
- **database_name:** specifies the database name.
- **table_name:** specifies the name of a local table in the database. It is used to map a distributed table to a local table.
- **sharding_key** (optional): specifies the sharding key, based on which a distributed table distributes data to each local table.

Example:

```
-- Create a ReplicatedMergeTree local table named test.
CREATE TABLE default.test ON CLUSTER default_cluster_1
(
  `EventDate` DateTime,
  `id` UInt64
)
ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/test', '{replica}')
PARTITION BY toYYYYMM(EventDate)
ORDER BY id

-- Create a distributed table named test_all based on the local table test.
CREATE TABLE default.test_all ON CLUSTER default_cluster_1
(
  `EventDate` DateTime,
  `id` UInt64
)
ENGINE = Distributed(default_cluster_1, default, test, rand())
```

Rules for creating a distributed table:

- When creating a distributed table, add **ON CLUSTER** *cluster_name* to the table creation statement so that the statement can be executed once on a ClickHouse instance and then distributed to all instances in the cluster for execution.
- Generally, a distributed table is named in the following format: *Local table name_all*. It forms a one-to-many mapping with local tables. Then, multiple local tables can be operated using the distributed table proxy.
- Ensure that the structure of a distributed table is the same as that of local tables. If they are inconsistent, no error is reported during table creation, but an exception may be reported during data query or insertion.

4.3 Creating a ClickHouse Table

ClickHouse implements the replicated table mechanism based on the ReplicatedMergeTree engine and ZooKeeper. When creating a table, you can specify an engine to determine whether the table is highly available. Shards and replicas of each table are independent of each other.

ClickHouse also implements the distributed table mechanism based on the Distributed engine. Views are created on all shards (local tables) for distributed query, which is easy to use. ClickHouse has the concept of data sharding, which is one of the features of distributed storage. That is, parallel read and write are used to improve efficiency.

The ClickHouse cluster table engine that uses Kunpeng as the CPU architecture does not support HDFS and Kafka.

Viewing cluster and Other Environment Parameters of ClickHouse

Step 1 Use the ClickHouse client to connect to the ClickHouse server by referring to [Using ClickHouse from Scratch](#).

Step 2 Query the cluster identifier and other information about the environment parameters.

```
select cluster,shard_num,replica_num,host_name from system.clusters;
```

```
SELECT  
  cluster,  
  shard_num,  
  replica_num,  
  host_name  
FROM system.clusters
```

cluster	shard_num	replica_num	host_name
default_cluster_1	1	1	node-master1dOnG
default_cluster_1	1	2	node-group-1tXED0001
default_cluster_1	2	1	node-master2OXQS
default_cluster_1	2	2	node-group-1tXED0002
default_cluster_1	3	1	node-master3QsRI
default_cluster_1	3	2	node-group-1tXED0003

6 rows in set. Elapsed: 0.001 sec.

Step 3 Query the shard and replica identifiers.

```
select * from system.macros;
```

```
SELECT *  
FROM system.macros
```

```
macro substitution
id      | 76
replica | 2
shard   | 3
-----|-----
3 rows in set. Elapsed: 0.001 sec.
```

----End

Creating a Local Replicated Table and a distributed Table

Step 1 Log in to the ClickHouse node using the client, for example, `clickhouse client --host node-master3QsRI --multiline --port 9440 --secure;`

 **NOTE**

node-master3QsRI is the value of **host_name** obtained in **Step 2** in [Viewing cluster and Other Environment Parameters of ClickHouse](#).

Step 2 Create a replicated table using the ReplicatedMergeTree engine.

For example, run the following commands to create a ReplicatedMergeTree table named **test** on the **default_cluster_1** node and in the **default** database:

```
CREATE TABLE default.test ON CLUSTER default_cluster_1
(
  `EventDate` DateTime,
  `id` UInt64
)
ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/test',
'{replica}')
PARTITION BY toYYYYMM(EventDate)
ORDER BY id;
```

The parameters are described as follows:

- The **ON CLUSTER** syntax indicates the distributed DDL, that is, the same local table can be created on all instances in the cluster after the statement is executed once.
- **default_cluster_1** is the cluster identifier obtained in **Step 2** in [Viewing cluster and Other Environment Parameters of ClickHouse](#).

CAUTION

ReplicatedMergeTree engine receives the following two parameters:

- Storage path of the table data in ZooKeeper

The path must be in the **/clickhouse** directory. Otherwise, data insertion may fail due to insufficient ZooKeeper quota.

To avoid data conflict between different tables in ZooKeeper, the directory must be in the following format:

/clickhouse/tables/{shard}/default/test, in which **/clickhouse/tables/{shard}** is fixed, *default* indicates the database name, and *test* indicates the name of the created table.

- Replica name: Generally, **{replica}** is used.

```
CREATE TABLE default.test ON CLUSTER default_cluster_1
(
  `EventDate` DateTime,
  `id` UInt64
)
ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/test', '{replica}')
PARTITION BY toYYYYMM(EventDate)
ORDER BY id
```

host	port	status	error	num_hosts_remaining	num_hosts_active
node-group-1tXED0002	9000	0		5	3
node-group-1tXED0003	9000	0		4	3
node-master1dOnG	9000	0		3	3

host	port	status	error	num_hosts_remaining	num_hosts_active
node-master3QsRI	9000	0		2	0
node-group-1tXED0001	9000	0		1	0
node-master2OXQS	9000	0		0	0

6 rows in set. Elapsed: 0.189 sec.

Step 3 Create a distributed table using the Distributed engine.

For example, run the following commands to create a distributed table named **test_all** on the **default_cluster_1** node and in the **default** database:

```
CREATE TABLE default.test_all ON CLUSTER default_cluster_1
```

```
(
  `EventDate` DateTime,
  `id` UInt64
)
ENGINE = Distributed(default_cluster_1, default, test, rand());
```

```
CREATE TABLE default.test_all ON CLUSTER default_cluster_1
(
  `EventDate` DateTime,
  `id` UInt64
)
ENGINE = Distributed(default_cluster_1, default, test, rand())
```

host	port	status	error	num_hosts_remaining	num_hosts_active
------	------	--------	-------	---------------------	------------------

```

┌─┐
│ node-group-1tXED0002      │ 9000 │ 0 │      │ 5 │ 0 │
│ node-master3QsRI         │ 9000 │ 0 │      │ 4 │ 0 │
│ node-group-1tXED0003     │ 9000 │ 0 │      │ 3 │ 0 │
│ node-group-1tXED0001     │ 9000 │ 0 │      │ 2 │ 0 │
│ node-master1dOnG         │ 9000 │ 0 │      │ 1 │ 0 │
│ node-master2OXQS         │ 9000 │ 0 │      │ 0 │ 0 │
└─┘

```

6 rows in set. Elapsed: 0.115 sec.

 **NOTE**

Distributed requires the following parameters:

- **default_cluster_1** is the cluster identifier obtained in [Step 2 in Viewing cluster and Other Environment Parameters of ClickHouse](#).
- **default** indicates the name of the database where the local table is located.
- **test** indicates the name of the local table. In this example, it is the name of the table created in [Step 2](#).
- (Optional) Sharding key

This key and the weight configured in the **config.xml** file determine the route for writing data to the distributed table, that is, the physical table to which the data is written. It can be the original data (for example, **site_id**) of a column in the table or the result of the function call, for example, **rand()** is used in the preceding SQL statement. Note that data must be evenly distributed in this key. Another common operation is to use the hash value of a column with a large difference, for example, **intHash64(user_id)**.

----End

ClickHouse Table Data Operations

Step 1 Log in to the ClickHouse node on the client. Example:

```
clickhouse client --host node-master3QsRI --multiline --port 9440 --secure;
```

 **NOTE**

node-master3QsRI is the value of **host_name** obtained in [Step 2 in Viewing cluster and Other Environment Parameters of ClickHouse](#).

Step 2 After creating a table by referring to [Creating a Local Replicated Table and a distributed Table](#), you can insert data to the local table.

For example, run the following command to insert data to the local table **test**:

```
insert into test values(toDateTime(now()), rand());
```

Step 3 Query the local table information.

For example, run the following command to query data information of the table **test** in [Step 2](#):

```
select * from test;
```

```
SELECT *
FROM test
```

EventDate	id
2020-11-05 21:10:42	1596238076


```
1 rows in set. Elapsed: 0.002 sec.
```

Step 4 Query the distributed table.

For example, the distributed table **test_all** is created based on table **test** in [Step 3](#). Therefore, the same data in table **test** can also be queried in table **test_all**.

select * from test_all;

```
SELECT *  
FROM test_all
```

EventDate	id
2020-11-05 21:10:42	1596238076

```
1 rows in set. Elapsed: 0.004 sec.
```

Step 5 Switch to the shard node with the same **shard_num** and query the information about the current table. The same table data can be queried.

For example, run the **exit;** command to exit the original node.

Run the following command to switch to the **node-group-1tXED0003** node:

clickhouse client --host node-group-1tXED0003 --multiline --port 9440 --secure;

 **NOTE**

The **shard_num** values of **node-group-1tXED0003** and **node-master3QsRI** are the same by performing [Step 2](#).

show tables;

```
SHOW TABLES
```

name
test
test_all

Step 6 Query the local table data. For example, run the following command to query data in table **test** on the **node-group-1tXED0003** node:

select * from test;

```
SELECT *  
FROM test
```

EventDate	id
2020-11-05 21:10:42	1596238076

```
1 rows in set. Elapsed: 0.005 sec.
```

Step 7 Switch to the shard node with different **shard_num** value and query the data of the created table.

For example, run the following command to exit the **node-group-1tXED0003** node:

exit;

Switch to the **node-group-1tXED0001** node. The **shard_num** values of **node-group-1tXED0001** and **node-master3QsRI** are different by performing [Step 2](#).

clickhouse client --host node-group-1tXED0001 --multiline --port 9440 --secure;

Query the local table **test**. Data cannot be queried on the different shard node because table **test** is a local table.

select * from test;

```
SELECT *
FROM test

Ok.
```

Query data in the distributed table **test_all**. The data can be queried properly.

select * from test_all;

```
SELECT *
FROM test

+-----+-----+-----+-----+
| EventDate |      id |
+-----+-----+
| 2020-11-05 21:12:19 | 3686805070 |
+-----+-----+

1 rows in set. Elapsed: 0.002 sec.
```

----End

4.4 ClickHouse Data Type

This section describes the data types of MRS ClickHouse.

For details about the ClickHouse data types, see [ClickHouse Data Types](#).

Table 4-2 ClickHouse data types

Category	Keyword	Data Type	Description
Data type	Int8	Int8	Value range: [-128, 127]
	Int16	Int16	Value range: [-32768, 32767]
	Int32	Int32	Value range: [-2147483648, 2147483647]
	Int64	Int64	Value range: [-9223372036854775808, 9223372036854775807]
Floating point type	Float32	Float	Similar to the Float type in C, a single-precision floating point number occupies 4 bytes in storage of a computer and is described in 32-bit binary mode.

Category	Keyword	Data Type	Description
	Float64	Double	Similar to the Double type in the C language, a double-precision floating point number occupies eight bytes in the machine and is described in 64-bit binary mode.
Decimal type	Decimal	Decimal	<p>A signed fixed-point number that can maintain precision during addition, subtraction, and multiplication operations. The following formats are supported:</p> <ul style="list-style-type: none"> • Decimal(P, S) • Decimal32(S) • Decimal64(S) • Decimal128(S) <p>NOTE</p> <ul style="list-style-type: none"> • P: precision. Valid range: [1:38]. It determines the number of decimal digits (including fractions) that can be contained. • S: scale. The value ranges from 0 to P, which determines the number of decimal places in the decimal part of a number.
String type	String	String	The character string can be of any length. It can contain any set of bytes, including empty bytes. Therefore, the string type can replace the VARCHAR, BLOB, and CLOB types in other DBMSs.

Category	Keyword	Data Type	Description
	FixedString	Fixed-length string	<p>When the length of the data happens to be N bytes, the FixedString type is efficient. In other cases, this may reduce efficiency. Examples of values that can be effectively stored in columns of the FixedString type:</p> <ul style="list-style-type: none"> • IP address represented in binary (IPv6 uses FixedString (16)) • Language code (ru_RU, en_US...) • Currency code • Hash value represented in binary mode (FixedString (16) for MD5 and FixedString (32) for SHA256)
Time and date type	Date	Date	The value is stored in two bytes, indicating the date value from 1970-01-01 (unsigned) to the current time. The time zone information is not stored in the date.
	DateTime	Timestamp	Stores Unix timestamps in four bytes (unsigned). Allows values within the same range as the date type to be stored. The minimum value is 1970-01-01 00:00:00. The timestamp type value is accurate to second (no leap second). The time zone uses the system time zone when the client or server is started.
	DateTime64	DateTime64	This type allows you to store a time value in the form of date plus time.
Boolean type	Boolean	Boolean	ClickHouse does not have a separate type to store Boolean values. The UInt8 type can be used. The value can be 0 or 1.

Category	Keyword	Data Type	Description
Array type	Array	Array	Array(T): array consisting of T-type elements. T The value can be of any type, including the array type. However, multi-dimensional arrays are not recommended. ClickHouse has limited support for multi-dimensional arrays. For example, you cannot store multidimensional arrays in a MergeTree table.
Tuple type	Tuple	Tuple	Tuple(T1, T2, ...), tuple, in which each element has a separate type and tuples cannot be stored in tables (except memory tables). They can be used for temporary column grouping. In a query, IN expressions and lambda functions with specific parameters can be used to group temporary columns.
Domains data type	Domains	Domains	Domains types are implementation-specific types: IPv4 is a Domains type that is binary compatible with the UInt32 type and is used to store the value of an IPv4 address. It provides more compact binary storage while supporting recognition of more readable input and output formats. IPv6 is a domain type that is binary compatible with the FixedString (16) type and is used to store the value of an IPv6 address. It provides more compact binary storage while supporting recognition of more readable input and output formats.
Enumerated type	Enum8	Enum8	Value range: [-128, 127] Enum stores the mapping of 'string'= integer, for example, Enum8('hello' = 1, 'world' = 2).

Category	Keyword	Data Type	Description
	Enum16	Enum16	Value range: [-32768, 32767]
Nullable type	Nullable	Nullable	Unless otherwise specified in the ClickHouse server configuration, NULL is the default value for any Nullable type. Nullable fields cannot be contained in table indexes. It can be stored together with the normal value of TypeName. For example, columns of the Nullable(Int8) type can store values of the Int8 type, while rows without values store NULL.
Nested type	nested	nested	A nested data structure is like a table within a cell. The parameters (column name and type) of the nested data structure are specified in the same way as in the CREATE TABLE query. Each table row can correspond to any number of rows in a nested data structure. Example: Nested (Name1 Type1, Name2 Type2, ...)

4.5 Common ClickHouse SQL Syntax

4.5.1 CREATE DATABASE: Creating a Database

This section describes the basic syntax and usage of the SQL statement for creating a ClickHouse database.

Basic Syntax

```
CREATE DATABASE [IF NOT EXISTS] database_name [ON CLUSTER ClickHouse cluster name]
```

NOTE

The syntax **ON CLUSTER *ClickHouse cluster name*** enables the Data Definition Language (DDL) statement to be executed on all instances in the cluster at a time. You can run the following statement to obtain the cluster name from the **cluster** field:

```
select cluster, shard_num, replica_num, host_name from system.clusters;
```

Example

```
-- Create a database named test.  
CREATE DATABASE test ON CLUSTER default_cluster;  
-- After the creation is successful, run the query command for verification.  
show databases;  
+----+  
| name |  
+----+  
| default |  
| system |  
| test |  
+----+
```

4.5.2 CREATE TABLE: Creating a Table

This section describes the basic syntax and usage of the SQL statement for creating a ClickHouse table.

Basic Syntax

- Method 1: Creating a table named **table_name** in the specified **database_name** database.

If the table creation statement does not contain **database_name**, the name of the database selected during client login is used by default.

```
CREATE TABLE [IF NOT EXISTS] [database_name.]table_name [ON CLUSTER  
ClickHouse cluster name]  
(  
name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1],  
name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2],  
...  
) ENGINE = engine_name()  
[PARTITION BY expr_list]  
[ORDER BY expr_list]
```

CAUTION

You are advised to use **PARTITION BY** to create table partitions when creating a ClickHouse table. The ClickHouse data migration tool migrates data based on table partitions. If you do not use **PARTITION BY** to create table partitions during table creation, the table data cannot be migrated on the GUI in [Using the ClickHouse Data Migration Tool](#).

- Method 2: Creating a table with the same structure as **database_name2.table_name2** and specifying a different table engine for the table

If no table engine is specified, the created table uses the same table engine as **database_name2.table_name2**.

```
CREATE TABLE [IF NOT EXISTS] [database_name.]table_name AS  
[database_name2.]table_name2 [ENGINE = engine_name]
```

- Method 3: Using the specified engine to create a table with the same structure as the result of the **SELECT** clause and filling it with the result of the **SELECT** clause

```
CREATE TABLE [IF NOT EXISTS] [database_name.]table_name ENGINE =  
engine_name AS SELECT ...
```

Example

```
-- Create a table named test in the default database and default_cluster cluster.  
CREATE TABLE default.test ON CLUSTER default_cluster  
(  
  `EventDate` DateTime,  
  `id` UInt64  
)  
ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/test', '{replica}')  
PARTITION BY toYYYYMM(EventDate)  
ORDER BY id
```

4.5.3 INSERT INTO: Inserting Data into a Table

This section describes the basic syntax and usage of the SQL statement for inserting data to a table in ClickHouse.

Basic Syntax

- Method 1: Inserting data in standard format
INSERT INTO [database_name.]table [(c1, c2, c3)] VALUES (v11, v12, v13), (v21, v22, v23), ...
- Method 2: Using the **SELECT** result to insert data
INSERT INTO [database_name.]table [(c1, c2, c3)] SELECT ...

Example

```
-- Insert data into the test2 table.  
insert into test2 (id, name) values (1, 'abc'), (2, 'bbbb');  
-- Query data in the test2 table.  
select * from test2;
```

id	name
1	abc
2	bbbb

4.5.4 SELECT: Querying Table Data

This section describes the basic syntax and usage of the SQL statement for querying table data in ClickHouse.

Basic Syntax

```
SELECT [DISTINCT] expr_list  
[FROM [database_name.]table | (subquery) | table_function] [FINAL]  
[SAMPLE sample_coeff]  
[ARRAY JOIN ...]  
[GLOBAL] [ANY|ALL|ASOF] [INNER|LEFT|RIGHT|FULL|CROSS] [OUTER|SEMI|  
ANTI] JOIN (subquery)|table (ON <expr_list>)|(USING <column_list>)  
[PREWHERE expr]  
[WHERE expr]
```



```
[GROUP BY expr_list] [WITH TOTALS]
[HAVING expr]
[ORDER BY expr_list] [WITH FILL] [FROM expr] [TO expr] [STEP expr]
[LIMIT [offset_value, ]n BY columns]
[LIMIT [n, ]m] [WITH TIES]
[UNION ALL ...]
[INTO OUTFILE filename]
[FORMAT format]
```

Example

```
-- View ClickHouse cluster information.
select * from system.clusters;
-- View the macros set for the current node.
select * from system.macros;
-- Check the database capacity.
select
sum(rows) as "Total number of rows",
formatReadableSize(sum(data_uncompressed_bytes)) as "Original size",
formatReadableSize(sum(data_compressed_bytes)) as "Compression size",
round(sum(data_compressed_bytes) / sum(data_uncompressed_bytes) * 100,
0) "Compression rate"
from system.parts;
-- Query the capacity of the test table. Add or modify the where clause based on the site
requirements.
select
sum(rows) as "Total number of rows",
formatReadableSize(sum(data_uncompressed_bytes)) as "Original size",
formatReadableSize(sum(data_compressed_bytes)) as "Compression size",
round(sum(data_compressed_bytes) / sum(data_uncompressed_bytes) * 100,
0) "Compression rate"
from system.parts
where table in ('test')
and partition like '2020-11-%'
group by table;
```

4.5.5 ALTER TABLE: Modifying a Table Structure

This section describes the basic syntax and usage of the SQL statement for modifying a table structure in ClickHouse.

Basic Syntax

```
ALTER TABLE [database_name].name [ON CLUSTER cluster] ADD|DROP|CLEAR|
COMMENT|MODIFY COLUMN ...
```

NOTE

ALTER supports only MergeTree, Merge, and Distributed engine tables.

Example

```
-- Add the test01 column to the t1 table.
ALTER TABLE t1 ADD COLUMN test01 String DEFAULT 'defaultvalue';
-- Query the modified table t1.
desc t1
┌─name──┬─type──┬─default_type─┬─default_expression─┐
```

```

comment | codec_expression | ttl_expression |
id | UInt8 | | |
name | String | | |
address | String | | |
test01 | String | DEFAULT | 'defaultvalue' |

```

-- Change the type of the **name** column in the **t1** table to UInt8.
ALTER TABLE t1 MODIFY COLUMN name UInt8;
-- Query the modified table **t1**.
desc t1

```

name | type | default_type | default_expression |
comment | codec_expression | ttl_expression |
id | UInt8 | | |
name | UInt8 | | |
address | String | | |
test01 | String | DEFAULT | 'defaultvalue' |

```

-- Delete the **test01** column from the **t1** table.
ALTER TABLE t1 DROP COLUMN test01;
-- Query the modified table **t1**.
desc t1

```

name | type | default_type | default_expression |
comment | codec_expression | ttl_expression |
id | UInt8 | | |
name | UInt8 | | |
address | String | | |

```

4.5.6 DESC: Querying a Table Structure

This section describes the basic syntax and usage of the SQL statement for querying a table structure in ClickHouse.

Basic Syntax

```
DESC|DESCRIBE TABLE [database_name.]table [INTO OUTFILE filename]
[FORMAT format]
```

Example

```

-- Query the t1 table structure.
desc t1;
name | type | default_type | default_expression |
comment | codec_expression | ttl_expression |
id | UInt8 | | |
name | UInt8 | | |
address | String | | |

```

4.5.7 DROP: Deleting a Table

This section describes the basic syntax and usage of the SQL statement for deleting a ClickHouse table.

Basic Syntax

```
DROP [TEMPORARY] TABLE [IF EXISTS] [database_name.]name [ON CLUSTER
cluster] [SYNC]
```

Example

```

-- Delete the t1 table.
drop table t1 SYNC;
```

 NOTE

When you delete a replication table, create a path on ZooKeeper to store related data. The default library engine of ClickHouse is the atomic database engine. After a table in the atomic database is deleted, it is not deleted immediately but deleted 480 seconds later. To resolve this issue, when deleting a table, add the **SYNC** field to the deletion command, for example, **drop table t1 SYNC**;

This issue does not occur when a local or distributed table is deleted. The **SYNC** field is not required in your deletion command, for example, **drop table t1**;

4.5.8 SHOW: Displaying Information About Databases and Tables

This section describes the basic syntax and usage of the SQL statement for displaying information about databases and tables in ClickHouse.

Basic Syntax

show databases

show tables

Example

```
-- Query database information.
```

```
show databases;
```

```
name
default
system
test
```

```
-- Query table information.
```

```
show tables;
```

```
name
t1
test
test2
test5
```

4.6 Migrating ClickHouse Data

4.6.1 Using ClickHouse to Import and Export Data

Using the ClickHouse Client to Import and Export Data

Use the ClickHouse client to import and export data.

- Importing data in CSV format

```
clickhouse client --host Host name or IP address of the ClickHouse instance  
--database Database name --port Port number --secure --  
format_csv_delimiter="CSV file delimiter" --query="INSERT INTO Table  
name FORMAT CSV" < Host path where the CSV file is stored
```

Example

```
clickhouse client --host 10.5.208.5 --database testdb --port 9440 --secure --format_csv_delimiter="," --  
query="INSERT INTO testdb.csv_table FORMAT CSV" < /opt/data
```

You need to create a table in advance.

- Exporting data in CSV format

 **CAUTION**

Exporting data files in CSV format may cause CSV injection. Exercise caution when performing this operation.

clickhouse client *--host* *Host name or IP address of the ClickHouse instance* *--database* *Database name* *--port* *Port number* **-m --secure --query="SELECT * FROM Table name" > CSV file export path**

Example

```
clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="SELECT * FROM test_table" > /opt/test
```

- Importing data in Parquet format

cat Parquet file | clickhouse client *--host* *Host name or IP address of the ClickHouse instance* *--database* *Database name* *--port* *Port number* **-m --secure --query="INSERT INTO Table name FORMAT Parquet"**

Example

```
cat /opt/student.parquet | clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="INSERT INTO parquet_tab001 FORMAT Parquet"
```

- Exporting data in Parquet format

clickhouse client *--host* *Host name or IP address of the ClickHouse instance* *--database* *Database name* *--port* *Port number* **-m --secure --query="select * from Table name FORMAT Parquet" > Parquet file export path**

Example

```
clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="select * from test_table FORMAT Parquet" > /opt/student.parquet
```

- Importing data in ORC format

cat ORC file path | clickhouse client *--host* *Host name or IP address of the ClickHouse instance* *--database* *Database name* *--port* *Port number* **-m --secure --query="INSERT INTO Table name FORMAT ORC"**

Example

```
cat /opt/student.orc | clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="INSERT INTO orc_tab001 FORMAT ORC"
# Data in the ORC file can be exported from HDFS. For example:
hdfs dfs -cat /user/hive/warehouse/hivedb.db/emp_orc/000000_0_copy_1 | clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="INSERT INTO orc_tab001 FORMAT ORC"
```

- Exporting data in ORC format

clickhouse client *--host* *Host name or IP address of the ClickHouse instance* *--database* *Database name* *--port* *Port number* **-m --secure --query="select * from Table name FORMAT ORC" > ORC file export path**

Example

```
clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="select * from csv_tab001 FORMAT ORC" > /opt/student.orc
```

- Importing data in JSON format

INSERT INTO Table name FORMAT JSONEachRow JSON string 1 JSON string 2

Example

```
INSERT INTO test_table001 FORMAT JSONEachRow {"PageViews":5,  
"UserID":"4324182021466249494", "Duration":146,"Sign":-1}  
{"UserID":"4324182021466249494","PageViews":6,"Duration":185,"Sign":1}
```

- Exporting data in JSON format

```
clickhouse client --host Host name or IP address of the ClickHouse instance  
--database Database name --port Port number -m --secure --  
query="SELECT * FROM Table name FORMAT JSON|JSONEachRow|  
JSONCompact|..." > JSON file export path
```

Example

Export JSON file.

```
clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="SELECT *  
FROM test_table FORMAT JSON" > /opt/test.json
```

Export json(JSONEachRow).

```
clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="SELECT *  
FROM test_table FORMAT JSONEachRow" > /opt/test_jsoneachrow.json
```

Export json(JSONCompact).

```
clickhouse client --host 10.5.208.5 --database testdb --port 9440 -m --secure --query="SELECT *  
FROM test_table FORMAT JSONCompact" > /opt/test_jsoncompact.json
```

4.6.2 Synchronizing Kafka Data to ClickHouse

This section describes how to create a Kafka table to automatically synchronize Kafka data to the ClickHouse cluster.

Prerequisites

- You have created a Kafka cluster. The Kafka client has been installed.
- You have created a ClickHouse cluster and installed the ClickHouse client. The ClickHouse and Kafka clusters are in the same VPC and can communicate with each other.

Constraints

Currently, ClickHouse cannot interconnect with Kafka clusters with security mode enabled.

Syntax of the Kafka Table

- **Syntax**

```
CREATE TABLE [IF NOT EXISTS] [db.]table_name [ON CLUSTER cluster]  
(  
  name1 [type1] [DEFAULT|MATERIALIZED|ALIAS expr1],  
  name2 [type2] [DEFAULT|MATERIALIZED|ALIAS expr2],  
  ...  
) ENGINE = Kafka()  
SETTINGS  
  kafka_broker_list = 'host1:port1,host2:port2',  
  kafka_topic_list = 'topic1,topic2,...',  
  kafka_group_name = 'group_name',  
  kafka_format = 'data_format';  
[kafka_row_delimiter = 'delimiter_symbol',]  
[kafka_schema = ''],  
[kafka_num_consumers = N]
```

- **Parameter description**

Table 4-3 Kafka table parameters

Parameter	Mandatory	Description
kafka_broker_list	Yes	A list of Kafka broker instances, separated by comma (,). For example, <i>IP address 1 of the Kafka broker instance:9092,IP address 2 of the Kafka broker instance:9092,IP address 3 of the Kafka broker instance:9092</i> . NOTE If the Kerberos authentication is enabled, parameter allow.everyone.if.no.acl.found must be set to true if port 21005 is used. Otherwise, an error will be reported. To obtain the IP address of the Kafka broker instance, perform the following steps:
kafka_topic_list	Yes	A list of Kafka topics.
kafka_group_name	Yes	A group of Kafka consumers, which can be customized.
kafka_format	Yes	Kafka message format, for example, JSONEachRow, CSV, and XML.
kafka_row_delimiter	No	Delimiter character, which ends a message.
kafka_schema	No	Parameter that must be used if the format requires a schema definition.
kafka_num_consumers	No	Number of consumers in per table. The default value is 1 . If the throughput of a consumer is insufficient, more consumers are required. The total number of consumers cannot exceed the number of partitions in a topic because only one consumer can be allocated to each partition.

How to Synchronize Kafka Data to ClickHouse

Step 1 Switch to the Kafka client installation directory. For details, see [Using the Kafka Client](#).

1. Log in to the node where the Kafka client is installed as the Kafka client installation user.
2. Run the following command to go to the client installation directory:
cd /opt/client
3. Run the following command to configure environment variables:
source bigdata_env

4. If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step.

For an MRS 3.1.0 cluster, run the **export CLICKHOUSE_SECURITY_ENABLED=true** command first.

kinit *Component service user*

- Step 2** Run the following command to create a Kafka topic. For details, see [Managing Kafka Topics](#).

```
kafka-topics.sh --topic kafkacktest2 --create --zookeeper IP address of the Zookeeper role instance:Port used by ZooKeeper to listen to client/kafka --partitions 2 --replication-factor 1
```

 **NOTE**

- **--topic** is the name of the topic to be created, for example, **kafkacktest2**.
- **--zookeeper** is the IP address of the node where the ZooKeeper role instances are located, which can be the IP address of any of the three role instances. You can obtain the IP address of the node by performing the following steps:
- **--partitions** and **--replication-factor** are the topic partitions and topic backup replicas, respectively. The number of the two parameters cannot exceed the number of Kafka role instances.
- To obtain the *Port used by ZooKeeper to listen to client*, log in to FusionInsight Manager, click **Cluster**, choose **Services > ZooKeeper**, and view the value of **clientPort** on the **Configuration** tab page. The default value is **24002**.

- Step 3** Log in to the ClickHouse client by referring to [Using ClickHouse from Scratch](#).

1. Run the following command to go to the client installation directory:

```
cd /opt/client
```

2. Run the following command to configure environment variables:

```
source bigdata_env
```

3. If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. The user must have the permission to create ClickHouse tables. Therefore, you need to bind the corresponding role to the user. For details, see [ClickHouse User and Permission Management](#). If Kerberos authentication is disabled for the current cluster, skip this step.

kinit *Component service user*

Example: **kinit clickhouseuser**

4. Run the following command to connect to the ClickHouse instance node to which data is to be imported:

```
clickhouse client --host IP address of the ClickHouse instance --user Login username --password --port ClickHouse port number --database Database name --multiline
```

Enter the user password.

- Step 4** Create a Kafka table in ClickHouse by referring to [Syntax of the Kafka Table](#). For example, the following table creation statement is used to create a Kafka table whose name is **kafka_src_tbl3**, topic name is **kafkacktest2**, and message format is **JSONEachRow** in the default database.

```
create table kafka_src_tbl3 on cluster default_cluster
(id UInt32, age UInt32, msg String)
ENGINE=Kafka()
SETTINGS
kafka_broker_list='IP address 1 of the Kafka broker instance:9092,IP address 2 of the Kafka broker
instance:9092,IP address 3 of the Kafka broker instance:9092',
kafka_topic_list='kafkacktest2',
kafka_group_name='cg12',
kafka_format='JSONEachRow';
```

Step 5 Create a ClickHouse replicated table, for example, the ReplicatedMergeTree table named **kafka_dest_tbl3**.

```
create table kafka_dest_tbl3 on cluster default_cluster
( id UInt32, age UInt32, msg String )
engine = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/kafka_dest_tbl3', '{replica}')
partition by age
order by id;
```

Step 6 Create a materialized view, which converts data in Kafka in the background and saves the data to the created ClickHouse table.

```
create materialized view consumer3 on cluster default_cluster to kafka_dest_tbl3 as select * from
kafka_src_tbl3;
```

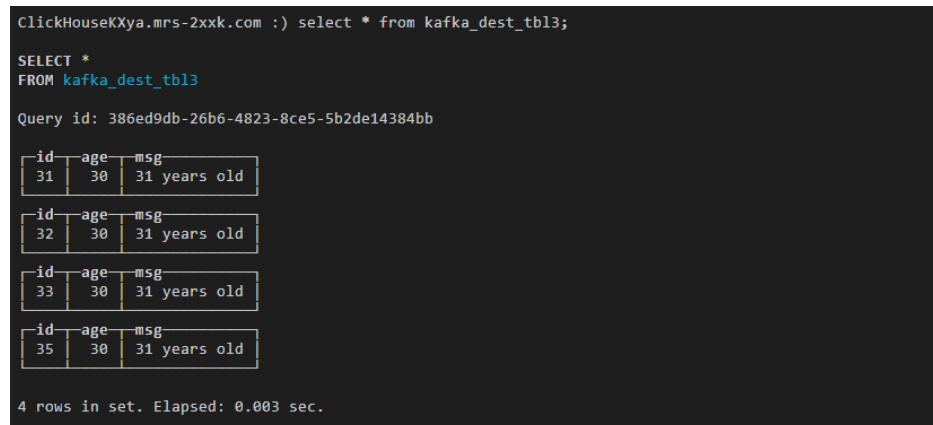
Step 7 Perform [Step 1](#) again to go to the Kafka client installation directory.

Step 8 Run the following command to send a message to the topic created in [Step 2](#):

```
kafka-console-producer.sh --broker-list IP address 1 of the kafka broker
instance:9092,IP address 2 of the kafka broker instance:9092,IP address 3 of the
kafka broker instance:9092 --topic kafkacktest2
>{"id":31, "age":30, "msg":"31 years old"}
>{"id":32, "age":30, "msg":"31 years old"}
>{"id":33, "age":30, "msg":"31 years old"}
>{"id":35, "age":30, "msg":"31 years old"}
```

Step 9 Use the ClickHouse client to log in to the ClickHouse instance node in [Step 3](#) and query the ClickHouse table data, for example, to query the replicated table **kafka_dest_tbl3**. It shows that the data in the Kafka message has been synchronized to this table.

```
select * from kafka_dest_tbl3;
```



```
ClickHouseKXya.mrs-2xxk.com :) select * from kafka_dest_tbl3;
SELECT *
FROM kafka_dest_tbl3
Query id: 386ed9db-26b6-4823-8ce5-5b2de14384bb

 id age msg
---
 31  30 31 years old
 32  30 31 years old
 33  30 31 years old
 35  30 31 years old

4 rows in set. Elapsed: 0.003 sec.
```

----End

4.6.3 Using the ClickHouse Data Migration Tool

The ClickHouse data migration tool can migrate some partitions of one or more partitioned MergeTree tables on several ClickHouseServer nodes to the same

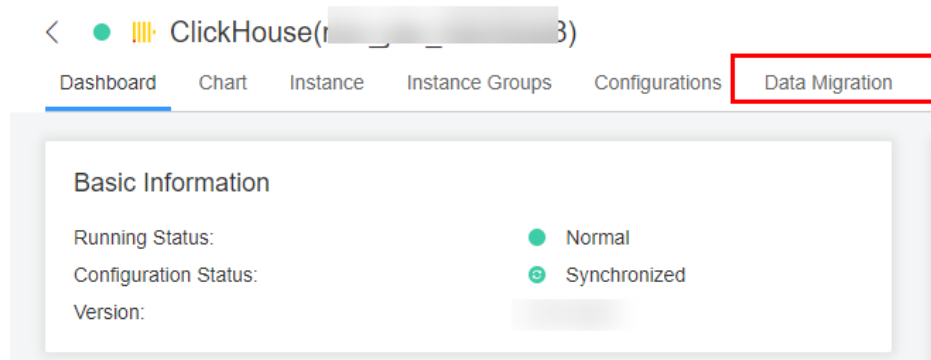
tables on other ClickHouseServer nodes. In the capacity expansion scenario, you can use this tool to migrate data from an original node to a new node to balance data after capacity expansion.

Prerequisites


- The ClickHouse and Zookeeper services are running properly. The ClickHouseServer instances on the source and destination nodes are normal.
- The destination node has the data table to be migrated and the table is a partitioned MergeTree table.
- Before creating a migration task, ensure that all tasks for writing data to a table to be migrated have been stopped. After the task is started, you can only query the table to be migrated and cannot write data to or delete data from the table. Otherwise, data may be inconsistent before and after the migration.
- The ClickHouse data directory on the destination node has sufficient space.

Procedure

Step 1 Log in to Manager and choose **Cluster > Services > ClickHouse**. On the ClickHouse service page, click the **Data Migration** tab.



Step 2 Click **Add Task**.



Step 3 On the page for creating a migration task, set the migration task parameters. For details, see [Table 4-4](#).





Table 4-4 Migration task parameters

Parameter	Description
Task Name	Enter a specific task name. The value can contain 1 to 50 characters, including letters, arrays, and underscores (_), and cannot be the same as that of an existing migration task.
Task Type	<ul style="list-style-type: none"> • Scheduled Task: When the scheduled task is selected, you can set Started to specify a time point later than the current time to execute the task. • Immediate task: The task is executed immediately after it is started.
Started	Set this parameter when Task Type is set to Scheduled Task . The valid value is a time point within 90 days from now.

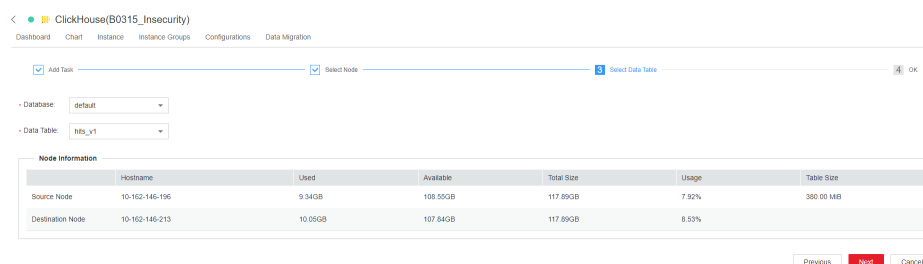
Step 4 On the **Select Node** page, specify **Source Node Host Name** and **Destination Node Host Name**, and click **Next**.



NOTE

- Only one host name can be entered in **Source Node Host Name** and **Destination Node Host Name**, respectively. Multi-node migration is not supported.
To obtain the parameter values, click the **Instance** tab on the ClickHouse service page and view the **Host Name** column of the current ClickHouseServer instance.
- **Maximum Bandwidth** is optional. If it is not specified, there is no upper limit. The maximum bandwidth can be set to **10000 MB/s**.


Step 5 On the **Select Data Table** page, click **Database**, select the database to be migrated on the source node, and select the data table to be migrated for **Data Table**. The data table drop-down list displays the partitioned MergeTree tables in the selected database. In the **Node Information** area, the space usage of the ClickHouse service data directory on the current source and destination nodes is displayed. Click **Next**.



Step 6 Confirm the task information and click **Submit**.

The data migration tool automatically calculates the partitions to be migrated based on the size of the data table. The amount of data to be migrated is the total size of the partitions to be migrated.

Step 7 After the migration task is submitted, click **Start** in the **Operation** column. If the task is an immediate task, the task starts to be executed. If the task is a scheduled task, the countdown starts.



Step 8 During the migration task execution, you can click **Cancel** to cancel the migration task that is being executed. If you cancel the task, the migrated data on the destination node will be rolled back.

You can choose **More > Details** to view the log information during the migration.

Step 9 After the migration is complete, choose **More > Results** to view the migration result and choose **More > Delete** to delete the directories related to the migration task on ZooKeeper and the source node.

----End

4.7 User Management and Authentication

4.7.1 ClickHouse User and Permission Management

User Permission Model

ClickHouse user permission management enables unified management of users, roles, and permissions on each ClickHouse instance in the cluster. You can use the permission management module of the Manager UI to create users, create roles, and bind the ClickHouse access permissions. User permissions are controlled by binding roles to users.

Resource management: [Table 4-5](#) lists the resources supported by ClickHouse permission management.

Resource permissions: [Table 4-6](#) lists the resource permissions supported by ClickHouse.

Table 4-5 Permission management objects supported by ClickHouse

Resource	Integration	Remarks
Database	Yes (level 1)	-
Table	Yes (level 2)	-
View	Yes (level 2)	Same as tables

Table 4-6 Resource permission list

Resource	Available Permission	Remarks
Database	CREATE	CREATE DATABASE/TABLE/VIEW/ DICTIONARY
Table/View	SELECT/INSERT	-

Prerequisites

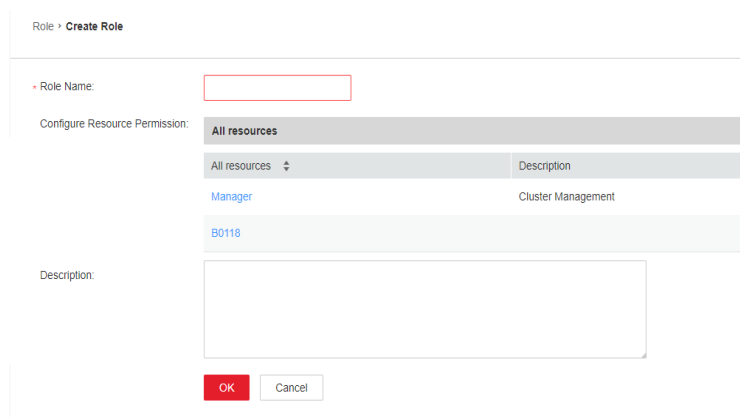
- The ClickHouse and Zookeeper services are running properly.
- The **ON CLUSTER** statement has been used to create a database or table in the cluster to ensure that the metadata of the database and table on each ClickHouse node is the same.

NOTE

After the permission is granted, it takes about 1 minute for the permission to take effect.

Adding the ClickHouse Role

Step 1 Log in to Manager and choose **System > Permission > Role**. On the **Role** page, click **Create Role**.



Step 2 On the **Create Role** page, specify **Role Name**. In the **Configure Resource Permission** area, click the cluster name. On the service list page that is displayed, click the ClickHouse service.

Determine whether to create a role with the ClickHouse administrator permissions based on service requirements.

NOTE

- The ClickHouse administrator has all the database operation permissions except the permissions to create, delete, and modify users and roles.
- Only the built-in user **clickhouse** of ClickHouse has the permission to manage users and roles.
- If yes, go to **Step 3**.

- If no, go to **Step 4**.

Role > **Create Role**

* Role Name:

Configure Resource Permission: **All resources > B0118 > ClickHouse**

View Name

SUPER_USER_GROUP

Clickhouse Scope

Description:

Step 3 Select **SUPER_USER_GROUP** and click **OK**.

Step 4 Click **Clickhouse Scope**. The ClickHouse database resource list is displayed. If you select **create**, the role has the create permission on the database.

Role > **Create Role**

* Role Name:

Configure Resource Permission: **All resources > B0118 > Clickhouse > Clickhouse Scope**

Resource Name	Resource Type	Permission
_temporary_and_external_tables	Database	<input type="checkbox"/> create
db1	Database	<input checked="" type="checkbox"/> <input type="checkbox"/>
db10	Database	<input checked="" type="checkbox"/> <input type="checkbox"/>
db2	Database	<input checked="" type="checkbox"/> <input type="checkbox"/>
db3	Database	<input type="checkbox"/>
db4	Database	<input type="checkbox"/>
db5	Database	<input type="checkbox"/>
db6	Database	<input type="checkbox"/>
db7	Database	<input type="checkbox"/>
db8	Database	<input type="checkbox"/>

Determine whether to grant the permission based on the service requirements.

- If yes, click **OK**.
- If no, go to **Step 5**.

Step 5 Click the resource name and select the *Database resource name to be operated*. On the displayed page, select **READ** (SELECT permission) or **WRITE** (INSERT permission) based on service requirements, and click **OK**.

Role > Create Role

Role Name:

Configure Resource Permission: All resources > B0118 > ClickHouse > Clickhouse Scope > db2

Resource Name	Resource Type	Permission	
		read	write
tb3	Table	<input type="checkbox"/>	<input checked="" type="checkbox"/>
tb4	Table	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Description:

----End

Adding a User and Binding the ClickHouse Role to the User

Step 1 Log in to Manager and choose **System > Permission > User** and click **Create**.

Step 2 Select **Human-Machine** for **User Type** and set **Password** and **Confirm Password** to the password of the user.

NOTE

- Username: The username cannot contain hyphens (-). Otherwise, the authentication will fail.
- Password: The password cannot contain special characters \$, , and #. Otherwise, the authentication will fail.

Step 3 In the **Role** area, click **Add**. In the displayed dialog box, select a role with the ClickHouse permission and click **OK** to add the role. Then, click **OK**.

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:

Step 4 Log in to the node where the ClickHouse client is installed and use the new username and password to connect to the ClickHouse service.

- Run the following command to go to the client installation directory:
cd /opt/Client installation directory
- Run the following command to configure environment variables:
source bigdata_env
- If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. The user must have the

permission to create ClickHouse tables. Therefore, you need to bind the corresponding role to the user. For details, see [Adding the ClickHouse Role](#). If Kerberos authentication is disabled for the current cluster, skip this step.

For an MRS 3.1.0 cluster, run the **export CLICKHOUSE_SECURITY_ENABLED=true** command first.

kinit *User added in [Step 1](#)*

- Log in to the system as the new user.

Cluster with Kerberos authentication disabled:

```
clickhouse client --host IP address of the ClickHouse instance --multiline --port ClickHouse port number --secure
```

Cluster with Kerberos authentication enabled:

```
clickhouse client --host IP address of the ClickHouse instance --user Username --password --port 9440 --secure
```

Enter the user password.

 **NOTE**

The user in normal mode is the default user, or you can create an administrator using the open source capability provided by the ClickHouse community. You cannot use the users created on FusionInsight Manager.

----End

Granting Permissions Using the Client in Abnormal Scenarios

By default, the table metadata on each node of the ClickHouse cluster is the same. Therefore, the table information on a random ClickHouse node is collected on the permission management page of Manager. If the **ON CLUSTER** statement is not used when databases or tables are created on some nodes, the resource may fail to be displayed during permission management, and permissions may not be granted to the resource. To grant permissions on the local table on a single ClickHouse node, perform the following steps on the background client.

 **NOTE**

The following operations are performed based on the obtained roles, database or table names, and IP addresses of the node where the corresponding ClickHouseServer instance is located.

- You can log in to FusionInsight Manager and choose **Cluster > Services > ClickHouse > Instance** to obtain the service IP address of the ClickHouseServer instance.
- The default system domain name is **hadoop.com**. Log in to FusionInsight Manager and choose **System > Permission > Domain and Mutual Trust**. The value of **Local Domain** is the system domain name. Change the letters to lowercase letters when running a command.

Step 1 Log in to the node where the ClickHouseServer instance is located as user **root**.

Step 2 Run the following command to obtain the path of the **clickhouse.keytab** file:

```
ls `${BIGDATA_HOME}/FusionInsight_ClickHouse_*/install/FusionInsight-ClickHouse-*/clickhouse/keytab/clickhouse.keytab
```

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

Step 4 Run the following command to go to the client installation directory:

```
cd /opt/client
```

Step 5 Run the following command to configure environment variables:

```
source bigdata_env
```

For an MRS 3.1.0 cluster with Kerberos authentication enabled, additionally run the following command:

```
export CLICKHOUSE_SECURITY_ENABLED=true
```

Step 6 Run the following command to connect to the ClickHouseServer instance:

If Kerberos authentication is enabled for the current cluster, run the following command:

```
clickhouse client --host IP address of the node where the ClickHouseServer instance is located --user clickhouse/hadoop.<System domain name> --password clickhouse.keytab path obtained in Step 2 --port ClickHouse port number --secure
```

If Kerberos authentication is disabled for the current cluster, run the following command:

```
clickhouse client --host IP address of the node where the ClickHouseServer instance is located --user clickhouse --port ClickHouse port number
```

Step 7 Run the following statement to grant permissions to a database:

In the syntax for granting permissions, *DATABASE* indicates the name of the target database, and *role* indicates the target role.

```
GRANT [ON CLUSTER cluster_name] privilege ON {DATABASE|TABLE} TO {user | role}
```

For example, grant user **testuser** the CREATE permission on database **t2**:

```
GRANT CREATE ON m2 to testuser;
```

Step 8 Run the following commands to grant permissions on the table or view. In the following command, *TABLE* indicates the name of the table or view to be operated, and *user* indicates the role to be operated.

Run the following command to grant the query permission on tables in a database:

```
GRANT SELECT ON TABLE TO user;
```

Run the following command to grant the write permission on tables in a database:

```
GRANT INSERT ON TABLE TO user;
```

Step 9 Run the following command to exit the client:

```
quit;
```

```
----End
```


4.7.2 Interconnecting ClickHouse With OpenLDAP for Authentication

ClickHouse can be interconnected with OpenLDAP. You can manage accounts and permissions in a centralized manner by adding the OpenLDAP server configuration and creating users on ClickHouse. You can use this method to import users from the OpenLDAP server to ClickHouse in batches.

This section applies only to MRS 3.1.0 or later.

Prerequisites

- The MRS cluster and ClickHouse instances are running properly, and the ClickHouse client has been installed.
- OpenLDAP has been installed and is running properly.

Creating a ClickHouse User for Interconnecting with the OpenLDAP Server

Step 1 Log in to Manager and choose **Cluster > Services > ClickHouse**. Click the **Configurations** tab and then **All Configurations**.

Step 2 Choose **ClickHouseServer(Role) > Customization**, and add the following OpenLDAP configuration parameters to the **clickhouse-config-customize** configuration item.

Table 4-7 OpenLDAP parameters

Parameter	Description	Example Value
ldap_servers.ldap_server_name.host	OpenLDAP server host name or IP address. This parameter cannot be empty.	localhost
ldap_servers.ldap_server_name.port	OpenLDAP server port number. If enable_tls is set to true , the default port number is 636 . Otherwise, the default port number is 389 .	636
ldap_servers.ldap_server_name.auth_dn_prefix	Prefix and suffix used to construct the DN to bind to.	uid=
ldap_servers.ldap_server_name.auth_dn_suffix	The generated DN will be constructed as a string in the following format: auth_dn_prefix + escape(user_name) + auth_dn_suffix . Use a comma (,) as the first non-space character of auth_dn_suffix .	,ou=Group,dc=node1,dc=com

Parameter	Description	Example Value
ldap_servers.ldap_server_name.enable_tls	A tag to trigger the use of the secure connection to the OpenLDAP server. <ul style="list-style-type: none"> Set it to no for the plaintext (ldap://) protocol (not recommended). Set it to yes for the LDAP over SSL/TLS (ldaps://) protocol. 	yes
ldap_servers.ldap_server_name.tls_require_cert	SSL/TLS peer certificate verification behavior. The value can be never , allow , try , or require .	allow

 NOTE

For details about other parameters, see [<ldap_servers> Parameters](#).

- Step 3** After the configuration is complete, click **Save**. In the displayed dialog box, click **OK**. After the configuration is saved, click **Finish**.
- Step 4** On Manager, click **Instance**, select a ClickHouseServer instance, and choose **More > Restart Instance**. In the displayed dialog box, enter the password and click **OK**. In the displayed **Restart instance** dialog box, click **OK**. Confirm that the instance is restarted successfully as prompted and click **Finish**.
- Step 5** Log in to the ClickHouseServer instance node and go to the `/${BIGDATA_HOME}/FusionInsight_ClickHouse_Version number/x_x_ClickHouseServer/etc` directory.
- Step 6** Run the following command to view the `config.xml` configuration file and check whether the OpenLDAP parameters are configured successfully:

cat config.xml

```
[root@k 3 etc]# cat config.xml
<vindex>
<ldap_servers>
  <ldap_server_name>
    <auth_dn_prefix>uid=</auth_dn_prefix>
    <port>636</port>
    <host>localhost</host>
    <enable_tls>yes</enable_tls>
    <tls_require_cert>allow</tls_require_cert>
    <auth_dn_suffix>,ou=Group,dc=node1,dc=com</auth_dn_suffix>
  </ldap_server_name>
</ldap_servers>
<zookeeper> ...
```

- Step 7** Log in to the node where the ClickHouseServer instance is located as user **root**.
- Step 8** Run the following command to obtain the path of the `clickhouse.keytab` file:

```
ls ${BIGDATA_HOME}/FusionInsight_ClickHouse_*/install/FusionInsight-ClickHouse-*/clickhouse/keytab/clickhouse.keytab
```

Step 9 Log in to the node where the client is installed as the client installation user.

Step 10 Run the following command to go to the ClickHouse client installation directory:

```
cd /opt/client
```

Step 11 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 12 Run the following command to connect to the ClickHouseServer instance:

- If Kerberos authentication is enabled for the current cluster, use **clickhouse.keytab** to connect to the ClickHouseServer instance.

```
clickhouse client --host IP address of the node where the ClickHouseServer instance is located --user clickhouse/hadoop.<System domain name> --password clickhouse.keytab path obtained in Step 8 --port ClickHouse port number
```

 **NOTE**

The default system domain name is **hadoop.com**. Log in to FusionInsight Manager and choose **System > Permission > Domain and Mutual Trust**. The value of **Local Domain** is the system domain name. Change the letters to lowercase letters when running a command.

- If Kerberos authentication is not enabled for the current cluster, connect to the ClickHouseServer instance as the ClickHouse administrator.

```
clickhouse client --host IP address of the node where the ClickHouseServer instance is located --user clickhouse --port ClickHouse port number
```

Step 13 Create a common user of OpenLDAP.

Run the following statement to create user **testUser** in cluster **default_cluster** and set **ldap_server** to the OpenLDAP server name in the **<ldap_servers>** tag in [Step 6](#). In this example, the name is **ldap_server_name**.

```
CREATE USER testUser ON CLUSTER default_cluster IDENTIFIED WITH ldap_server BY 'ldap_server_name';
```

testUser indicates an existing username in OpenLDAP. Change it based on the site requirements.

Step 14 Log out of the client, and then log in to the client as the new user to check whether the configuration is successful.

```
exit;
```

```
clickhouse client --host IP address of the ClickHouseServer instance --user testUser --password --port ClickHouse port number
```

Enter the password of testUser.

```
----End
```

<ldap_servers> Parameters

- **host**
OpenLDAP server host name or IP address. This parameter is mandatory and cannot be empty.
- **port**
Port number of the OpenLDAP server. If **enable_tls** is set to **true**, the default value is **636**. Otherwise, the value is **389**.
- **auth_dn_prefix, auth_dn_suffix**
Prefix and suffix used to construct the DN to bind to.
The generated DN will be constructed as a string in the following format: **auth_dn_prefix + escape(user_name) + auth_dn_suffix**.
Note that you should use a comma (,) as the first non-space character of **auth_dn_suffix**.
- **enable_tls**
A tag to trigger the use of the secure connection to the OpenLDAP server. Set it to **no** for the plaintext (ldap://) protocol (not recommended). Set it to **yes** for LDAP over SSL/TLS (ldaps://) protocol (recommended and default).
- **tls_minimum_protocol_version**
Minimum protocol version of SSL/TLS. The value can be **ssl2, ssl3, tls1.0, tls1.1, or tls1.2** (default).
- **tls_require_cert**
SSL/TLS peer certificate verification behavior. The value can be **never, allow, try, or require** (default).
- **tls_cert_file**
Certificate file.
- **tls_key_file**
Certificate key file.
- **tls_ca_cert_file**
CA certificate file.
- **tls_ca_cert_dir**
Directory where the CA certificate is stored.
- **tls_cipher_suite**
Allowed encryption suite.

4.8 Backing Up and Restoring ClickHouse Data Using a Data File

Scenario

This section describes how to back up data by exporting ClickHouse data to a CSV file and restore data using the CSV file.

Prerequisites

- You have installed the ClickHouse client.
- You have created a user with related permissions on ClickHouse tables on Manager.
- You have prepared a server for backup.

Backing Up Data

Step 1 Log in to the node where the client is installed as the client installation user.

Step 2 Run the following command to go to the client installation directory:

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create ClickHouse tables. If Kerberos authentication is disabled for the current cluster, skip this step.

1. Run the following command if it is an MRS 3.1.0 cluster:

```
export CLICKHOUSE_SECURITY_ENABLED=true
```

2. **kinit** *Component service user*

Example: **kinit clickhouseuser**

Step 5 Run the ClickHouse client command to export the ClickHouse table data to be backed up to a specified directory.

```
clickhouse client --host Host name/Instance IP address --secure --port 9440 --query="Table query statement" > Path of the exported CSV file
```

The following shows an example of backing up data in the **test** table to the **default_test.csv** file on the ClickHouse instance **10.244.225.167**.

```
clickhouse client --host 10.244.225.167 --secure --port 9440 --query="select * from default.test FORMAT CSV" > /opt/clickhouse/default_test.csv
```

Step 6 Upload the exported CSV file to the backup server.

```
----End
```

Restoring Data

Step 1 Upload the backup data file on the backup server to the directory where the ClickHouse client is located.

For example, upload the **default_test.csv** backup file to the **/opt/clickhouse** directory.

Step 2 Log in to the node where the client is installed as the client installation user.

Step 3 Run the following command to go to the client installation directory:

```
cd /opt/client
```

Step 4 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 5 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create ClickHouse tables. If Kerberos authentication is disabled for the current cluster, skip this step.

1. Run the following command if it is an MRS 3.1.0 cluster:

```
export CLICKHOUSE_SECURITY_ENABLED=true
```

2. **kinit** *Component service user*

Example: **kinit clickhouseuser**

Step 6 Run the ClickHouse client command to log in to the ClickHouse cluster.

```
clickhouse client --host Host name/Instance IP address --secure --port 9440
```

Step 7 Create a table with the format corresponding to the CSV file.

```
CREATE TABLE [IF NOT EXISTS] [database_name.]table_name [ON CLUSTER  
Cluster name]  
  
(  
  
  name1 [type1] [DEFAULT|materialized|ALIAS expr1],  
  name2 [type2] [DEFAULT|materialized|ALIAS expr2],  
  ...  
)  
ENGINE = engine
```

Step 8 Import the content in the backup file to the table created in [Step 7](#) to restore data.

```
clickhouse client --host Host name/Instance IP address --secure --port 9440 --  
query="insert into Table name FORMAT CSV" < CSV file path
```

The following shows an example of restoring data from the **default_test.csv** backup file to the **test_cpy** table on the ClickHouse instance **10.244.225.167**.

```
clickhouse client --host 10.244.225.167 --secure --port 9440 --query="insert  
into default.test_cpy FORMAT CSV" < /opt/clickhouse/default_test.csv
```

----End

4.9 ClickHouse Log Overview

Log Description

Log path: The default storage path of ClickHouse log files is as follows: **\$ {BIGDATA_LOG_HOME}/clickhouse**

Log archive rule: The automatic log compression and archiving function are enabled. By default, when the size of logs exceeds 100 MB, logs are automatically compressed into a log file named in the format of *<Original log file name>*.

[No.].gz. A maximum of 10 latest compressed files are reserved by default. The number of compressed files can be configured on Manager.

Table 4-8 ClickHouse log list

Log Type	Log File Name	Description
Run log	/var/log/Bigdata/clickhouse/clickhouseServer/clickhouse-server.err.log	Path of ClickHouseServer error log files
	/var/log/Bigdata/clickhouse/clickhouseServer/checkService.log	Path of key ClickHouseServer run log files
	/var/log/Bigdata/clickhouse/clickhouseServer/clickhouse-server.log	
	/var/log/Bigdata/clickhouse/balance/start.log	Path of ClickHouseBalancer startup log files
	/var/log/Bigdata/clickhouse/balance/error.log	Path of ClickHouseBalancer error log files
	/var/log/Bigdata/clickhouse/balance/access_http.log	Path of ClickHouseBalancer run log files
Data migration log	/var/log/Bigdata/clickhouse/migration/ <i>Data migration task name</i> /clickhouse-copier_{timestamp}_{processId}/copier.log	Run log generated when you use the migration tool by referring to Using the ClickHouse Data Migration Tool
	/var/log/Bigdata/clickhouse/migration/ <i>Data migration task name</i> /clickhouse-copier_{timestamp}_{processId}/copier.err.log	Error log generated when you use the migration tool by referring to Using the ClickHouse Data Migration Tool

Log Level

Table 4-9 describes the log levels supported by ClickHouse.

Levels of run logs are error, warning, trace, information, and debug from the highest to the lowest priority. Run logs of equal or higher levels are recorded. The higher the specified log level, the fewer the logs recorded.

Table 4-9 Log levels

Level	Description
error	Logs of this level record error information about system running
warning	Logs of this level record exception information about the current event processing

Level	Description
trace	Logs of this level record trace information about the current event processing
information	Logs of this level record normal running status information about the system and events
debug	Logs of this level record system running and debugging information

To modify log levels, perform the following operations:

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster > Services > ClickHouse > Configurations**.
- Step 3** Select **All Configurations**.
- Step 4** On the menu bar on the left, select the log menu of the target role.
- Step 5** Select a desired log level.
- Step 6** Click **Save**. Then, click **OK**.

----End

 **NOTE**

The configurations take effect immediately without the need to restart the service.

Log Format

The following table lists the ClickHouse log format:

Table 4-10 Log formats

Log Type	Format	Example
ClickHouse run log	<i><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></i>	2021.02.23 15:26:30.691301 [6085] {} <Error> DynamicQueryHandler: Code: 516, e.displayText() = DB::Exception: default: Authentication failed: password is incorrect or there is no user with such name, Stack trace (when copying this message, always include the lines below): 0. Poco::Exception::Exception(std::__1::basic_string<char, std::__1::char_traits<char>, std::__1::allocator<char> > const&, int) @ 0x1250e59c

4.10 ClickHouse Performance Tuning

4.10.1 Solution to the "Too many parts" Error in Data Tables

Troubleshooting

1. Log in to the ClickHouse client and check whether abnormal merge exists.
select database, table, elapsed, progress, merge_type from system.merges;
2. Do not perform the INSERT operation too frequently, do not insert a small amount of data, and increase the interval for inserting data.
3. The data table partitions are not properly allocated. As a result, too many partitions are generated and data tables need to be re-partitioned.
4. If the MERGE operations are not triggered or slow, adjust the following parameters to accelerate them.

For details, see [Accelerating Merge Operations](#).

Configuration Item	Reference Value
max_threads	Number of CPU cores x 2
background_pool_size	Number of CPU cores
merge_max_block_size	The value is an integer multiple of 8192 and is adjusted based on the CPU and memory resources.
cleanup_delay_period	Set this parameter to a value that is appropriately less than the default value 30.

Changing the Value of parts_to_throw_insert

CAUTION

Increase the value of this parameter only in special scenarios. This configuration acts as a warning for potential issues to some extent. If the cluster hardware resources are insufficient and this configuration is not adjusted properly, potential service issues cannot be detected in a timely manner, which may cause other faults and increase the difficulty of fault recovery.

Log in to FusionInsight Manager and choose **Cluster > Services > ClickHouse**. Click **Configurations** then **All Configurations**. Click **ClickHouseServer > Customization**, find the **clickhouse-config-customize** parameter, add the following configuration, save it, and restart the service.

Name	Value
merge_tree.parts_to_throw_insert	Memory of ClickHouse instances/32 GB x 300 (conservative estimation)

Verify the modification.

Log in to the ClickHouse client and run the **select * from system.merge_tree_settings where name = 'parts_to_throw_insert'**; command.

4.10.2 Accelerating Merge Operations

To accelerate background tasks, adjust the ZooKeeper service configuration first. Otherwise, the ClickHouse service and background tasks will be abnormal due to insufficient ZooKeeper resources such as znodes.

1. Adjust the ZooKeeper configuration. Log in to FusionInsight Manager and choose **Cluster > Services > Zookeeper**. Click **Configurations** then **All Configurations**. Click **quorumpeer > System**, change the value of **GC_OPTS** according to the following table, save the configuration, and roll restart the ZooKeeper service.

Configuration Item	Reference Value	Description
GC_OPTS	Xmx = (Memory size of master nodes - 16 GB) x 0.65 (conservative estimation)	JVM parameter used for garbage collection (GC). This parameter is valid only when GC_PROFILE is set to custom . Ensure that the GC_OPT parameter is set correctly. Otherwise, the process will fail to be started. CAUTION Exercise caution when modifying this item. If this parameter is set incorrectly, the service will be unavailable.

2. Adjust the ClickHouse configuration. On FusionInsight Manager, choose **Cluster > Services > ClickHouse**. Click **Configurations** then **All Configurations**. Click **ClickHouse > Zookeeper**, modify the following parameters, and save the configuration. You do not need to restart the service.

Configuration Item	Reference Value	Description
clickhouse.zookeeper.quota.node.conut	Xmx/4 GB x 1,500,000	<p>Node quantity quota of ClickHouse in the top directory on ZooKeeper.</p> <p>This parameter cannot be set to 0, but can be set to a minimum value of -1. -1 indicates that there is no limit on the number of ClickHouse nodes in the top directory.</p> <p>CAUTION If the quantity quota is less than the actual value of the current ZooKeeper directory, the configuration can be saved but does not take effect and an alarm is reported on the GUI.</p>
clickhouse.zookeeper.quota.size	Xmx/4 GB x 1 GB	<p>Capacity quota of ClickHouse in the top directory on ZooKeeper.</p> <p>CAUTION If the quantity quota is less than the actual value of the current ZooKeeper directory, the configuration can be saved but does not take effect and an alarm is reported on the GUI.</p>

4.10.3 Accelerating TTL Operations

When TTL is triggered in ClickHouse, a large amount of CPU and memory are consumed.

Log in to FusionInsight Manager and choose **Cluster > Services > ClickHouse**. Click **Configurations** then **All Configurations**. Click **ClickHouseServer > Customization**, find the **clickhouse-config-customize** parameter, add the following parameters, save the configuration, and restart the service.

Configuration Item	Reference Value	Description
merge_tree.max_replicated_merges_with_ttl_in_queue	Half of number of CPU cores	Number of tasks that allow TTL to merge parts concurrently in the ReplicatedMergeTree queue.
merge_tree.max_number_of_merges_with_ttl_in_pool	Number of CPU cores	The thread pool that allows TTL to merge parts in the ReplicatedMergeTree queue.

 NOTE

Do not modify these configurations when the cluster writes heavily. Idle threads need to be reserved for regular Merge operations to avoid the "Too many parts" issue.

4.11 ClickHouse FAQ

4.11.1 How Do I Do If the Disk Status Displayed in the System.disks Table Is fault or abnormal?

Question

How do I do if the disk status displayed in the System.disks table is fault or abnormal?

Answer

This problem is caused by I/O errors on the disk. To rectify the fault, perform the following steps:

- Method 1: Log in to FusionInsight Manager and check whether an alarm is generated indicating that the disk I/O is abnormal. If yes, replace the faulty disk by referring to the alarm help.
- Method 2: Log in to FusionInsight Manager and restart the ClickHouse instance to restore the disk status.

 NOTE

If an I/O error occurs but the disk is not replaced, the disk status will still turn to fault or abnormal.

4.11.2 How Do I Migrate Data from Hive/HDFS to ClickHouse?

Question

How do I migrate Hive/HDFS data to ClickHouse?

Answer

You can export data from Hive as CSV files and import the CSV files to ClickHouse.

1. Export data from Hive as CSV files.

```
hive -e "select * from db_hive.student limit 1000" | tr "\t" "," > /data/bigdata/hive/student.csv;
```

2. Import the CSV files to the **student_hive** table in the default database of ClickHouse.

```
clickhouse --client --port 9002 --password password -m --query='INSERT INTO default.student_hive FORMAT CSV' < /data/bigdata/hive/student.csv
```


4.11.4 An Error Is Reported in Logs When the Auxiliary ZooKeeper or Replica Data Is Used to Synchronize Table Data

Question

An error is reported in logs when the auxiliary ZooKeeper or replica data is used to synchronize table data.

```
DB::Exception: Cannot parse input: expected 'quorum:' before: 'merge_type: 2'...' and "Too many parts (315). Merges are processing significantly slower than inserts..."
```

Answer

The versions of replication table replicas are inconsistent, causing compatibility issues. The table schema contains TTL statements. TTL_DELETE is added in versions later than ClickHouse 20.9, which cannot be identified in earlier versions. This issue occurs when the replication table replica of a later version is elected as the leader.

You can modify the **config.xml** file of ClickHouse of a later version to avoid such an issue. Ensure that the replication table replicas are the same as those of ClickHouse.

4.11.5 How Do I Grant the Select Permission at the Database Level to ClickHouse Users?

Procedure

Step 1 Log in to the node where the ClickHouse client is installed in the MRS cluster and run the following commands:

```
su - omm
```

```
source {Client installation directory}/bigdata_env
```

```
kinit Component user (You do not need to run the kinit command for normal clusters.)
```

```
clickhouse client --host IP address of the ClickHouse node --port 9000 -m --user clickhouse --password 'Password of the ClickHouse user'
```

NOTE

View the password of the ClickHouse user.

Log in to FusionInsight Manager and choose **Cluster > Services > ClickHouse**. Click **Instance** and click any ClickHouseServer role name. Go to the **Dashboard** tab page of ClickHouseServer, click the **users.xml** file in **Configuration File** area, and view the password of the ClickHouse user.

Step 2 You can use either of the following methods to create a role with the read-only permission for a specified database:

Method 1

1. Creating a role with the read-only permission for a specified database (the **default** database is used as an example)

```
create role ck_role on cluster default_cluster;  
GRANT SELECT ON default.* TO ck_role on cluster default_cluster;
```

2. Creating a common user

```
CREATE USER user_01 on cluster default_cluster IDENTIFIED WITH  
PLAINTEXT_PASSWORD BY 'password';
```

3. Granting the read-only permission role to a common user

```
GRANT ck_role to user_01 on cluster default_cluster;
```

4. Viewing user permissions

```
show grants for user_01;  
select * from system.grants where role_name = 'ck_role';
```

Method 2

Creating a user with the read-only permission for a specified database

1. Creating a user:

```
CREATE USER user_01 on cluster default_cluster IDENTIFIED WITH  
PLAINTEXT_PASSWORD BY 'password';
```

2. Granting the query permission on a specified database to the created user:

```
grant select on default.* to user_01 on cluster default_cluster;
```

3. Querying user permissions:

```
select * from system.grants where user_name = 'user_01';
```

----End

5 Using DBService

5.1 DBService Log Overview

Log Description

Log path: The default storage path of DBService log files is `/var/log/Bigdata/dbservice`.

- GaussDB: `/var/log/Bigdata/dbservice/DB` (GaussDB run log directory), `/var/log/Bigdata/dbservice/scriptlog/gaussdbinstall.log` (GaussDB installation log), and `/var/log/gaussdbuninstall.log` (GaussDB uninstallation log).
- HA: `/var/log/Bigdata/dbservice/ha/runlog` (HA run log directory) and `/var/log/Bigdata/dbservice/ha/scriptlog` (HA script log directory)
- DBServer: `/var/log/Bigdata/dbservice/healthCheck` (Directory of service and process health check logs)
`/var/log/Bigdata/dbservice/scriptlog` (run log directory), `/var/log/Bigdata/audit/dbservice/` (audit log directory)

Log archive rule: The automatic DBService log compression function is enabled. By default, when the size of logs exceeds 1 MB, logs are automatically compressed into a log file named in the following format: `<Original log file name>-[No.].gz`. A maximum of 20 latest compressed files are reserved.

NOTE

Log archive rules cannot be modified.

Table 5-1 DBService log list

Type	Log File Name	Description
DBServer run log	dbservice_serviceCheck.log	Run log file of the service check script

Type	Log File Name	Description
	dbservice_processCheck.log	Run log file of the process check script
	backup.log	Run logs of backup and restoration operations (The DBService backup and restoration operations need to be performed.)
	checkHaStatus.log	Log file of HA check records
	cleanupDBService.log	Uninstallation log file (You need to uninstall DBService logs.)
	componentUserManager.log	Log file that records the adding and deleting operations on the database by users (Services that depend on DBService need to be added.)
	install.log	Installation log file
	preStartDBService.log	Pre-startup log file
	start_dbserver.log	DBServer startup operation log file (DBService needs to be started.)
	stop_dbserver.log	DBServer stop operation log file (DBService needs to be stopped.)
	status_dbserver.log	Log file of the DBServer status check (You need to execute the \$DBSERVICE_HOME/sbin/status-dbserver.sh script.)
	modifyPassword.log	Run log file of changing the DBService password script. (You need to execute the \$DBSERVICE_HOME/sbin/modifyDBPwd.sh script.)

Type	Log File Name	Description
	modifyDBPwd_YYYY-MM-DD.log	Run log file that records the DBService password change tool (You need to execute the \$DBSERVICE_HOME/sbin/modifyDBPwd.sh script.)
	dbserver_switchover.log	Log for DBServer to execute the active/standby switchover script (the active/standby switchover needs to be performed)
GaussDB run log	gaussdb.log	Log file that records database running information
	gs_ctl-current.log	Log file that records operations performed by using the gs_ctl tool
	gs_guc-current.log	Log file that records operations, mainly parameter modification performed by using the gs_guc tool
	gaussdbinstall.log	GaussDB installation log file
	gaussdbuninstall.log	GaussDB uninstallation log file
HA script run log	floatip_ha.log	Log file that records the script of floating IP addresses
	gaussDB_ha.log	Log file that records the script of GaussDB resources
	ha_monitor.log	Log file that records the HA process monitoring information
	send_alarm.log	Alarm sending log file
	ha.log	HA run log file

Type	Log File Name	Description
DBService audit log	dbservice_audit.log	Audit log file that records DBService operations, such as backup and restoration operations

Log Format

The following table lists the DBService log formats.

Table 5-2 Log format

Type	Format	Example
Run log	[<yyyy-MM-dd HH:mm:ss> <Log level>: [< Name of the script that generates the log. Line number >]: < Message in the log>	[2020-12-19 15:56:42] INFO [postinstall.sh:653] Is cloud flag is false. (main)
Audit log	[<yyyy-MM-dd HH:mm:ss,SSS> UserName:<Username> UserIP:<User IP address> Operation:<Operation content> Result:<Operation results> Detail:<Detailed information>	[2020-05-26 22:00:23] UserName:omm UserIP:192.168.10.21 Operation:DBService data backup Result: SUCCESS Detail: DBService data backup is successful.

6 Using Flink

6.1 Using Flink from Scratch

This section describes how to use Flink to run wordcount jobs.

Prerequisites

- Flink has been installed in an MRS cluster.
- The cluster runs properly and the client has been correctly installed, for example, in the **/opt/hadoopclient** directory. The client directory in the following operations is only an example. Change it to the actual installation directory.

Using the Flink Client (MRS 3.x or Later)

Step 1 Install a client.

The following uses installing a Flink client on a node in the cluster as an example:

1. Log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, click **More**, and select **Download Client**.
2. In the **Download Cluster Client** dialog box that is displayed, select **Complete Client** for **Select Client Type**, select the platform type that matches the architecture of the node where the client is to install, select **Save to Path**, and click **OK**.
 - The generated file is stored in the **/tmp/FusionInsight-Client** directory on the active management node by default.
 - The name of the client software package is in the follow format: **FusionInsight_Cluster_<Cluster ID>_Services_Client.tar**. In this section, the cluster ID **1** is used as an example. Replace it with the actual cluster ID.
3. Log in to the server where the client is to be installed as the client installation user.
4. Go to the directory where the installation package is stored and run the following commands to decompress the package:

```
cd /tmp/FusionInsight-Client
```

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

5. 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.tar.sha256
```

```
FusionInsight_Cluster_1_Services_ClientConfig.tar: OK
```

6. Decompress the obtained installation file.

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

7. Go to the directory where the installation package is stored, and run the following command to install the client to a specified directory (absolute path), for example, **/opt/hadoopclient**:

```
cd /tmp/FusionInsight-Client/
```

```
FusionInsight_Cluster_1_Services_ClientConfig
```

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

The client is installed if information similar to the following is displayed:

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

Step 2 Log in to the node where the client is installed as the client installation user.

Step 3 Run the following command to go to the client installation directory:

```
cd /opt/hadoopclient
```

Step 4 Run the following command to initialize environment variables:

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

Step 5 Perform the following operations if Kerberos authentication is enabled for the cluster. Otherwise, skip these operations.

1. Prepare a user for submitting Flink jobs.


Log in to FusionInsight Manager and choose **System > Permission > Role**. Click **Create Role** and configure **Role Name** and **Description**. In **Configure Resource Permission**, choose *Name of the desired cluster* > **Flink** and select **FlinkServer Admin Privilege**. Then click **OK**.

Choose **System > Permission > User** and click **Create User**. Configure **Username**, set **User Type** to **Human-Machine**, configure **Password** and **Confirm Password**, click **Add** next to **User Group** to add the **hadoop**, **yarnviewgroup**, and **hadooppmanager** user groups as needed, click **Add** next to **Role** to add the **System_administrator**, **default**, and the created role, and click **OK**. (If you create a Flink job user for the first time, log in to FusionInsight Manager as the user and change the password.)

2. Log in to Manager and download the authentication credential.

Log in to FusionInsight Manager. Choose **System > Permission > User**. In the **Operation** column of the created user, click **More** and select **Download Authentication Credential**.

Figure 6-1 Downloading the authentication credential



- Decompress the downloaded authentication credential package and copy the obtained file to a directory on the client node, for example, **/opt/hadoopclient/Flink/flink/conf**. If the client is installed on a node outside the cluster, copy the obtained file to the **/etc/** directory on this node.
- Add the service IP address of the node where the client is installed and IP address of the master node to the **jobmanager.web.access-control-allow-origin** and **jobmanager.web.allow-access-address** configuration items in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** file. Use commas (,) to separate the IP addresses.

```
jobmanager.web.access-control-allow-origin: xx.xx.xxx.xxx,xx.xx.xxx.xxx,xx.xx.xxx.xxx
jobmanager.web.allow-access-address: xx.xx.xxx.xxx,xx.xx.xxx.xxx,xx.xx.xxx.xxx
```

NOTE

To obtain the service IP address of the node where the client is installed, perform the following operations:

- Node inside the cluster:
 - In the navigation tree of the MRS management console, choose **Clusters > Active Clusters**, select a cluster, and click its name to switch to the cluster details page.
 - On the **Nodes** tab page, view the IP address of the node where the client is installed.
 - Node outside the cluster: IP address of the ECS where the client is installed.
- Configure security authentication by adding the **keytab** path and username in the **/opt/hadoopclient/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/hadoopclient/Flink/flink/conf/user.keytab
security.kerberos.login.principal: test
```

- In the **bin** directory of the Flink client, run the following command to perform security hardening and configure a password used for submitting jobs:

```
cd /opt/hadoopclient/Flink/flink/bin
sh generate_keystore.sh
```

The script automatically replaces the SSL value in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** file.

 NOTE

After authentication and encryption, the **flink.keystore** and **flink.truststore** files are generated in the **conf** directory on the Flink client and the following configuration items are set to the default values in the **flink-conf.yaml** file:

- Set **security.ssl.keystore** to the absolute path of the **flink.keystore** file.
- Set **security.ssl.truststore** to the absolute path of the **flink.truststore** file.
- Set **security.cookie** to a random password automatically generated by the **generate_keystore.sh** script.
- By default, **security.ssl.encrypt.enabled** is set to **false** in the **flink-conf.yaml** file by default. The **generate_keystore.sh** script sets **security.ssl.key-password**, **security.ssl.keystore-password**, and **security.ssl.truststore-password** to the password entered when the **generate_keystore.sh** script is called.

7. Configure paths for the client to access the **flink.keystore** and **flink.truststore** files.

- Relative path (recommended):

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. Create a directory, for example, **ssl**, in **/opt/hadoopclient/Flink/flink/conf/**.

```
cd /opt/hadoopclient/Flink/flink/conf/  
mkdir ssl
```

- ii. Move the **flink.keystore** and **flink.truststore** files to the **/opt/hadoopclient/Flink/flink/conf/ssl/** directory.

```
mv flink.keystore ssl/  
mv flink.truststore ssl/
```

- iii. Change the values of the following parameters to relative paths in the **flink-conf.yaml** file:

```
security.ssl.keystore: ssl/flink.keystore  
security.ssl.truststore: ssl/flink.truststore
```

- Absolute path:

After the **generate_keystore.sh** script is executed, the file path of **flink.keystore** and **flink.truststore** is automatically set to the absolute path **/opt/hadoopclient/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.

Step 6 Run a wordcount job.

NOTICE

To submit or run jobs on Flink, the user must have the following permissions:

- If Ranger authentication is enabled, the current user must belong to the **hadoop** group or the user has been granted the **/flink** read and write permissions in Ranger.
- If Ranger authentication is disabled, the current user must belong to the **hadoop** group.

- For a normal cluster (Kerberos authentication disabled), you can submit jobs in either of the following ways:
 - 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/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
```
 - Run the following command to submit a single job on Yarn:

```
flink run -m yarn-cluster /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
```
- For a security cluster (Kerberos authentication enabled), you can submit jobs in either of the following ways based on the paths of the **flink.keystore** and **flink.truststore** files:
 - If the **flink.keystore** and **flink.truststore** files are stored in the relative path:
 - Run the following command in the directory at the same level as **ssl** to start the session and submit the job in the session:

```
ssl
```

 is a relative path. For example, if **ssl** is in **opt/hadoopclient/Flink/flink/conf/**, run the command in the **opt/hadoopclient/Flink/flink/conf/** directory.

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

```
cd /opt/hadoopclient/Flink/flink/conf
flink run -m yarn-cluster -yt ssl/ /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
```
 - If the **flink.keystore** and **flink.truststore** files are stored in the absolute path:
 - Run the following commands to start a session and submit a job in the session:

```
cd /opt/hadoopclient/Flink/flink/conf
yarn-session.sh -nm "session-name" -d
flink run /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
```


- Run the following command to submit a single job on Yarn:
flink run -m yarn-cluster /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

Step 7 After the job has been successfully submitted, the following information is displayed on the client:

Figure 6-2 Job submitted successfully on Yarn

```
[root@node-master1kszp ~]# flink run -m yarn-cluster /opt/client/Flink/flink/examples/streaming/WordCount.jar
2019-07-10 16:30:11,090 [WARN | main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory (DomainSocketFactory.java:212)
2019-07-10 16:30:11,090 [WARN | main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory (DomainSocketFactory.java:118)
Starting execution of program
Executing WordCount example with default input data set.
Use --input to specify file input.
Printing result to stdout. Use --output to specify output path.
Program execution finished
Job with JobID c0831921e0e1ef2bba24b51a5b61d has finished.
Job Runtime: 7953 ms
```

Figure 6-3 Session started successfully

```
[root@node-master1kszp ~]# hive> yarn-session.sh -m "test4dce" -d
2019-07-26 09:17:08,919 [WARN | main] | Unable to load native-hadoop library for your platform... using builtin-java classes where applicable | org.apache.hadoop.util.NativeCodeLoader (NativeCodeLoader.java:62)
(DomainSocketFactory.java:118) [WARN | main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory
Flink JobManager is now running on node-ana-corehdp:32598 with leader id b9b5ab8-1983-435f-bb90-ad128fd1d46b.
JobManager Web Interface: http://192.168.2.01:4769
[root@node-master1kszp ~]#
```

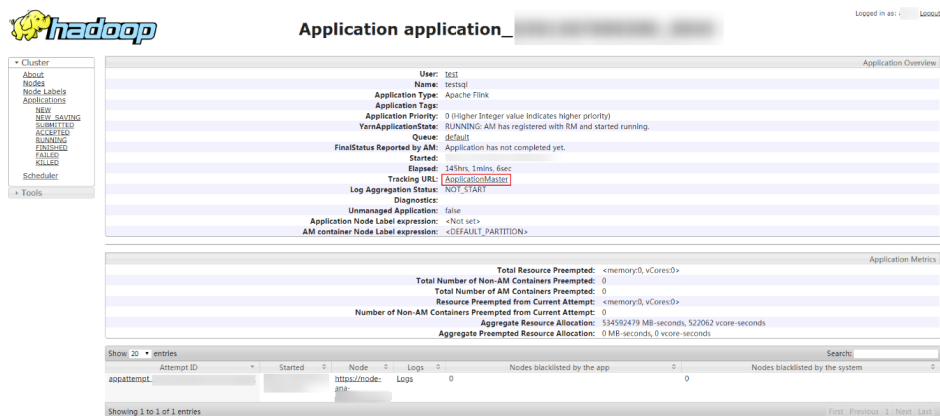
Figure 6-4 Job submitted successfully in the session

```
[root@node-master1kszp ~]# flink run -m yarn-cluster /opt/client/Flink/flink/examples/streaming/WordCount.jar
YARN properties set default parallelism to 3
2019-07-26 09:19:20,548 [WARN | main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory
(DomainSocketFactory.java:118) [WARN | main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory
Starting execution of program
Executing WordCount example with default input data set.
Use --input to specify file input.
Printing result to stdout. Use --output to specify output path.
Program execution finished
Job with JobID 5b0bc18d0563fd3d792a19163c267c3c3 has finished.
Job Runtime: 5509 ms
[root@node-master1kszp ~]#
```

Step 8 Go to the native Yarn service page, find the application of the job, and click the application name to go to the job details page.

- If the job is not completed, click **Tracking URL** to go to the native Flink page and view the job running information.
- If the job submitted in a session has been completed, you can click **Tracking URL** to log in to the native Flink service page to view job information.

Figure 6-5 application



----End

Using the Flink Client (Versions Earlier Than MRS 3.x)

Step 1 Install a client.

The following uses installing a Flink client on a core node as an example:

1. Log in to MRS Manager and choose **Services > Download Client** to download the client installation package to the active management node.
2. Use the IP address to search for the active management node, and log in to the active management node over VNC.
3. Log in to the active management node, and run the following command to switch the user:

```
sudo su - omm
```

4. On the MRS management console, view the IP address on the **Nodes** tab page of the specified cluster.

Record the IP address of the core node where the client is to be used.

5. On the active management node, run the following command to copy the client installation package to the core node:

```
scp -p /tmp/MRS-client/MRS_Services_Client.tar IP address of the core node:/opt/client
```

6. Log in to the core node as user **root**.

Master nodes support Cloud-Init. The preset username for Cloud-Init is **root**, and the password is the one set during cluster creation.

7. Run the following commands to install the client:

```
cd /opt/client
```

```
tar -xvf MRS_Services_Client.tar
```

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

```
cd /opt/client/MRS_Services_ClientConfig
```

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

For example, run the following command:

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

Step 2 Log in to the node where the client is installed as the client installation user.

Step 3 Run the following command to go to the client installation directory:

```
cd /opt/hadoopclient
```

Step 4 Run the following command to initialize environment variables:

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

Step 5 If Kerberos authentication is enabled for the cluster, perform the following steps. If not, skip this whole step.

1. Prepare a user for submitting Flink jobs.
2. Log in to Manager and download the authentication credential.

Log in to Manager of the cluster. For details, see [Accessing MRS Manager \(Versions Earlier Than MRS 3.x\)](#). Choose **System Settings > User Management**. In the **Operation** column of the row that contains the added user, choose **More > 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/hadoopclient/Flink/flink/conf**. If the client is installed on a node outside the cluster, copy the obtained file to the **/etc/** directory on this node.
4. Configure security authentication by adding the **keytab** path and username in the **/opt/hadoopclient/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/hadoopclient/Flink/flink/conf/user.keytab
security.kerberos.login.principal: test
```

5. In the **bin** directory of the Flink client, run the following command to perform security hardening and configure a password used for submitting jobs:

```
cd /opt/hadoopclient/Flink/flink/bin
sh generate_keystore.sh
```

This script automatically replaces the SSL value in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** file. For MRS 2.x or earlier, external SSL is disabled for security clusters by default. To enable external SSL, configure it by referring to "Authentication and Encryption" and run this script again.

NOTE

- You do not need to manually generate the **generate_keystore.sh** script.
 - After authentication and encryption, the generated **flink.keystore**, **flink.truststore**, and **security.cookie** items are automatically filled in the corresponding configuration items in **flink-conf.yaml**.
6. Configure paths for the client to access the **flink.keystore** and **flink.truststore** files.
 - Relative path (recommended):

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. Create a directory, for example, **ssl**, in **/opt/hadoopclient/Flink/flink/conf/**.


```
cd /opt/hadoopclient/Flink/flink/conf/
mkdir ssl
```
 - ii. Move the **flink.keystore** and **flink.truststore** files to the **/opt/hadoopclient/Flink/flink/conf/ssl/** directory.


```
mv flink.keystore ssl/
mv flink.truststore ssl/
```
 - iii. Change the values of the following parameters to relative paths in the **flink-conf.yaml** file:


```
security.ssl.internal.keystore: ssl/flink.keystore
security.ssl.internal.truststore: ssl/flink.truststore
```
 - Absolute path:

After the **generate_keystore.sh** script is executed, the file path of **flink.keystore** and **flink.truststore** is automatically set to the absolute path **/opt/hadoopclient/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.

Step 6 Run a wordcount job.

NOTICE

To submit or run jobs on Flink, the user must have the following permissions:

- If Ranger authentication is enabled, the current user must belong to the **hadoop** group or the user has been granted the **/flink** read and write permissions in Ranger.
- If Ranger authentication is disabled, the current user must belong to the **hadoop** group.

-
- For a normal cluster (Kerberos authentication disabled), you can submit jobs in either of the following ways:
 - 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/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
```
 - Run the following command to submit a single job on Yarn:

```
flink run -m yarn-cluster /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
```
 - For a security cluster (Kerberos authentication enabled), you can submit jobs in either of the following ways based on the paths of the **flink.keystore** and **flink.truststore** files:
 - If the **flink.keystore** and **flink.truststore** files are stored in the relative path:
 - Run the following command in the directory at the same level as **ssl** to start the session and submit the job in the session:

```
ssl
```

 is a relative path. For example, if **ssl** is in **opt/hadoopclient/Flink/flink/conf/**, run the command in the **opt/hadoopclient/Flink/flink/conf/** directory.

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

```
cd /opt/hadoopclient/Flink/flink/conf
flink run -m yarn-cluster -yt ssl/ /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
```
 - If the **flink.keystore** and **flink.truststore** files are stored in the absolute path:

- Run the following commands to start a session and submit a job in the session:


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


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

Step 7 After the job has been successfully submitted, the following information is displayed on the client:

Figure 6-6 Job submitted successfully on Yarn

```
[root@node-master1kz2P ~]# flink run -m yarn-cluster /opt/client/Flink/flink/examples/streaming/WordCount.jar
2019-07-10 16:30:11,090 | WARN | [main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory (DomainSocketFactory.java:212)
2019-07-10 16:30:11,090 | WARN | [main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory (DomainSocketFactory.java:118)
Starting execution of program
Executing WordCount example with default input data set.
Use --input to specify file input.
Printing result to stdout. Use --output to specify output path.
Program execution finished
Job with JobID c042b1921e0a1efezbba24b51a5be1d has finished.
Job Runtime: 7953 ms
```

Figure 6-7 Session started successfully

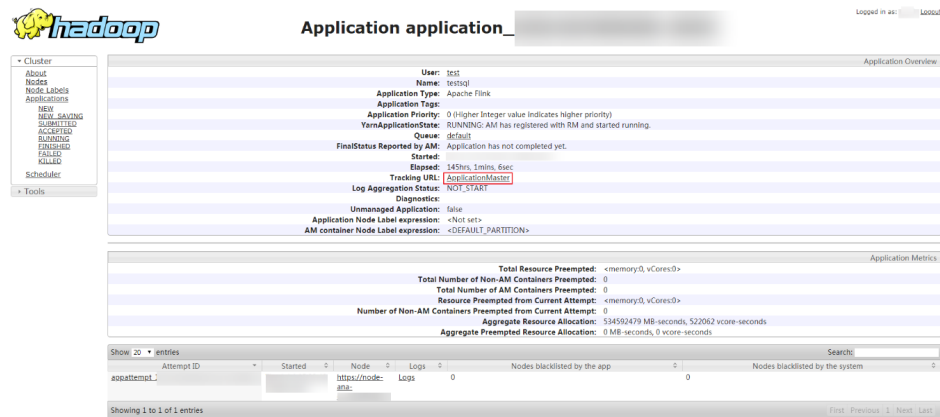
```
[root@node-master1kz2P Hive]# yarn-session.sh -nm "test4doc" -d
2019-07-26 09:17:08,619 | WARN | [main] | Unable to load native-hadoop library for your platform... using builtin-java classes where applicable | org.apache.hadoop.util.NativeCodeLoader (NativeCodeLoader.java:91)
2019-07-26 09:17:08,886 | WARN | [main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory (DomainSocketFactory.java:118)
Flink JobManager is now running on node-ana-corehdp:32586 with leader id b9b5ab0-1983-435f-bb90-ad128fd1d46b.
JobManager Web Interface: http://192.168.2.61:47897
[root@node-master1kz2P Hive]#
```

Figure 6-8 Job submitted successfully in the session

```
[root@node-master1kz2P Hive]# flink run /opt/client/Flink/flink/examples/streaming/WordCount.jar
YARN properties set default parallelism to 3
2019-07-26 09:19:20,248 | WARN | [main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory (DomainSocketFactory.java:118)
2019-07-26 09:19:20,548 | WARN | [main] | The short-circuit local reads feature cannot be used because libhadoop cannot be loaded. | org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory (DomainSocketFactory.java:118)
Starting execution of program
Executing WordCount example with default input data set.
Use --input to specify file input.
Printing result to stdout. Use --output to specify output path.
Program execution finished
Job with JobID Shubc3960563f3d792a19163c2e7c3c3 has finished.
Job Runtime: 5906 ms
[root@node-master1kz2P Hive]#
```

- Step 8** Go to the native Yarn service page, find the application of the job, and click the application name to go to the job details page.
- If the job is not completed, click **Tracking URL** to go to the native Flink page and view the job running information.
 - If the job submitted in a session has been completed, you can click **Tracking URL** to log in to the native Flink service page to view job information.

Figure 6-9 Application



----End

6.2 Viewing Flink Job Information

You can view Flink job information on the Yarn web UI.

Prerequisites

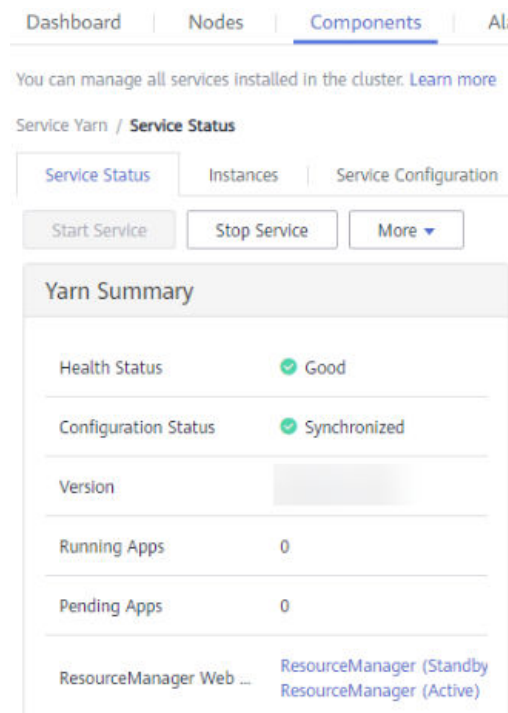
The Flink service has been installed in a cluster.

Accessing the Yarn Web UI

Step 1 Go to the Yarn service page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components > Yarn > Yarn Summary**.

Figure 6-10 Yarn summary



NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later: Log in to FusionInsight Manager, choose **Cluster**, click the name of the target cluster, choose **Services** > **Yarn**, and click **Dashboard**.

Step 2 Click the link next to **ResourceManager WebUI** to go to the Yarn web UI page.

----End

6.3 Configuring Flink Service Parameters

Description

All Flink parameters are configurable on the client. You are advised to modify the **flink-conf.yaml** configuration file on the client. If Flink parameters are modified on FusionInsight Manager, you need to download and install the client again after the configuration is complete.

- Configuration file path: *Client installation path*/**Flink/flink/conf/flink-conf.yaml**
- File configuration format: *Key:Value*

Example: **taskmanager.heap.size: 1024mb**

A space is required between *Key:* and *Value.*

Configurations

The following describes the parameters:

- **JobManager & TaskManager:**
JobManager and TaskManager are main components of Flink, which is used for different security and performance scenarios. The configuration items include communication ports, memory management, and connection retry.
- **Blob server:**
The Blob server on the JobManager node is used to receive JAR files uploaded by users on the client, send JAR files to TaskManager, and transfer log files. The configuration items include ports, SSL, retries, and concurrency.
- **Distributed Coordination (via Akka):**
The Akka actor model is the basis of communications between a Flink client and JobManager, JobManager and TaskManager, as well as TaskManagers. Related parameters can be configured based on the network environment or optimization policy. The configuration items include the timeout settings for message sending and waiting and the Akka listening mechanism DeathWatch.
- **SSL:**
To configure a secure Flink cluster, you need to configure SSL-related configuration items, including SSL, certificate, password, and encryption algorithm.
- **Network communication (via Netty):**
When Flink runs a job, data transmission and reverse pressure detection between tasks depend on Netty. In certain environments, **Netty** parameters should be configured. For advanced optimization, you can adjust some Netty configuration items. The default configuration can meet the requirements of concurrent and high-throughput tasks in a large-scale cluster.
- **JobManager Web Frontend:**
When JobManager is started, the web server is started in the same process. You can access the web server to obtain information about the current Flink cluster, including JobManager, TaskManager, and jobs running in the cluster. Configuration items of the web server include the port, temporary directory, display items, error redirection, and security-related items.
- **File Systems:**
When a task is running, a result file is created. You can configure the file creation behavior, including the file overwriting strategy and directory creation.
- **State Backend:**
Flink enables HA and job exception, as well as job pause and recovery during version upgrade. Flink depends on state backend to store job states and on the restart strategy to restart a job. You can configure state backend and the restart strategy. Configuration items include the state backend type, storage path, and restart strategy.
- **Kerberos-based Security:**
Kerberos-related configuration items must be configured in Flink's security mode. The configuration items include keytab and principal of Kerberos.
- **HA:**

The HA mode of Flink depends on ZooKeeper. So, ZooKeeper configuration items must be configured, including the ZooKeeper address, path, and security authentication.

- Environment:**

In scenarios raising special requirements on JVM configuration, users can use configuration items to transfer JVM parameters to the client, JobManager, and TaskManager.
- Yarn:**

Flink runs on a Yarn cluster and JobManager runs on ApplicationMaster. Some configuration parameters of JobManager depend on YARN. You can configure YARN-related configurations to enable Flink to better run on YARN. The configuration items include the memory, virtual kernel, and port of YARN containers.
- Pipeline:**

The Netty connection is used among multiple jobs to reduce latency. In this case, NettySink is used on the server and NettySource is used on the client for data transmission. Configuration items include NettySink information storing path, range of NettySink listening port, whether to enable SSL encryption, domain of the network used for NettySink monitoring.

JobManager & TaskManager

Table 6-1 JobManager & TaskManager parameters

Parameter	Description	Default Value	Mandatory	Remarks
taskmanager.memory.size	Amount of heap memory of the Java virtual machine (JVM) that TaskManager reserves for sorting, hash tables, and caching of intermediate results. If unspecified, the memory manager will take a fixed ratio with respect to the size of JVM as specified by taskmanager.memory.fraction . The unit is MB.	0	No	Only versions earlier than MR S 3.x
taskmanager.registration.initial-backoff	Initial interval between two consecutive registration attempts. The unit is ms/s/m/h/d. The time value and unit are separated by half-width spaces. ms/s/m/h/d indicates millisecond, second, minute, hour, and day, respectively.	500 ms	No	

Parameter	Description	Default Value	Mandatory	Remarks
taskmanager.registration.refused-backoff	Retry interval when a registration connection is rejected by JobManager.	5 min	No	
taskmanager.rpc.port	IPC port range of TaskManager	32326-32390	No	Applicable to all versions
taskmanager.memory.segment-size	Size of the memory buffer used by the memory manager and network stack. The unit is bytes.	32768	No	
taskmanager.data.port	Data exchange port range of TaskManager	32391-32455	No	
taskmanager.data.ssl.enabled	Whether to enable secure sockets layer (SSL) encryption for data transfer between TaskManagers. This parameter is valid only when the global switch security.ssl is enabled.	false	No	
taskmanager.numberOfTaskSlots	Number of slots occupied by TaskManager. Generally, the value is configured as the number of cores of the physical machine. In yarn-session mode, the value can be transmitted by only the -s parameter. In yarn-cluster mode, the value can be transmitted by only the -ys parameter.	1	No	
parallelism.default	Default degree of parallelism, which is used for jobs for which the degree of parallelism is not specified	1	No	
task.cancellation.interval	Interval between two successive task cancellation attempts. The unit is millisecond.	30000	No	
client.rpc.port	Akka system listening port on the Flink client.	32651-32720	No	

Parameter	Description	Default Value	Mandatory	Remarks
jobmanager.heap.size	Size of the heap memory of JobManager. In yarn-session mode, the value can be transmitted by only the -jm parameter. In yarn-cluster mode, the value can be transmitted by only the -yjm parameter. If the value is smaller than yarn.scheduler.minimum-allocation-mb in the Yarn configuration file, the Yarn configuration value is used. The unit is B/KB/MB/GB/TB.	1024mb	No	
taskmanager.heap.size	Size of the heap memory of TaskManager. In yarn-session mode, the value can be transmitted by only the -tm parameter. In yarn-cluster mode, the value can be transmitted by only the -ytm parameter. If the value is smaller than yarn.scheduler.minimum-allocation-mb in the Yarn configuration file, the Yarn configuration value is used. The unit is B/KB/MB/GB/TB.	1024mb	No	
taskmanager.network.numberOfBuffers	Number of TaskManager network transmission buffer stacks. If an error indicates insufficient system buffer, increase the parameter value.	2048	No	
taskmanager.debug.memory.startLoggingThread	Enable this item for debugging Flink memory and garbage collection (GC)-related problems. TaskManager periodically collects memory and GC statistics, including the current utilization of heap and off-heap memory pools and GC time.	false	No	
taskmanager.debug.memory.logIntervalMs	Interval at which TaskManager periodically collects memory and GC statistics.	0	No	

Parameter	Description	Default Value	Mandatory	Remarks
taskmanager.maxRegistrationDuration	Maximum duration of TaskManager registration on JobManager. If the actual duration exceeds the value, TaskManager is disabled.	5 min	No	
taskmanager.initial-registration-pause	Initial interval between two consecutive registration attempts. The value must contain a time unit (ms/s/min/h/d), for example, 5 seconds. The time value and unit are separated by half-width spaces. ms/s/m/h/d indicates millisecond, second, minute, hour, and day, respectively.	500 ms	No	
taskmanager.max-registration-pause	Maximum registration retry interval in case of TaskManager registration failures. The unit is ms/s/m/h/d.	30 s	No	
taskmanager.refused-registration-pause	Retry interval when a TaskManager registration connection is rejected by JobManager. The unit is ms/s/m/h/d.	10 s	No	
classloader.resolve-order	Class resolution policies defined when classes are loaded from user codes, which means whether to first check the user code JAR file (child-first) or the application class path (parent-first). The default setting indicates that the class is first loaded from the user code JAR file, which means that the user code JAR file can contain and load dependencies that are different from those used by Flink.	child-first	No	
slot.idle.timeout	Timeout for an idle slot in Slot Pool, in milliseconds.	50000	No	
slot.request.timeout	Timeout for requesting a slot from Slot Pool, in milliseconds.	300000	No	

Parameter	Description	Default Value	Mandatory	Remarks
task.cancellation.timeout	Timeout of task cancellation, in milliseconds. If a task cancellation times out, a fatal TaskManager error may occur. If this parameter is set to 0 , no error is reported when a task cancellation times out.	180000	No	
taskmanager.network.detailed-metrics	Indicates whether to enable the detailed metrics monitoring of network queue lengths.	false	No	
taskmanager.network.memory.buffers-per-channel	Maximum number of network buffers used by each output/input channel (sub-partition/incoming channel). In credit-based flow control mode, this indicates how much credit is in each input channel. It should be configured with at least 2 buffers to deliver good performance. One buffer is used to receive in-flight data in the sub-partition, and the other for parallel serialization.	2	No	
taskmanager.network.memory.floating-buffers-per-gate	Number of extra network buffers used by each output gate (result partition) or input gate, indicating the amount of floating credit shared among all input channels in credit-based flow control mode. Floating buffers are distributed based on the backlog feedback (real-time output buffers in sub-partitions) and can help mitigate back pressure caused by unbalanced data distribution among sub-partitions. Increase this value if the round-trip time between nodes is long and/or the number of machines in the cluster is large.	8	No	

Parameter	Description	Default Value	Mandatory	Remarks
taskmanager.network.memory.fraction	Ratio of JVM memory used for network buffers, which determines how many streaming data exchange channels a TaskManager can have at the same time and the extent of channel buffering. Increase this value or the values of taskmanager.network.memory.min and taskmanager.network.memory.max if the job is rejected or a warning indicating that the system does not have enough buffers is received. Note that the values of taskmanager.network.memory.min and taskmanager.network.memory.max may overwrite this value.	0.1	No	
taskmanager.network.memory.max	Maximum memory size of the network buffer. The value must contain a unit (B/KB/MB/GB/TB).	1 GB	No	
taskmanager.network.memory.min	Minimum memory size of the network buffer. The value must contain a unit (B/KB/MB/GB/TB).	64 MB	No	
taskmanager.network.request-backoff.initial	Minimum backoff for partition requests of input channels.	100	No	
taskmanager.network.request-backoff.max	Maximum backoff for partition requests of input channels.	10000	No	
taskmanager.registration.timeout	Timeout for TaskManager registration. TaskManager will be terminated if it is not successfully registered within the specified time. The value must contain a time unit (ms/s/min/h/d).	5 min	No	
resourcemanager.taskmanager.timeout	Timeout interval for releasing an idle TaskManager, in milliseconds.	30000	No	

Blob server

Table 6-2 Blob server parameters

Parameter	Description	Default Value	Mandatory
blob.server.port	Blob server port	32456-32520	No
blob.service.ssl.enabled	Indicates whether to enable the encryption for the blob transmission channel. This parameter is valid only when the global switch security.ssl is enabled.	true	Yes
blob.fetch.retries	Number of times that TaskManager tries to download blob files from JobManager.	50	No
blob.fetch.num-concurrent	Number of concurrent tasks for downloading blob files supported by JobManager.	50	No
blob.fetch.backlog	Number of blob files, such as .jar files, to be downloaded in the queue supported by JobManager. The unit is count.	1000	No
library-cache-manager.cleanup.interval	Interval at which JobManager deletes the JAR files stored on the HDFS when the user cancels the Flink job. The unit is second. This parameter is available only to MRS 3.x or later.	3600	No

Distributed Coordination (via Akka)

Table 6-3 Distributed Coordination parameters

Parameter	Description	Default Value	Mandatory	Remarks
akka.ask.timeout	Timeout duration of Akka asynchronous and block requests. If a Flink timeout failure occurs, this value can be increased. Timeout occurs when the machine processing speed is slow or the network is blocked. The unit is ms/s/m/h/d.	10s	No	Applicable to all versions
akka.lookup.timeout	Timeout duration for JobManager actor object searching. The unit is ms/s/m/h/d.	10s	No	
akka.framesize	Maximum size of the message transmitted between JobManager and TaskManager. If a Flink error occurs because the message exceeds this limit, the value can be increased. The unit is b/B/KB/MB.	10485760b	No	
akka.watch.heartbeat.interval	Heartbeat interval at which the Akka DeathWatch mechanism detects disconnected TaskManager. If TaskManager is frequently and incorrectly marked as disconnected due to heartbeat loss or delay, the value can be increased. The unit is ms/s/m/h/d.	10s	No	
akka.watch.heartbeat.pause	Acceptable heartbeat pause for Akka DeathWatch mechanism. A small value indicates that irregular heartbeat is not accepted. The unit is ms/s/m/h/d.	60s	No	
akka.watch.threshold	DeathWatch failure detection threshold. A small value may mark normal TaskManager as failed and a large value increases failure detection time.	12	No	

Parameter	Description	Default Value	Mandatory	Remarks
akka.tcp.timeout	Timeout duration of Transmission Control Protocol (TCP) connection request. If TaskManager connection timeout occurs frequently due to the network congestion, the value can be increased. The unit is ms/s/m/h/d.	20s	No	
akka.throughput	Number of messages processed by Akka in batches. After an operation, the processing thread is returned to the thread pool. A small value indicates the fair scheduling for actor message processing. A large value indicates improved overall performance but lowered scheduling fairness.	15	No	
akka.log.lifecycle.events	Switch of Akka remote time logging, which can be enabled for debugging.	false	No	
akka.startup-timeout	Timeout interval before a remote component fails to be started. The value must contain a time unit (ms/s/min/h/d).	The value is the same as the value of akka.ask.timeout .	No	
akka.ssl.enabled	Switch of Akka communication SSL. This parameter is valid only when the global switch security.ssl is enabled.	true	Yes	
akka.client-socket-worker-pool.pool-size-factor	Factor that is used to determine the thread pool size. The pool size is calculated based on the following formula: $\text{ceil}(\text{available processors} * \text{factor})$. The size is bounded by the pool-size-min and pool-size-max values.	1.0	No	Applicable only to MR S 3.x or later
akka.client-socket-worker-pool.pool-size-max	Maximum number of threads calculated based on the factor.	2	No	

Parameter	Description	Default Value	Mandatory	Remarks
akka.client-socket-worker-pool.pool-size-min	Minimum number of threads calculated based on the factor.	1	No	
akka.client.timeout	Timeout duration of the client. The value must contain a time unit (ms/s/min/h/d).	60s	No	
akka.server-socket-worker-pool.pool-size-factor	Factor that is used to determine the thread pool size. The pool size is calculated based on the following formula: $\text{ceil}(\text{available processors} * \text{factor})$. The size is bounded by the pool-size-min and pool-size-max values.	1.0	No	
akka.server-socket-worker-pool.pool-size-max	Maximum number of threads calculated based on the factor.	2	No	
akka.server-socket-worker-pool.pool-size-min	Minimum number of threads calculated based on the factor.	1	No	

SSL

Table 6-4 SSL parameters

Parameter	Description	Default Value	Mandatory	Remarks
security.ssl.internal.enabled	Whether to enable SSL for internal communication. This parameter is automatically configured based on the security mode of the cluster.	<ul style="list-style-type: none"> Security mode: true Normal mode: false 	Yes	Applicable only to versions earlier than MRS 3.x
security.ssl.internal.keystore	Java keystore file.	-	Yes	
security.ssl.internal.keystore-password	Password used to decrypt the keystore file.	-	Yes	
security.ssl.internal.key-password	Password used to decrypt the server key in the keystore file.	-	Yes	
security.ssl.internal.truststore	truststore file containing the public CA certificates.	-	Yes	
security.ssl.internal.truststore-password	Password used to decrypt the truststore file.	-	Yes	
security.ssl.rest.enabled	Whether to enable SSL for external communication. This parameter is automatically configured based on the security mode of the cluster.	<ul style="list-style-type: none"> Security mode: true Normal mode: false 	Yes	
security.ssl.rest.keystore	Java keystore file.	-	Yes	
security.ssl.rest.keystore-password	Password used to decrypt the keystore file.	-	Yes	
security.ssl.rest.key-password	Password used to decrypt the server key in the keystore file.	-	Yes	
security.ssl.rest.truststore	truststore file containing the public CA certificates.	-	Yes	
security.ssl.rest.truststore-password	Password used to decrypt the truststore file.	-	Yes	

Parameter	Description	Default Value	Mandatory	Remarks
security.ssl.protocol	SSL transmission protocol version.	TLSv1.2	Yes	Applicable to all versions
security.ssl.algorithms	Supported SSL standard algorithm.	TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	Yes	
security.ssl.enabled	Whether to enable SSL for internal communication. This parameter is automatically configured based on the cluster installation mode.	<ul style="list-style-type: none"> Security mode: true Normal mode: false 	Yes	Applicable only to MRS 3.x or later
security.ssl.keystore	Java keystore file.	-	Yes	
security.ssl.keystore-password	Password used to decrypt the keystore file.	-	Yes	
security.ssl.key-password	Password used to decrypt the server key in the keystore file.	-	Yes	
security.ssl.truststore	truststore file containing the public CA certificates.	-	Yes	
security.ssl.truststore-password	Password used to decrypt the truststore file.	-	Yes	

Network communication (via Netty)

Table 6-5 Network communication parameters

Parameter	Description	Default Value	Mandatory
taskmanager.network.netty.num-arenas	Number of Netty memory blocks.	1	No
taskmanager.network.netty.server.numThreads	Number of Netty server threads	1	No
taskmanager.network.netty.client.numThreads	Number of Netty client threads	1	No
taskmanager.network.netty.client.connectTimeoutSec	Netty client connection timeout duration. The unit is second.	120	No
taskmanager.network.netty.sendReceiveBufferSize	Size of Netty sending and receiving buffers. This defaults to the system buffer size (<code>cat /proc/sys/net/ipv4/tcp_[rw]mem</code>) and is 4 MB in modern Linux. The unit is bytes.	4096	No
taskmanager.network.netty.transport	Netty transport type, either nio or epoll	nio	No

JobManager Web Frontend

Table 6-6 JobManager Web Frontend parameters

Parameter	Description	Default Value	Mandatory	Remarks
jobmanager.web.port	The web port. Value range: 32261–32325.	32261-32325	No	Applicable only to versions earlier than MR S 3.x
jobmanager.web.allow-access-address	Web access whitelist. IP addresses are separated by commas (,). Only whitelisted IP addresses can access the web.	*	Yes	Applicable to all versions
flink.security.enable	<p>When installing a Flink cluster, you are required to select security mode or normal mode.</p> <ul style="list-style-type: none"> If security mode is selected, this parameter is automatically set to true. If normal mode is selected, this parameter is automatically set to false. <p>For an installed Flink cluster, you can view the configured value to determine whether the cluster is in security or normal mode.</p>	Automatic configuration	No	Applicable only to MR S 3.x or later
rest.bind-port	The web port. Value range: 32261–32325.	32261-32325	No	
jobmanager.web.history	Number of recent jobs to be displayed.	5	No	

Parameter	Description	Default Value	Mandatory	Remarks
jobmanager.web.checkpoints.disable	Whether to disable checkpoint statistics.	false	No	
jobmanager.web.checkpoints.history	Number of checkpoint statistical records.	10	No	
jobmanager.web.backpressure.cleanup-interval	Interval for clearing unaccessed backpressure records. The unit is millisecond.	600000	No	
jobmanager.web.backpressure.refresh-interval	Interval for updating backpressure records. The unit is millisecond.	60000	No	
jobmanager.web.backpressure.num-samples	Number of stack tracing records for reverse pressure calculation.	100	No	
jobmanager.web.backpressure.delay-between-samples	Sampling interval for reverse pressure calculation. The unit is millisecond.	50	No	
jobmanager.web.ssl.enabled	Whether SSL encryption is enabled for web transmission. This parameter is valid only when the global switch security.ssl is enabled.	false	Yes	
jobmanager.web.accesslog.enable	Switch to enable or disable web operation logs. The log is stored in webaccess.log .	true	Yes	
jobmanager.web.x-frame-options	Value of the HTTP security header X-Frame-Options . The value can be SAMEORIGIN , DENY , or ALLOW-FROM uri .	DENY	Yes	
jobmanager.web.cache-directive	Whether the web page can be cached.	no-store	Yes	
jobmanager.web.expires-time	Expiration duration of web page cache. The unit is millisecond.	0	Yes	

Parameter	Description	Default Value	Mandatory	Remarks
jobmanager.web.access-control-allow-origin	Web page same-origin policy that prevents cross-domain attacks.	*	Yes	
jobmanager.web.refresh-interval	Web page refresh interval. The unit is millisecond.	3000	Yes	
jobmanager.web.logout-timer	Automatic logout interval when no operation is performed. The unit is millisecond.	600000	Yes	
jobmanager.web.403-redirect-url	Web page access error 403. If 403 error occurs, the page switch to a specified page.	Automatic configuration	Yes	
jobmanager.web.404-redirect-url	Web page access error 404. If 404 error occurs, the page switch to a specified page.	Automatic configuration	Yes	
jobmanager.web.415-redirect-url	Web page access error 415. If 415 error occurs, the page switch to a specified page.	Automatic configuration	Yes	
jobmanager.web.500-redirect-url	Web page access error 500. If 500 error occurs, the page switch to a specified page.	Automatic configuration	Yes	
rest.await-leader-timeout	Time of the client waiting for the leader address. The unit is millisecond.	30000	No	
rest.client.max-content-length	Maximum content length that the client handles (unit: bytes).	104857600	No	
rest.connection-timeout	Maximum time for the client to establish a TCP connection (unit: ms).	15000	No	
rest.idleness-timeout	Maximum time for a connection to stay idle before failing (unit: ms).	300000	No	
rest.retry.delay	The time that the client waits between retries (unit: ms).	3000	No	

Parameter	Description	Default Value	Mandatory	Remarks
rest.retry.max-attempts	The number of retry times if a retrievable operator fails.	20	No	
rest.server.max-content-length	Maximum content length that the server handles (unit: bytes).	104857600	No	
rest.server.numThreads	Maximum number of threads for the asynchronous processing of requests.	4	No	
web.timeout	Timeout for web monitor (unit: ms).	10000	No	

File Systems

Table 6-7 File Systems parameters

Parameter	Description	Default Value	Mandatory
fs.overwrite-files	Whether to overwrite the existing file by default when the file is written.	false	No
fs.output.always-create-directory	<p>When the degree of parallelism (DOP) of file writing programs is greater than 1, a directory is created under the output file path and different result files (each parallel write program) are stored in the directory.</p> <ul style="list-style-type: none"> If this parameter is set to true, a directory is created for the writing program whose DOP is 1 and a result file is stored in the directory. If this parameter is set to false, the file of the writing program whose DOP is 1 is created directly in the output path and no directory is created. 	false	No

State Backend

Table 6-8 State Backend parameters

Parameter	Description	Default Value	Mandatory
state.backend.fs.checkpointdir	Path when the backend is set to filesystem . The path must be accessible by JobManager. Only the local mode is supported. In the cluster mode, use an HDFS path.	hdfs:///flink/checkpoints	No
state.savepoints.dir	Savepoint storage directory used by Flink to restore and update jobs. When a savepoint is triggered, the metadata of the savepoint is saved to this directory.	hdfs:///flink/savepoint	Mandatory in security mode
restart-strategy	Default restart policy, which is used for jobs for which no restart policy is specified. <ul style="list-style-type: none"> fixed-delay failure-rate none 	none	No
restart-strategy.fixed-delay.attempts	The retry times of the fixed-delay policy.	<ul style="list-style-type: none"> If the checkpoint is enabled, the default value is the value of Integer.MAX_VALUE. If the checkpoint is disabled, the default value is 3. 	No

Parameter	Description	Default Value	Mandatory
restart-strategy.fixed-delay.delay	Retry interval when the fixed-delay strategy is used. The unit is ms/s/m/h/d.	<ul style="list-style-type: none"> If the checkpoint is enabled, the default value is 10s. If the checkpoint is disabled, the default value is the value of akka.ask.timeout. 	No
restart-strategy.failure-rate.max-failures-per-interval	Maximum number of restart times in a specified period before a job fails when the fault rate policy is used.	1	No
restart-strategy.failure-rate.failure-rate-interval	Retry interval when the failure-rate strategy is used. The unit is ms/s/m/h/d.	60 s	No
restart-strategy.failure-rate.delay	Retry interval when the failure-rate strategy is used. The unit is ms/s/m/h/d.	The default value is the same as the value of akka.ask.timeout . For details, see Distributed Coordination (via Akka) .	No

Kerberos-based Security

Table 6-9 Kerberos-based security parameters

Parameter	Description	Default Value	Mandatory
security.kerberos.login.keytab	Keytab file path. This parameter is a client parameter.	Configure the parameter based on actual service requirements.	Yes
security.kerberos.login.principal	A parameter on the client. If security.kerberos.login.keytab and security.kerberos.login.principal are both set, keytab certificate is used by default.	Configure the parameter based on actual service requirements.	No
security.kerberos.login.contexts	Contexts of the jass file generated by Flink. This parameter is a server parameter.	Client, KafkaClient	Yes

HA

Table 6-10 HA parameters

Parameter	Description	Default Value	Mandatory
high-availability	<p>Whether HA is enabled. Only the following two modes are supported currently:</p> <ul style="list-style-type: none"> • none: Only a single JobManager is running. The checkpoint is disabled for JobManager. • ZooKeeper: <ul style="list-style-type: none"> - In non-Yarn mode, multiple JobManagers are supported and the leader JobManager is elected. - In Yarn mode, only one JobManager exists. 	zookeeper	No

Parameter	Description	Default Value	Mandatory
high-availability.zookeeper.quorum	ZooKeeper quorum address.	Automatic configuration	No
high-availability.zookeeper.path.root	Root directory that Flink creates on ZooKeeper, storing metadata required in HA mode.	/flink	No
high-availability.storageDir	Directory for storing JobManager metadata of state backend. ZooKeeper stores only pointers to actual data.	hdfs:///flink/recovery	No
high-availability.zookeeper.client.session-timeout	Session timeout duration on the ZooKeeper client. The unit is millisecond.	60000	No
high-availability.zookeeper.client.connection-timeout	Connection timeout duration on the ZooKeeper client. The unit is millisecond.	15000	No
high-availability.zookeeper.client.retry-wait	Retry waiting time on the ZooKeeper client. The unit is millisecond.	5000	No
high-availability.zookeeper.client.max-retry-attempts	Maximum retry times on the ZooKeeper client.	3	No
high-availability.job.delay	Delay of job restart when JobManager recovers. This parameter is available only to MRS 3.x or later.	The default value is the same as the value of akka.ask.timeout .	No

Parameter	Description	Default Value	Mandatory
high-availability.zookeeper.client.acl	Configure the ACL (open creator) of the ZooKeeper node. The ACL is automatically configured based on the security mode of the cluster.	<ul style="list-style-type: none"> Security mode: The default value is creator. Non-security mode: The default value is open. 	Yes
zookeeper.sasl.disable	Whether to enable SASL authentication. This parameter is automatically configured based on the security mode of the cluster.	<ul style="list-style-type: none"> Security mode: false Non-security mode: true 	Yes
zookeeper.sasl.service-name	<ul style="list-style-type: none"> If the ZooKeeper server configures a service whose name is different from ZooKeeper, this configuration item can be set. If service names on the client and server are inconsistent, authentication fails. 	zookeeper	Yes

Environment

Table 6-11 Environment parameters

Parameter	Description	Default Value	Mandatory
env.java.opts	JVM parameter, which is transferred to the startup script, JobManager, TaskManager, and Yarn client. For example, transfer remote debugging parameters.	-Xloggc:<LOG_DIR>/gc.log -XX:+PrintGCDetails -XX:-OmitStackTraceInFastThrow -XX:+PrintGCTimeStamps -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=20 -XX:GCLogFileSize=20M -Djdk.tls.ephemeralDHKeySize=2048 -Djava.library.path=\${HADOOP_COMMON_HOME}/lib/native -Djava.net.preferIPv4Stack=true -Djava.net.preferIPv6Addresses=false -Dbeetle.application.home.path=/opt/xxx/Bigdata/common/runtime/security/config	No

Yarn

Table 6-12 YARN parameters

Parameter	Description	Default Value	Mandatory
yarn.maximum-failed-containers	Maximum number of containers the system is going to reallocate in case of a container failure of TaskManager. The default value is the number of TaskManagers when the Flink cluster is started.	5	No

Parameter	Description	Default Value	Mandatory
yarn.application-attempts	Number of ApplicationMaster restarts. The value is the maximum value in the validity interval that is set to Akka's timeout in Flink. After the restart, the IP address and port number of ApplicationMaster will change and you will need to connect to the client manually.	2	No
yarn.heartbeat-delay	Time between heartbeats with the ApplicationMaster and Yarn ResourceManager in seconds. The unit is second.	5	No
yarn.containers.vcores	Number of virtual cores of each Yarn container	Number of TaskManager slots	No
yarn.application-master.port	ApplicationMaster port number setting. A port number range is supported.	32586-32650	No

Pipeline

It applies to MRS 3.x or later.

Table 6-13 Pipeline parameters

Parameter	Description	Default Value	Mandatory
nettyconnector.registerserver.topic.storage	Path (on a third-party server) to information about IP address, port numbers, and concurrency of NettySink. ZooKeeper is recommended for storage.	/flink/nettyconnector	No. However, if pipeline is enabled, the feature is mandatory.

Parameter	Description	Default Value	Mandatory
nettyconnector.sinkserver.port.range	Port range of NettySink.	28444-28843	No. However, if pipeline is enabled, the feature is mandatory.
nettyconnector.ssl.enabled	Whether SSL encryption for the communication between NettySink and NettySource is enabled. For details about the encryption key and protocol, see SSL .	false	No. However, if pipeline is enabled, the feature is mandatory.
nettyconnector.message.delimiter	Delimiter used to configure the message sent by NettySink to the NettySource, which is 2-4 bytes long, and cannot contain \n, #, or space.	The default value is \$_.	No. However, if pipeline is enabled, the feature is mandatory.

6.4 Configuring Flink Security Features

6.4.1 Security Features

Flink Authentication and Encryption

- In a Flink cluster, all components support authentication.
 - The kerberos authentication is supported between Flink cluster components and external components, such as YARN, HDFS, and ZooKeeper.
 - The security cookie authentication between Flink cluster components, for example, Flink client and JobManager, JobManager and TaskManager, and TaskManager and TaskManager, are supported.
- SSL encrypted transmission is supported by components of a Flink cluster and between components in a cluster, such as Flink client and JobManager, JobManager and TaskManager, as well as TaskManager and TaskManager.

For details, see [Authentication and Encryption](#).

ACL Control

In HA mode, ACL control is supported.

In HA mode of Flink, ZooKeeper can be used to manage clusters and discover services. Zookeeper supports SASL ACL control. Only users who have passed the SASL (Kerberos) authentication have the permission to operate files on ZooKeeper. To enable SASL ACL control, perform following configurations in the Flink configuration file.

```
high-availability.zookeeper.client.acl: creator  
zookeeper.sasl.disable: false
```

For details about configuration items, see [HA](#).

Web Security

Security of the Flink web UI is hardened, the whitelist filtering feature is provided, the Flink web UI is accessible only through a YARN proxy, and security headers are enhanced. In Flink clusters, listening ports of components can be configured.

- Encoding rules
 - Description: The same encoding mode is used on the web service client and server to prevent garbled characters and to implement input verification.
 - Security hardening: Response messages of web servers are encoded using UTF-8.
- Whitelist-based filter of IP addresses
 - Note: IP filter must be added to the web server to filter unauthorized requests from the source IP address and prevent unauthorized login.
 - Security hardening: Add **jobmanager.web.allow-access-address** to enable the IP filter. By default, only YARN users are supported.

NOTE

After the client is installed, you need to add the IP address of the client node to the **jobmanager.web.allow-access-address** configuration item.

- Preventing sending the absolute paths to the client
 - Note: If an absolute path is sent to a client, the directory structure of the server is exposed, increasing the risk that attackers know and attack the system.
 - Security hardening: If the Flink configuration file contains a parameter starting with a slash (/), the first-level directory is deleted.
- Same-origin policy

It applies to MRS 3.x or later.

 - If two URL protocols have same hosts and ports, they are of the same origin. Protocols of different origins cannot access each other, unless the source of the visitor is specified on the host of the service to be visited.
 - Security hardening: The default value of the header of the response header **Access-Control-Allow-Origin** is the IP address of

ResourceManager on Yarn clusters. If the IP address is not from Yarn, mutual access is not allowed.

- Preventing sensitive information disclosure
It applies to MRS 3.x or later.
 - Web pages containing sensitive data must not be cached, to avoid leakage of sensitive information or data crosstalk among users who visit the internet through the proxy server.
 - Security hardening: Add **Cache-control**, **Pragma**, **Expires** security header. The default value is **Cache-Control: no-store**, **Pragma: no-cache**, and **Expires: 0**. The security hardening stops contents interacted between Flink and web server from being cached.
- Anti-hijacking
It applies to MRS 3.x or later.
 - Since hotlinking and clickjacking use framing technologies, security hardening is required to prevent attacks.
 - Security hardening: Add **X-Frame-Options** security header to specify whether the browser will load the pages from **iframe**, **frame** or **object**. The default value is **X-Frame-Options: DENY**, indicating that no pages can be nested to **iframe**, **frame** or **object**.
- Logging calls of the Web Service APIs
It applies to MRS 3.x or later.
 - Calls of the **Flink webmonitor restful** APIs are logged.
 - The **jobmanager.web.accesslog.enable** can be added in the **access log**. The default value is **true**. Logs are stored in a separate **webaccess.log** file.
- Cross-site request forgery prevention
It applies to MRS 3.x or later.
 - In **Browser/Server** applications, CSRF must be prevented for operations involving server data modification, such as adding, modifying, and deleting. The CSRF forces end users to execute non-intended operations on the current web application.
 - Security hardening: Only two post APIs, one delete API, and get interfaces are reserve for modification requests. All other APIs are deleted.
- Troubleshooting:
It applies to MRS 3.x or later.
 - When the application is abnormal, exception information is filtered, logged, and returned to the client.
 - Security hardening: A default error message page to filter information and log detailed error information. Four configuration parameters are added to ensure that the error page is switched to a specified URL provided by FusionInsight, preventing exposure of unnecessary information.

Table 6-14 Parameters

Parameter	Description	Default Value	Mandatory
jobmanager.web.403-redirect-url	Web page access error 403. If 403 error occurs, the page switch to a specified page.	-	Yes
jobmanager.web.404-redirect-url	Web page access error 404. If 404 error occurs, the page switch to a specified page.	-	Yes
jobmanager.web.415-redirect-url	Web page access error 415. If 415 error occurs, the page switch to a specified page.	-	Yes
jobmanager.web.500-redirect-url	Web page access error 500. If 500 error occurs, the page switch to a specified page.	-	Yes

- HTML5 security

It applies to MRS 3.x or later.

- HTML5 is a next generation web development specification that provides new functions and extend the labels for developers. These new labels and functions increase the attack surface and pose attack risks (such as cross-domain resource sharing, client storage, WebWorker, WebRTC, and WebSocket).
- Security hardening: Add the **Access-Control-Allow-Origin** parameter. For example, if you want to enable the cross-domain resource sharing, configure the **Access-Control-Allow-Origin** parameter of the HTTP response header.

 **NOTE**

Flink does not involve security risks of functions such as storage on the client, WebWorker, WebRTC, and WebSocket.

Security Statement

- All security functions of Flink are provided by the open source community or self-developed. Security features that need to be configured by users, such as authentication and SSL encrypted transmission, may affect performance.
- As a big data computing and analysis platform, Flink does not detect sensitive information. Therefore, you need to ensure that the input data is not sensitive.
- You can evaluate whether configurations are secure as required.
- For any security-related problems, contact O&M support.

6.4.2 Authentication and Encryption

Security Authentication

Flink supports the following authentication modes:

- Kerberos authentication: It is used between the Flink Yarn client and Yarn ResourceManager, JobManager and ZooKeeper, JobManager and HDFS, TaskManager and HDFS, Kafka and TaskManager, as well as TaskManager and ZooKeeper.
- Security cookie authentication: Security cookie authentication is used between Flink Yarn client and JobManager, JobManager and TaskManager, as well as TaskManager and TaskManager.
- Internal authentication of Yarn: The Internal authentication mechanism of Yarn is used between Yarn ResourceManager and ApplicationMaster (AM).

 **NOTE**

- Flink JobManager and Yarn ApplicationMaster are in the same process.
- If Kerberos authentication is enabled for the user's cluster, Kerberos authentication is required.
- For versions earlier than MRS 3.x, Flink does not support security cookie authentication.

Table 6-15 Authentication modes

Authen- tication Mode	Descrip- tion	Configuration Method
Kerbero s authent ication	Current ly, only keytab authent ication mode is support ed.	<ol style="list-style-type: none"> 1. Download the user keytab file from FusionInsight Manager and save the keytab file to a folder on the host where the Flink client is located. 2. Configure the following parameters in the flink-conf.yaml file: <ol style="list-style-type: none"> a. Keytab path security.kerberos.login.keytab: /home/flinkuser/keytab/abc222.keytab Note: /home/flinkuser/keytab/abc222.keytab indicates the user directory. b. Principal name security.kerberos.login.principal: abc222 c. In HA mode, if ZooKeeper is configured, the Kerberos authentication configuration items must be configured as follows: zookeeper.sasl.disable: false security.kerberos.login.contexts: Client d. If you want to perform Kerberos authentication between Kafka client and Kafka broker, set the value as follows: security.kerberos.login.contexts: Client,KafkaClient

Authentication Mode	Description	Configuration Method
Security cookie authentication	-	<p>1. In the bin directory of the Flink client, run the generate_keystore.sh script to generate security cookie, flink.keystore, and flink.truststore. Run the sh generate_keystore.sh command and enter the user-defined password. The password cannot contain #.</p> <p>NOTE After the script is executed, the flink.keystore and flink.truststore files are generated in the conf directory on the Flink client. In the flink-conf.yaml file, default values are specified for following parameters:</p> <ul style="list-style-type: none"> • Set security.ssl.keystore to the absolute path of the flink.keystore file. • Set security.ssl.truststore to the absolute path of the flink.truststore file. • Set security.cookie to a random password automatically generated by the generate_keystore.sh script. • By default, security.ssl.encrypt.enabled: false is set in the flink-conf.yaml file by default. The generate_keystore.sh script sets security.ssl.key-password, security.ssl.keystore-password, and security.ssl.truststore-password to the password entered when the generate_keystore.sh script is called. <p>2. Check whether Security Cookie is enabled. That is, check security.enable: true in the flink-conf.yaml file and check whether security cookie is configured. An example is as follows:</p> <pre>security.cookie: ae70acc9-9795-4c48- ad35-8b5adc8071744f605d1d-2726-432e-88ae-dd39bfec40a9</pre> <p>NOTE The validity period of the SSL certificate obtained by using the generate_keystore.sh script preset on the MRS client is 5 years.</p> <p>To disable the default SSL authentication mode, set security.ssl.enabled to false in the flink-conf.yaml file and comment out security.ssl.key-password, security.ssl.keystore-password, security.ssl.keystore, security.ssl.truststore-password, and security.ssl.truststore.</p>

Authen tication Mode	Descrip tion	Configuration Method
Internal authent ication of Yarn	This authent ication mode does not need to be configu red by the user.	-

 **NOTE**

One Flink cluster supports only one user. One user can create multiple Flink clusters.

Encrypted Transmission

Flink supports three encrypted transmission modes:

- Encrypted transmission inside Yarn: It is used between the Flink Yarn client and Yarn ResourceManager, as well as Yarn ResourceManager and JobManager.
- SSL transmission: SSL transmission is used between Flink Yarn client and JobManager, JobManager and TaskManager, as well as TaskManager and TaskManager.
- Encrypted transmission inside Hadoop: The internal encrypted transmission mode of Hadoop used between JobManager and HDFS, TaskManager and HDFS, JobManager and ZooKeeper, as well as TaskManager and ZooKeeper.

 **NOTE**

Configuration about SSL encrypted transmission is mandatory while configuration about encryption of Yarn and Hadoop is not required.

To configure SSL encrypted transmission, configure the following parameters in the **flink-conf.yaml** file on the client:

1. Enable SSL and configure the SSL encryption algorithm. For MRS 3.x or later, see [Table 6-16](#). Modify the parameters as required.

Table 6-16 Parameter description

Parameter	Example Value	Description
security.ssl.enabled	true	Enable SSL.

Parameter	Example Value	Description
akka.ssl.enabled	true	Enable Akka SSL.
blob.service.ssl.enabled	true	Enable SSL for the Blob channel.
taskmanager.data.ssl.enabled	true	Enable SSL transmissions between TaskManagers.
security.ssl.algorithms	TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	Configure the SSL encryption algorithm.

For versions earlier than MRS 3.x, see [Table 6-17](#).

Table 6-17 Parameter description

Parameter	Example Value	Description
security.ssl.internal.enabled	true	Enable internal SSL.
akka.ssl.enabled	true	Enable Akka SSL.
blob.service.ssl.enabled	true	Enable SSL for the Blob channel.
taskmanager.data.ssl.enabled	true	Enable SSL transmissions between TaskManagers.
security.ssl.algorithms	TLS_RSA_WITH_AES128_CBC_SHA256	Configure the SSL encryption algorithm.

For versions earlier than MRS 3.x, the following parameters in [Table 6-18](#) do not exist in the default Flink configuration of MRS. If you want to enable SSL for external connections, add the following parameters. After SSL for external connection is enabled, the native Flink page cannot be accessed using a Yarn proxy, because the Yarn open-source version cannot process HTTPS requests using a proxy. However, you can create a Windows VM in the same VPC of the cluster and access the native Flink page from the VM.

Table 6-18 Parameter description

Parameter	Example Value	Description
security.ssl.rest.enabled	true	Enable external SSL. If this parameter is set to true , set the related parameters by referring to Table 6-18 .
security.ssl.rest.keystore	\${path}/flink.keystore	Path for storing the keystore .
security.ssl.rest.keystore-password	-	A user-defined password of keystore .
security.ssl.rest.key-password	-	A user-defined password of the SSL key.
security.ssl.rest.truststore	\${path}/flink.truststore	Path for storing the truststore .
security.ssl.rest.truststore-password	-	A user-defined password of truststore .

 **NOTE**

Enabling SSL for data transmission between TaskManagers may pose great impact on the system performance.

- In the **bin** directory of the Flink client, run the **sh generate_keystore.sh <password>** command. For details, see [Authentication and Encryption](#). The configuration items in [Table 6-19](#) are set by default for MRS 3.x or later. You can also configure them manually.

Table 6-19 Parameter description

Parameter	Example Value	Description
security.ssl.keystore	\${path}/flink.keystore	Path for storing the keystore . flink.keystore indicates the name of the keystore file generated by the generate_keystore.sh* tool.
security.ssl.keystore-password	-	A user-defined password of keystore .
security.ssl.key-password	-	A user-defined password of the SSL key.
security.ssl.truststore	\${path}/flink.truststore	Path for storing the truststore . flink.truststore indicates the name of the truststore file generated by the generate_keystore.sh* tool.

Parameter	Example Value	Description
security.ssl.truststore-password	-	A user-defined password of truststore .

For versions earlier than MRS 3.x, the *generate_keystore.sh* command is generated automatically, and the configuration items in [Table 6-20](#) are set by default. You can also configure them manually.

Table 6-20 Parameter description

Parameter	Example Value	Description
security.ssl.internal.keystore	\${path}/flink.keystore	Path for storing the keystore . flink.keystore indicates the name of the keystore file generated by the generate_keystore.sh* tool.
security.ssl.internal.keystore-password	-	A user-defined password of keystore .
security.ssl.internal.keystore-key-password	-	A user-defined password of the SSL key.
security.ssl.internal.truststore	\${path}/flink.truststore	Path for storing the truststore . flink.truststore indicates the name of the truststore file generated by the generate_keystore.sh* tool.
security.ssl.internal.truststore-password	-	A user-defined password of truststore .

For versions earlier than MRS 3.x, if SSL for external connections is enabled, that is, **security.ssl.rest.enabled** is set to **true**, you need to configure the parameters listed in [Table 6-21](#).

Table 6-21 Parameters

Parameter	Example Value	Description
security.ssl.rest.enabled	true	Enable external SSL. If this parameter is set to true , set the related parameters by referring to Table 6-21 .
security.ssl.rest.keystore	\${path}/flink.keystore	Path for storing the keystore .

Parameter	Example Value	Description
security.ssl.rest.keystore-password	-	A user-defined password of keystore .
security.ssl.rest.key-password	-	A user-defined password of the SSL key.
security.ssl.rest.truststore	\${path}/flink.truststore	Path for storing the truststore .
security.ssl.rest.truststore-password	-	A user-defined password of truststore .

 **NOTE**

The **path** directory is a user-defined directory for storing configuration files of the SSL keystore and truststore.

3. Configure the paths for the client to access the **keystore** or **truststore** file.
 - Relative path (recommended)

Perform the following steps to set the file paths of **flink.keystore** and **flink.truststore** to relative paths and ensure that the directory where the Flink client command is executed can directly access the relative paths.

 - i. In the **conf/** directory of the Flink client, create a directory, for example, **ssl**.


```
cd / Flink client directory/Flink/flink/conf/
mkdir ssl
```
 - ii. Move the **flink.keystore** and **flink.truststore** files to the new paths.


```
mv flink.keystore ssl/
mv flink.truststore ssl/
```
 - iii. Change the values of the following parameters to relative paths in the **flink-conf.yaml** file:


```
vi / Flink client directory/Flink/flink/conf/flink-conf.yaml
security.ssl.keystore: ssl/flink.keystore
security.ssl.truststore: ssl/flink.truststore
```
 - Absolute path

After the **generate_keystore.sh** script is executed, the **flink.keystore** and **flink.truststore** file paths are automatically set to absolute paths in the **flink-conf.yaml** file by default. In this case, you need to place the **flink.keystore** and **flink.truststore** files in the **conf** directory to the absolute paths of the Flink client and each Yarn node, respectively.

6.4.3 Configuring Kafka

Sample project data of Flink is stored in Kafka. A user with Kafka permission can send data to Kafka and receive data from it.

- Step 1** Ensure that clusters, including HDFS, Yarn, Flink, and Kafka are installed.

Step 2 Create a topic.

- Run Linux command line to create a topic. Before running commands, ensure that the kinit command, for example, **kinit flinkuser**, is run for authentication.

NOTE

To create a Flink user, you need to have the permission to create Kafka topics. The format of the command is shown as follows, in which **{zkQuorum}** indicates ZooKeeper cluster information and the format is *IP.port*, and **{Topic}** indicates the topic name.

bin/kafka-topics.sh --create --zookeeper {zkQuorum}/kafka --replication-factor 1 --partitions 5 --topic {Topic}

Assume the topic name is **topic 1**. The command for creating this topic is displayed as follows:

```
/opt/client/Kafka/kafka/bin/kafka-topics.sh --create --zookeeper
10.96.101.32:2181,10.96.101.251:2181,10.96.101.177:2181,10.91.8.160:2181/kafka --replication-factor
1 --partitions 5 --topic topic1
```

NOTE

The ZooKeeper cluster information is as follows:

- Service IP address of the ZooKeeper quorumpeer instance:
Log in to FusionInsight Manager and choose **Cluster > Services > ZooKeeper**. On the page that is displayed, click the **Instance** tab and view the service IP addresses of all nodes where the quorumpeer instances reside.
 - Port number of the ZooKeeper client:
Log in to FusionInsight Manager and choose **Cluster > Services > ZooKeeper**. On the page that is displayed, click the **Configurations** tab. On this tab page, view the value of **clientPort**.
- Configure the permission of the topic on the server.
Set the **allow.everyone.if.no.acl.found** parameter of Kafka Broker to **true**.

Step 3 Perform the security authentication.

The Kerberos authentication, SSL encryption authentication, or Kerberos + SSL authentication mode can be used.

NOTE

For versions earlier than MRS 3.x, only Kerberos authentication is supported.

- **Kerberos authentication**

- Client configuration

In the Flink configuration file **flink-conf.yaml**, add configurations about Kerberos authentication. For example, add **KafkaClient** in **contexts** as follows:

```
security.kerberos.login.keytab: /home/demo/keytab/flinkuser.keytab
security.kerberos.login.principal: flinkuser
security.kerberos.login.contexts: Client,KafkaClient
security.kerberos.login.use-ticket-cache: false
```

NOTE

For versions earlier than MRS 3.x, set **security.kerberos.login.keytab** to **/home/demo/flink/release/keytab/flinkuser.keytab**.

- Running parameter

Running parameters about the **SASL_PLAINTEXT** protocol are as follows:

```
--topic topic1 --bootstrap.servers 10.96.101.32:21007 --security.protocol SASL_PLAINTEXT --  
sas.l.kerberos.service.name kafka //10.96.101.32:21007 indicates the IP:port of the Kafka server.
```

- **SSL encryption**

- Configure the server.

Log in to FusionInsight Manager, choose **Cluster > Services > Kafka > Configurations**, and set **Type** to **All**. Search for **ssl.mode.enable** and set it to **true**.

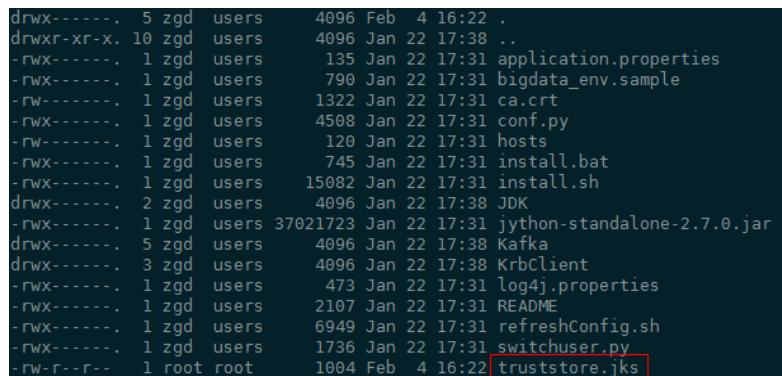
- Configure the client.

- Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Kafka > More > Download Client** to download Kafka client.
- Use the **ca.crt** certificate file in the client root directory to generate the **truststore** file for the client.

Run the following command:

```
keytool -noprompt -import -alias myservcert -file ca.crt -keystore truststore.jks
```

The command execution result is similar to the following:



```
dwx----- 5 zgd users 4096 Feb 4 16:22 .  
dwxr-xr-x 10 zgd users 4096 Jan 22 17:38 ..  
-rw----- 1 zgd users 135 Jan 22 17:31 application.properties  
-rw----- 1 zgd users 790 Jan 22 17:31 bigdata_env.sample  
-rw----- 1 zgd users 1322 Jan 22 17:31 ca.crt  
-rw----- 1 zgd users 4508 Jan 22 17:31 conf.py  
-rw----- 1 zgd users 120 Jan 22 17:31 hosts  
-rw----- 1 zgd users 745 Jan 22 17:31 install.bat  
-rw----- 1 zgd users 15082 Jan 22 17:31 install.sh  
dwx----- 2 zgd users 4096 Jan 22 17:38 JDK  
-rw----- 1 zgd users 37021723 Jan 22 17:31 jython-standalone-2.7.0.jar  
dwx----- 5 zgd users 4096 Jan 22 17:38 Kafka  
dwx----- 3 zgd users 4096 Jan 22 17:38 KrbClient  
-rw----- 1 zgd users 473 Jan 22 17:31 log4j.properties  
-rw----- 1 zgd users 2107 Jan 22 17:31 README  
-rw----- 1 zgd users 6949 Jan 22 17:31 refreshConfig.sh  
-rw----- 1 zgd users 1736 Jan 22 17:31 switchuser.py  
-rw-r--r-- 1 root root 1004 Feb 4 16:22 truststore.jks
```

- Run parameters.

The value of **ssl.truststore.password** must be the same as the password you entered when creating **truststore**. Run the following command to run parameters:

```
--topic topic1 --bootstrap.servers 10.96.101.32:9093 --security.protocol SSL --  
ssl.truststore.location /home/zgd/software/FusionInsight_Kafka_ClientConfig/truststore.jks  
--ssl.truststore.password XXX
```

- **Kerberos+SSL encryption**

After completing preceding configurations of the client and server of Kerberos and SSL, modify the port number and protocol type in running parameters to enable the Kerberos+SSL encryption mode.

```
--topic topic1 --bootstrap.servers 10.96.101.32:21009 --security.protocol SASL_SSL --  
sas.l.kerberos.service.name kafka --ssl.truststore.location /home/zgd/software/  
FusionInsight_Kafka_ClientConfig/truststore.jks --ssl.truststore.password XXX
```

----End

6.4.4 Configuring Pipeline

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

1. Configure files.
 - **nettyconnector.registerserver.topic.storage:** (Mandatory) Configures the path (on a third-party server) to information about IP address, port numbers, and concurrency of NettySink. For example:
nettyconnector.registerserver.topic.storage: /flink/nettyconnector
 - **nettyconnector.sinkserver.port.range:** (Mandatory) Configures the range of port numbers of NettySink. For example:
nettyconnector.sinkserver.port.range: 28444-28843
 - **nettyconnector.ssl.enabled:** Configures whether to enable SSL encryption between NettySink and NettySource. The default value is **false**. For example:
nettyconnector.ssl.enabled: true
2. Configure security authentication.
 - SASL authentication of ZooKeeper depends on the HA configuration in the **flink-conf.yaml** file.
 - SSL configurations such as keystore, truststore, keystore password, truststore password, and password inherit from **flink-conf.yaml**. For details, see [Encrypted Transmission](#).

6.5 Configuring and Developing a Flink Visualization Job

6.5.1 Introduction to Flink Web UI

Flink web UI provides a web-based visual development platform. You only need to compile SQL statements to develop jobs, slashing the job development threshold. In addition, the exposure of platform capabilities allows service personnel to compile SQL statements for job development to quickly respond to requirements, greatly reducing the Flink job development workload.

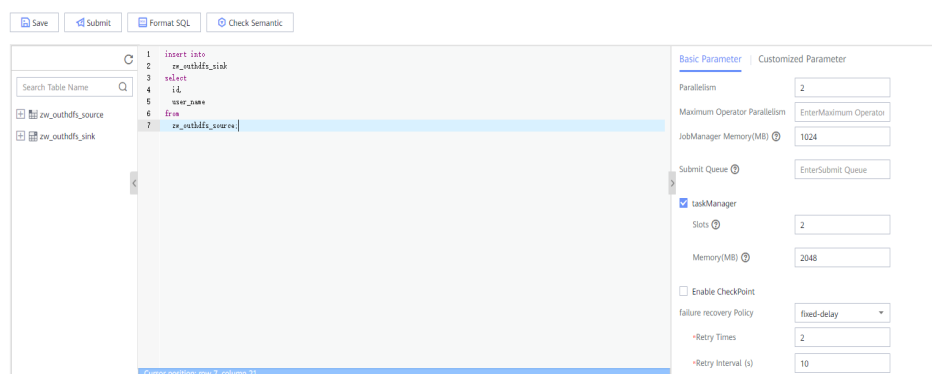
 **NOTE**

This section applies to only MRS 3.1.0 or later.

Flink Web UI Features

The Flink web UI has the following features:

- Enterprise-class visual O&M: GUI-based O&M management, job monitoring, and standardization of Flink SQL statements for job development.



- Quick cluster connection: After configuring the client and user credential key file, you can quickly access a cluster using the cluster connection function.
- Quick data connection: You can access a component by configuring the data connection function. If **Data Connection Type** is set to **HDFS**, you need to create a cluster connection. If **Authentication Mode** is set to **KERBEROS** for other data connection types, you need to create a cluster connection. If **Authentication Mode** is set to **SIMPLE**, you do not need to create a cluster connection.

 **NOTE**

If **Data Connection Type** is set to **Kafka**, **Authentication Type** cannot be set to **KERBEROS**.

- Visual development platform: The input/output mapping table can be customized to meet the requirements of different input sources and output destinations.
- Easy to use GUI-based job management

Job Management



Job Name	Type	Status	Kind	Description	Created by	Operation
uff	Flink SQL	 Saved	stream		flinkuser	Start Develop Stop Delete Edit More
test2	Flink Jar	 Stop	stream		flinkuser	Start Develop Stop Delete Edit More
test	Flink SQL	 Stop	stream		flinkuser	Start Develop Stop Delete Edit More

Key Web UI Capabilities

Table 6-22 shows the key capabilities provided by Flink web UI.

Table 6-22 Key web UI capabilities

Item	Description
Batch-Stream convergence	<ul style="list-style-type: none"> • Batch jobs and stream jobs can be processed with a unified set of Flink SQL statements.
Flink SQL kernel capabilities	<ul style="list-style-type: none"> • Flink SQL supports customized window size, stream compute within 24 hours, and batch processing beyond 24 hours. • Flink SQL supports reading data from Kafka and HDFS, writing data to Kafka and HDFS. • A job can define multiple Flink SQL jobs, and multiple metrics can be combined into one job for computing. If a job contains same primary keys as well as same inputs and outputs, the job supports the computing of multiple windows. • The AVG, SUM, COUNT, MAX, and MIN statistical methods are supported.

Item	Description
Flink SQL functions on the console	<ul style="list-style-type: none"> ● Cluster connection management allows you to configure clusters where services such as Kafka and HDFS reside. ● Data connection management allows you to configure services such as Kafka and HDFS. ● Data table management allows you to define data tables accessed by SQL statements and generate DDL statements. ● Flink SQL job definition allows you to verify, parse, optimize, convert a job into a Flink job, and submit the job for running based on the entered SQL statements.
Flink job visual management	<ul style="list-style-type: none"> ● Stream jobs and batch jobs can be defined in a visual manner. ● Job resources, fault recovery policies, and checkpoint policies can be configured in a visual manner. ● Status monitoring of stream and batch jobs are supported. ● The Flink job O&M is enhanced, including redirection of the native monitoring page.
Performance and reliability	<ul style="list-style-type: none"> ● Stream processing supports 24-hour window aggregation computing and millisecond-level performance. ● Batch processing supports 90-day window aggregation computing, which can be completed in minutes. ● Invalid data of stream processing and batch processing can be filtered out. ● When HDFS data is read, the data can be filtered based on the calculation period in advance. ● If the job definition platform is faulty or the service is degraded, jobs cannot be redefined, but the computing of existing jobs is not affected. ● The automatic restart mechanism is provided for job failures. You can configure restart policies.

Flink Web UI Application Process

The Flink web UI application process is shown as follows:

Figure 6-11 Flink web UI application process




Table 6-23 Description of the Flink web UI application process

Step	Description	Reference
Creating an application	Applications can be used to isolate different upper-layer services.	Creating an Application
Creating a cluster connection	Different clusters can be accessed by configuring the cluster connection.	Creating a Cluster Connection
Creating a Data Connection	Through data connections, you can access different data services, including HDFS and Kafka.	Creating a Data Connection
Creating a stream table	Data tables can be used to define basic attributes and parameters of source tables, dimension tables, and output tables.	Creating a Stream Table
Creating a SQL/JAR job (stream/batch job)	APIs can be used to define Flink jobs, including Flink SQL and Flink Jar jobs.	Creating a Job

Step	Description	Reference
Managing jobs	A created job can be managed, including starting, developing, stopping, deleting, and editing the job.	Creating a Job

6.5.2 Flink Web UI Permission Management

To access and use the Flink web UI to perform service operations, you need to assign FlinkServer-related permissions to users. User **admin** on FusionInsight Manager does not have FlinkServer service operation permissions.

Applications (tenants) in FlinkServer are the maximum management scope, including cluster connection management, data connection management, application management, stream table management, and job management.

There are three types of resource permissions for FlinkServer, as shown in [Table 6-24](#).

Table 6-24 FlinkServer resource permissions

Name	Description	Remarks
FlinkServer administrator permission	Users who have the permission can edit and view all applications.	This is the highest-level permission of FlinkServer. If you have the FlinkServer administrator permission, you have the permission on all applications by default.
Application edit permission	Users who have the permission can create, edit, and delete cluster connections and data connections. They can also create stream tables as well as create and run jobs.	In addition, users who have the permission can view current applications.
Application view permission	Users who have the permission can view applications.	-

6.5.3 Creating a FlinkServer Role

Create and configure a FlinkServer role on Manager as an MRS cluster administrator. A FlinkServer role can be configured with the FlinkServer administrator permission and the permissions to edit and view applications.

You need to set permissions for the specified user in FlinkServer so that they can update, query, and delete data.

Prerequisites

The cluster administrator has planned permissions based on service requirements.

Procedure

Step 1 Log in to Manager.

Step 2 Choose **System > Permission > Role**.

Step 3 On the displayed page, click **Create Role** and specify **Role Name** and **Description**.

Step 4 Set **Configure Resource Permission**.

FlinkServer permissions are as follows:

- **FlinkServer Admin Privilege:** highest-level permission. Users with the permission can perform service operations on all FlinkServer applications.
- **FlinkServer Application:** Users can set **application view** and **applications management** permissions on applications.

Table 6-25 Setting a role

Scenario	Role Authorization
Setting the FlinkServer administrator permission	In Configure Resource Permission , choose <i>Name of the desired cluster</i> > Flink and select FlinkServer Admin Privilege .
Setting FlinkServer application permissions	In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Flink > FlinkServer Application . In the Permission column, select application view or applications management .

Step 5 Click **OK**. Return to role management page.

Step 6 (Optional) Create a user with FlinkServer-related permissions.

After the FlinkServer role is created, create a FlinkServer user and bind the user to the configured FlinkServer role and user group. For details, see [Creating a User](#).

----End

6.5.4 Accessing the Flink Web UI

Scenario

After Flink is installed in an MRS cluster, you can connect to clusters and data as well as manage stream tables and jobs using the Flink web UI.

This section describes how to access the Flink web UI in an MRS cluster.

NOTE

You are advised to use Google Chrome 50 or later to access the Flink web UI. The Internet Explorer may be incompatible with the Flink web UI.

Impact on the System

Site trust must be added to the browser when you access Manager and the Flink web UI for the first time. Otherwise, the Flink web UI cannot be accessed.

Procedure

Step 1 Log in to FusionInsight Manager as a user with **FlinkServer Admin Privilege**. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Services > Flink**.

Step 2 On the right of **Flink WebUI**, click the link to access the Flink web UI.

The Flink web UI provides the following functions:

- System management:
 - Cluster connection management allows you to create, view, edit, test, and delete a cluster connection.
 - Data connection management allows you to create, view, edit, test, and delete a data connection. Data connection types include HDFS and Kafka.
 - Application management allows you to create, view, and delete an application.
- Stream table management allows you to create, view, edit, and delete a stream table.
- Job management allows you to create, view, start, develop, edit, stop, and delete a job.

----End

6.5.5 Creating an Application

Scenario

Applications can be used to isolate different upper-layer services.

Creating an Application

Step 1 Access the Flink web UI as a user with **FlinkServer Admin Privilege**. For details, see [Accessing the Flink Web UI](#).

Step 2 Choose **System Management > Application Management**.

Step 3 Click **Create Application**. On the displayed page, set parameters and click **OK**.

After the application is created, you can switch to the application to be operated in the upper left corner of the Flink web UI and develop jobs.

----End

6.5.6 Creating a Cluster Connection

Scenario

Different clusters can be accessed by configuring the cluster connection.

Creating a Cluster Connection

- Step 1** Access the Flink web UI. For details, see [Accessing the Flink Web UI](#).
- Step 2** Choose **System Management > Cluster Connection Management**. The **Cluster Connection Management** page is displayed.
- Step 3** Click **Create Cluster Connection**. On the displayed page, set parameters by referring to [Table 6-26](#) and click **OK**. After the cluster connection is created, you can edit, test, or delete the cluster connection in the **Operation** column.

Table 6-26 Parameters for creating a cluster connection

Parameter	Description
Cluster Connection Name	Enter the name of the cluster connection.
Description	Enter the description of the cluster connection name.
FusionInsight HD Version	Set a cluster version.
Secure Version	<ul style="list-style-type: none"> If the secure version is used, select Yes for a security cluster. Enter the username and upload the user credential. If not, select No.
Username	<p>The user must have the minimum permissions for accessing services in the cluster.</p> <p>This parameter is available only when Secure Version is set to Yes.</p>
Client Profile	Client profile of the cluster, in TAR format.
User Credential	<p>User authentication credential in FusionInsight Manager in TAR format.</p> <p>This parameter is available only when Secure Version is set to Yes.</p> <p>Files can be uploaded only after the username is entered.</p>

NOTE

To obtain the cluster client configuration files, perform the following steps:

1. Log in to FusionInsight Manager and choose **Cluster > Dashboard**.
2. Choose **More > Download Client > Configuration Files Only**, select a platform type, and click **OK**.

To obtain the user credential, perform the following steps:

1. Log in to FusionInsight Manager and click **System**.
2. In the **Operation** column of the user, choose **More > Download Authentication Credential**, select a cluster, and click **OK**.

----End

6.5.7 Creating a Data Connection

Scenario

You can use data connections to access different data services. Currently, FlinkServer supports HDFS and Kafka data connections.

Creating a Data Connection

- Step 1** Access the Flink web UI. For details, see [Accessing the Flink Web UI](#).
- Step 2** Choose **System Management > Data Connection Management**. The **Data Connection Management** page is displayed.
- Step 3** Click **Create Data Connection**. On the displayed page, select a data connection type, enter information by referring to [Table 6-27](#), and click **OK**. After the data connection is created, you can edit, test, or delete the data connection in the **Operation** column.

Table 6-27 Parameters for creating a data connection

Parameter	Description	Example Value
Data Connection Type	Type of the data connection, which can be HDFS or Kafka .	-
Data Connection Name	Name of the data connection.	-
Cluster Connection	Cluster connection name in configuration management. This parameter is mandatory for HDFS data connections.	-
Kafka broker	Connection information about Kafka broker instances. The format is <i>IP address.Port number</i> . Use commas (,) to separate multiple instances. This parameter is mandatory for Kafka data connections.	192.168.0.1:21005,192.168.0.2:21005
Authentication Mode	<ul style="list-style-type: none"> • SIMPLE: indicates that the connected service is in non-security mode and does not need to be authenticated. • KERBEROS: indicates that the connected service is in security mode and the Kerberos protocol for security authentication is used for authentication. 	-

----End

6.5.8 Creating a Stream Table

Scenario

Data tables can be used to define basic attributes and parameters of source tables, dimension tables, and output tables.

Creating a Stream Table

- Step 1** Access the Flink web UI. For details, see [Accessing the Flink Web UI](#).
- Step 2** Click **Table Management**. The table management page is displayed.
- Step 3** Click **Create Stream Table**. On the stream table creation page, set parameters by referring to [Table 6-28](#) and click **OK**. After the stream table is created, you can edit or delete the stream table in the **Operation** column.

Table 6-28 Parameters for creating a stream table

Parameter	Description	Remarks
Stream/ Table Name	Stream/Table name	Example: flink_sink
Description	Stream/Table description information	-
Mapping Table Type	Flink SQL does not provide the data storage function. Table creation is actually the creation of mapping for external data tables or storage. The value can be Kafka or HDFS .	-
Type	Includes data source table Source and data result table Sink . Tables included in different mapping table types are as follows: <ul style="list-style-type: none"> ● Kafka: Source and Sink ● HDFS: Source and Sink 	-
Data Connection	Name of the data connection	-
Topic	Kafka topic to be read. Multiple Kafka topics can be read. Use separators to separate topics. This parameter is available when Mapping Table Type is set to Kafka .	-
File Path	HDFS directory or a single file path to be transferred. This parameter is available when Mapping Table Type is set to HDFS .	Example: /user/sqoop/ or /user/sqoop/example.csv

Parameter	Description	Remarks
Code	Codes corresponding to different mapping table types are as follows: <ul style="list-style-type: none"> • Kafka: CSV and JSON • HDFS: CSV 	-
Prefix	When Mapping Table Type is set to Kafka , Type is set to Source , and Code is set to JSON , this parameter indicates the hierarchical prefixes of multi-layer nested JSON, which are separated by commas (,).	For example, data,info indicates that the content under data and info in the nested JSON file is used as the data input in JSON format.
Separator	Has different meanings when Mapping Table Type is set to the following values: It is used as the separator of specified CSV fields. This parameter is available only when Code is set to CSV .	Example: comma (,)
Row Separator	Line break in the file, including \r , \n , and \r\n . This parameter is available when Mapping Table Type is set to HDFS .	-
Column Separator	Field separator in the file. This parameter is available when Mapping Table Type is set to HDFS .	Example: comma (,)
Stream Table Structure	Stream/Table structure, including Name and Type .	-
Proctime	System time, which is irrelevant to the data timestamp. That is, the time when the calculation is complete in Flink operators. This parameter is available when Type is set to Source .	-
Event Time	Time when an event is generated, that is, the timestamp generated during data generation. This parameter is available when Type is set to Source .	-

----End

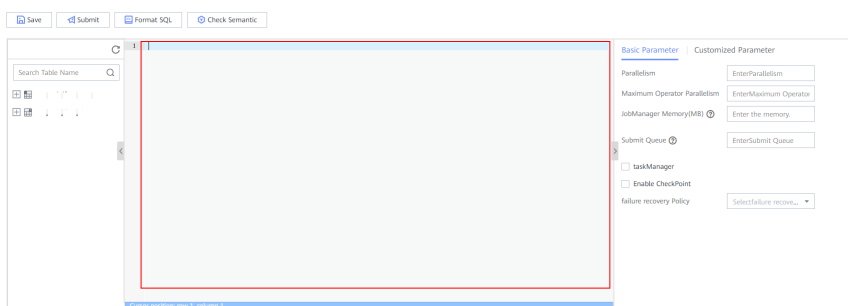
6.5.9 Creating a Job

Scenario

Define Flink jobs, including Flink SQL and Flink JAR jobs.

Creating a Job

- Step 1** Access the Flink web UI. For details, see [Accessing the Flink Web UI](#).
- Step 2** Click **Job Management**. The job management page is displayed.
- Step 3** Click **Create Job**. Create a Flink SQL job or Flink Jar job, enter job information, and click **OK**. The job is created and the job development page is displayed.
- Step 4** (Optional) To develop a job immediately, configure the job on the job development page.
 - Creating a Flink SQL job
 - a. Develop the job on the job development page.



- b. Click **Check Semantic** to check the input content and click **Format SQL** to format SQL statements.
- c. After the job SQL statements are developed, set basic and customized parameters as required by referring to [Table 6-29](#) and click **Save**.

Table 6-29 Basic parameters

Parameter	Description
Parallelism	Number of parallel jobs
Maximum Operator Parallelism	Maximum degree of parallelism of operators
JobManager Memory (MB)	Memory of JobManager The minimum value is 512 .
Submit Queue	Queue to which a job is submitted. If this parameter is not set, the default queue is used.

Parameter	Description
taskManager	<p>taskManager running parameters include:</p> <ul style="list-style-type: none"> ▪ Slots: If this parameter is left blank, the default value 1 is used. ▪ Memory (MB): The minimum value is 512.
Enable CheckPoint	<p>Whether to enable CheckPoint. After CheckPoint is enabled, you need to configure the following information:</p> <ul style="list-style-type: none"> ▪ Time Interval (ms): This parameter is mandatory. ▪ Mode: This parameter is mandatory. The options are EXACTLY_ONCE and AT_LEAST_ONCE. ▪ Minimum Interval (ms): The minimum value is 10. ▪ Timeout Duration: The minimum value is 10. ▪ Maximum Parallelism: The value must be a positive integer containing a maximum of 64 characters. ▪ Whether to clean up: This parameter can be set to Yes or No. ▪ Whether to enable incremental checkpoints: This parameter can be set to Yes or No.
Failure Recovery Policy	<p>Failure recovery policy of a job. The options are as follows. For details, see Flink Restart Policy.</p> <ul style="list-style-type: none"> ▪ fixed-delay: You need to configure Retry Times and Retry Interval (s). ▪ failure-rate: You need to configure Max Retry Times, Interval (min), and Retry Interval (s). ▪ none

- d. Click **Submit** in the upper left corner to submit the job.
- Creating a Flink JAR job
 - a. Click **Select** to upload a local JAR file and set parameters by referring to [Table 6-30](#) or add customized parameters.

Table 6-30 Parameter configuration

Parameter	Description
Local .jar File	Upload a local JAR file. The size of the file cannot exceed 10 MB.
Main Class	Main-Class type. <ul style="list-style-type: none"> ▪ Default: By default, the class name is specified based on the Mainfest file in the JAR file. ▪ Specify: Manually specify the class name.
Type	Class name. This parameter is available when Main Class is set to Specify .
Class Parameter	Class parameters of Main-Class (parameters are separated by spaces).
Parallelism	Number of parallel jobs
JobManager Memory (MB)	Memory of JobManager The minimum value is 512 .
Submit Queue	Queue to which a job is submitted. If this parameter is not set, the default queue is used.
taskManager	taskManager running parameters include: <ul style="list-style-type: none"> ▪ Slots: If this parameter is left blank, the default value 1 is used. ▪ Memory (MB): The minimum value is 512.

b. Click **Save** to save the configuration and click **Submit** to submit the job.

Step 5 Return to the job management page. You can view information about the created job, including job name, type, status, kind, and description.

After a job is created, you can start, develop, stop, edit, and delete the job, view job details, and rectify checkpoint faults in the **Operation** column of the job.

 **NOTE**

- To read files related to the submitted job on the node as another user, ensure that the user and the user who submitted the job belong to the same user group and the user has been assigned the FlinkServer application management role. For example, **application view** is selected by referring to [Creating a FlinkServer Role](#).
- You can view details about jobs in the **Running** state.
- You can rectify checkpoint faults for jobs in the **Running failed**, **Running succeeded**, or **Stop** state.

----End

6.6 Flink Log Overview

Log Description

Log path:

- Run logs of a Flink job: `${BIGDATA_DATA_HOME}/hadoop/data${i}/nm/containerlogs/application_${appid}/container_${Scontid}`

NOTE

The logs of executing tasks are stored in the preceding path. After the execution is complete, the Yarn configuration determines whether these logs are gathered to the HDFS directory.

- FlinkResource run logs: `/var/log/Bigdata/flink/flinkResource`

Log archive rules:

1. FlinkResource run logs:
 - By default, service logs are backed up each time when the log size reaches 20 MB. A maximum of 20 logs can be reserved without being compressed.

NOTE

For versions earlier than MRS 3.x, The executor logs are backed up each time when the log size reaches 30 MB. A maximum of 20 logs can be reserved without being compressed.

- You can set the log size and number of compressed logs on the Manager page or modify the corresponding configuration items in **log4j-cli.properties**, **log4j.properties**, and **log4j-session.properties** in *Client installation directory/Flink/flink/conf/* on the client.

Table 6-31 FlinkResource log list

Type	Name	Description
FlinkResource run logs	checkService.log	Health check log
	kinit.log	Initialization log
	postinstall.log	Service installation log
	prestart.log	Prestart script log
	start.log	Startup log

2. FlinkServer service logs and audit logs.
 - By default, FlinkServer service logs and audit logs are backed up each time when the log size reaches 100 MB. The service logs are stored for a maximum of 30 days, and audit logs are stored for a maximum of 90 days.
 - You can set the log size and number of compressed logs on the Manager page or modify the corresponding configuration items in **log4j-**

cli.properties, **log4j.properties**, and **log4j-session.properties** in *Client installation directory/Flink/flink/conf/* on the client.

Table 6-32 FlinkServer log list

Type	Name	Description
FlinkServer run logs	checkService.log	Health check log
	checkFlinkServer.log	Health check log of FlinkServer
	localhost_access_log..yyyy-mm-dd.txt	URL log of FlinkServer
	start_thrift_server.out	Thrift server startup log
	thrift_server_thriftServer_xxx.log.last	
	cleanup.log	Cleanup log file for instance installation and uninstallation
	flink-omm-client-IP.log	Job startup log
	flinkserver_yyyymmdd-x.log.gz	Service archive log
	flinkserver.log	Service log
	flinkserver---pidxxx-gc.log.x.current	GC log
	kinit.log	Initialization log
	postinstall.log	Service installation log
	prestart.log	Prestart script log
	start.log	Startup log
	stop.log	Stop log
	catalina.yyyy-mm-dd.log	Tomcat run log
	catalina.out	
	host-manager.yyyy-mm-dd.log	
	localhost.yyyy-mm-dd.log	
	manager.yyyy-mm-dd.log	
FlinkServer audit logs	flinkserver_audit_yyyymmdd-x.log.gz	Audit archive log
	flinkserver_audit.log	Audit log

Log Level

Table 6-33 describes the log levels supported by Flink. The priorities of log levels are ERROR, WARN, INFO, and DEBUG in descending order. Logs whose levels are higher than or equal to the specified level are printed. The number of printed logs decreases as the specified log level increases.

Table 6-33 Log levels

Level	Description
ERROR	Error information about the current event processing
WARN	Exception information about the current event processing
INFO	Normal running status information about the system and events
DEBUG	System information and system debugging information

To modify log levels, perform the following steps:

- Step 1** Go to the **All Configurations** page of Flink by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.

----End

NOTE

- After the configuration is complete, you do not need to restart the service. Download the client again for the configuration to take effect.
- You can also change the configuration items corresponding to the log level in **log4j-cli.properties**, **log4j.properties**, and **log4j-session.properties** in *Client installation directory*/**Flink/flink/conf/** on the client.
- When a job is submitted using a client, a log file is generated in the **log** folder on the client. The default umask value is **0022**. Therefore, the default log permission is **644**. To change the file permission, you need to change the umask value. For example, to change the umask value of user **omm**:
 - Add **umask 0026** to the end of the **/home/omm/.baskrc** file.
 - Run the **source /home/omm/.baskrc** command to make the file permission take effect.

Log Format

Table 6-34 Log formats

Type	Format	Example
Run log	<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>	2019-06-27 21:30:31,778 INFO [flink-akka.actor.default-dispatcher-3] TaskManager container_e10_1498290698388_0004_02_0000 07 has started. org.apache.flink.yarn.YarnFlinkResourceManager (FlinkResourceManager.java:368)

6.7 Flink Performance Tuning

6.7.1 Memory Configuration Optimization

Scenarios

The computing of Flink depends on memory. If the memory is insufficient, the performance of Flink will be greatly deteriorated. One solution is to monitor garbage collection (GC) to evaluate the memory usage. If the memory becomes the performance bottleneck, optimize the memory usage according to the actual situation.

If **Full GC** is frequently reported in the Container GC on the Yarn that monitors the node processes, the GC needs to be optimized.

NOTE

In the `env.java.opts` configuration item of the `conf/flink-conf.yaml` file on the client, add the `-Xloggc:<LOG_DIR>/gc.log -XX:+PrintGCDetails -XX:-OmitStackTraceInFastThrow -XX:+PrintGCTimeStamps -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=20 -XX:GCLogFileSize=20M` parameter. The GC log is configured by default.

Procedure

- Optimize GC.
Adjust the ratio of tenured generation memory to young generation memory. In the `conf/flink-conf.yaml` configuration file on the client, add the `-XX:NewRatio` parameter to the `env.java.opts` configuration item. For example, `-XX:NewRatio=2` indicates that ratio of tenured generation memory to young generation memory is 2:1, that is, the young generation memory occupies one third and tenured generation memory occupies two thirds.
- When developing Flink applications, optimize the partitioning or grouping operation of `DataStream`.

- If partitioning causes data skew, partitions need to be optimized.
- Do not perform concurrent operations, because some operations, WindowAll for example, to DataStream do not support parallelism.
- Do not use set keyBy to string type.

6.7.2 Configuring DOP

Scenario

The degree of parallelism (DOP) indicates the number of tasks to be executed concurrently. It determines the number of data blocks after the operation. Configuring the DOP will optimize the number of tasks, data volume of each task, and the host processing capability.

Query the CPU and memory usage. If data and tasks are not evenly distributed among nodes, increase the DOP for even distribution.

Procedure

Configure the DOP at one of the following layers (the priorities of which are in the descending order) based on the actual memory, CPU, data, and application logic conditions:

- Operator

Call the **setParallelism()** method to specify the DOP of an operator, data source, and sink. For example:

```
final StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();

DataStream<String> text = [...];
DataStream<Tuple2<String, Integer>> wordCounts = text
    .flatMap(new LineSplitter())
    .keyBy(0)
    .timeWindow(Time.seconds(5))
    .sum(1).setParallelism(5);

wordCounts.print();

env.execute("Word Count Example");
```

- Execution environment

Flink runs in the execution environment which defines a default DOP for operators, data source and data sink.

Call the **setParallelism()** method to specify the default DOP of the execution environment. Example:

```
final StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
env.setParallelism(3);
DataStream<String> text = [...];
DataStream<Tuple2<String, Integer>> wordCounts = [...];
wordCounts.print();
env.execute("Word Count Example");
```

- Client

Specify the DOP when submitting jobs to Flink on the client. If you use the CLI client, specify the DOP using the **-p** parameter. Example:

```
./bin/flink run -p 10 ../examples/*WordCount-java*.jar
```

- System

On the Flink client, modify the **parallelism.default** parameter in the **flink-conf.yaml** file under the conf to specify the DOP for all execution environments.

6.7.3 Configuring Process Parameters

Scenario

In Flink on Yarn mode, there are JobManagers and TaskManagers. JobManagers and TaskManagers schedule and run tasks.

Therefore, configuring parameters of JobManagers and TaskManagers can optimize the execution performance of a Flink application. Perform the following steps to optimize the Flink cluster performance.

Procedure

Step 1 Configure JobManager memory.

JobManagers are responsible for task scheduling and message communications between TaskManagers and ResourceManagers. JobManager memory needs to be increased as the number of tasks and the DOP increases.

JobManager memory needs to be configured based on the number of tasks.

- When running the **yarn-session** command, add the **-jm MEM** parameter to configure the memory.
- When running the **yarn-cluster** command, add the **-yjm MEM** parameter to configure the memory.

Step 2 Configure the number of TaskManagers.

Each core of a TaskManager can run a task at the same time. Increasing the number of TaskManagers has the same effect as increasing the DOP. Therefore, you can increase the number of TaskManagers to improve efficiency when there are sufficient resources.

Step 3 Configure the number of TaskManager slots.

Multiple cores of a TaskManager can process multiple tasks at the same time. This has the same effect as increasing the DOP. However, the balance between the number of cores and the memory must be maintained, because all cores of a TaskManager share the memory.

- When running the **yarn-session** command, add the **-s NUM** parameter to configure the number of slots.
- When running the **yarn-cluster** command, add the **-ys NUM** parameter to configure the number of slots.

Step 4 Configure TaskManager memory.

TaskManager memory is used for task execution and communication. A large-size task requires more resources. In this case, you can increase the memory.

- When running the **yarn-session** command, add the **-tm MEM** parameter to configure the memory.

- When running the **yarn-cluster** command, add the **-ytm MEM** parameter to configure the memory.

----End

6.7.4 Optimizing the Design of Partitioning Method

Scenarios

The divide of tasks can be optimized by optimizing the partitioning method. If data skew occurs in a certain task, the whole execution process is delayed. Therefore, when designing the partitioning method, ensure that partitions are evenly assigned.

Procedure

Partitioning methods are as follows:

- **Random partitioning:** randomly partitions data.
`dataStream.shuffle();`
- **Rebalancing (round-robin partitioning):** evenly partitions data based on round-robin. The partitioning method is useful to optimize data with data skew.
`dataStream.rebalance();`
- **Rescaling:** assign data to downstream subsets in the form of round-robin. The partitioning method is useful if you want to deliver data from each parallel instance of a data source to subsets of some mappers without the using `rebalance ()`, that is, the complete rebalance operation.
`dataStream.rescale();`
- **Broadcast:** broadcast data to all partitions.
`dataStream.broadcast();`
- **User-defined partitioning:** use a user-defined partitioner to select a target task for each element. The user-defined partitioning allows user to partition data based on a certain feature to achieve optimized task execution.

The following is an example:

```
// fromElements builds simple Tuple2 stream
DataStream<Tuple2<String, Integer>> dataStream = env.fromElements(Tuple2.of("hello",1),
Tuple2.of("test",2), Tuple2.of("world",100));

// Defines the key value used for partitioning. Adding one to the value equals to the id.
Partitioner<Tuple2<String, Integer>> strPartitioner = new Partitioner<Tuple2<String, Integer>>() {
    @Override
    public int partition(Tuple2<String, Integer> key, int numPartitions) {
        return (key.f0.length() + key.f1) % numPartitions;
    }
};

// The Tuple2 data is used as the basis for partitioning.

dataStream.partitionCustom(strPartitioner, new KeySelector<Tuple2<String, Integer>, Tuple2<String,
Integer>>() {
    @Override
    public Tuple2<String, Integer> getKey(Tuple2<String, Integer> value) throws Exception {
        return value;
    }
}).print();
```

6.7.5 Configuring the Netty Network Communication

Scenarios

The communication of Flink is based on Netty network. The network performance determines the data switching speed and task execution efficiency. Therefore, the performance of Flink can be optimized by optimizing the Netty network.

Procedure

In the **conf/flink-conf.yaml** file on the client, change configurations as required. Exercise caution when changing default values, because default values are optimal.

- **taskmanager.network.netty.num-arenas**: Specifies the number of arenas of Netty. The default value is **taskmanager.numberOfTaskSlots**.
- **taskmanager.network.netty.server.numThreads** and **taskmanager.network.netty.client.numThreads**: specify the number of threads on the client and server. The default value is **taskmanager.numberOfTaskSlots**.
- **taskmanager.network.netty.client.connectTimeoutSec**: specifies the timeout interval for connection of TaskManager client. The default value is **120s**.
- **taskmanager.network.netty.sendReceiveBufferSize**: specifies the buffer size of the Netty network. The default value is the buffer size (cat /proc/sys/net/ipv4/tcp_[rw]mem) of the system and the value is usually 4 MB.
- **taskmanager.network.netty.transport**: specifies the transmission method of the Netty network. The default value is **nio**. The value can only be **nio** and **epoll**.

6.7.6 Experience Summary

Avoiding Data Skew

If data skew occurs (certain data volume is extremely large), the execution time of tasks is inconsistent even though no GC is performed.

- Redefine keys. Use keys of smaller granularity to optimize the task size.
- Modify the DOP.
- Call the rebalance operation to balance data partitions.

Setting Timeout Interval for the Buffer

- During the execution of tasks, data is exchanged through network. You can set the **setBufferTimeout** parameter to specify a buffer timeout interval for data exchanging among different servers.
- If **setBufferTimeout** is set to **-1**, the refreshing operation is performed when the buffer is full to maximize the throughput. If **setBufferTimeout** is set to **0**, the refreshing operation is performed each time data is received to minimize the delay. If **setBufferTimeout** is set to a value greater than **0**, the refreshing operation is performed after the buffer times out.

The following is an example:

```
env.setTimeout(timeoutMillis);  
env.generateSequence(1,10).map(new MyMapper()).setTimeout(timeoutMillis);
```

6.8 Common Flink Shell Commands

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

Before running the Flink shell commands, perform the following steps:

- Step 1** Install the Flink client in `/opt/client`.
- Step 2** Run the following command to initialize environment variables:
`source /opt/client/bigdata_env`
- Step 3** If Kerberos authentication has been enabled for the cluster, configure client authentication by referring to [Step 5](#). If Kerberos authentication is disabled, skip this step.
- Step 4** Run the related commands according to [Table 6-35](#).

Table 6-35 Flink Shell commands

Command	Description	Description
yarn-session.sh	<p>-at,--applicationType <arg>: Defines the Yarn application type.</p> <p>-D <property=value>: Configures dynamic parameter.</p> <p>-d,--detached: Disables the interactive mode and starts a separate Flink Yarn session.</p> <p>-h,--help: Displays the help information about the Yarn session CLI.</p> <p>-id,--applicationId <arg>: Binds to a running Yarn session.</p> <p>-j,--jar <arg>: Sets the path of the user's JAR file.</p> <p>-jm,--jobManagerMemory <arg>: Sets the JobManager memory.</p> <p>-m,--jobmanager <arg>: Address of the JobManager (master) to which to connect. Use this parameter to connect to a specified JobManager.</p> <p>-nl,--nodeLabel <arg>: Specifies the nodeLabel of the Yarn application.</p> <p>-nm,--name <arg>: Customizes a name for the application on Yarn.</p> <p>-q,--query: Queries available Yarn resources.</p> <p>-qu,--queue <arg>: Specifies a Yarn queue.</p> <p>-s,--slots <arg>: Sets the number of slots for each TaskManager.</p> <p>-t,--ship <arg>: specifies the directory of the file to be sent.</p> <p>-tm,--taskManagerMemory <arg>: sets the TaskManager memory.</p> <p>-yd,--yarn detached: starts Yarn in the detached mode.</p> <p>-z,--zookeeperNamespace <args>: specifies the namespace of ZooKeeper.</p> <p>-h: Gets help information.</p>	Start a resident Flink cluster to receive tasks from the Flink client.

Command	Description	Description
flink run	<p>-c,--class <classname>: Specifies a class as the entry for running programs.</p> <p>-C,--classpath <url>: Specifies classpath.</p> <p>-d,--detached: Runs a job in the detached mode.</p> <p>-files,--dependencyFiles <arg>: File on which the Flink program depends.</p> <p>-n,--allowNonRestoredState: A state that cannot be restored can be skipped during restoration from a snapshot point in time. For example, if an operator in the program is deleted, you need to add this parameter when restoring the snapshot point.</p> <p>-m,--jobmanager <host:port>: Specifies the JobManager.</p> <p>-p,--parallelism <parallelism>: Specifies the job DOP, which will overwrite the DOP parameter in the configuration file.</p> <p>-q,--sysoutLogging: Disables the function of outputting Flink logs to the console.</p> <p>-s,--fromSavepoint <savepointPath>: Specifies a savepoint path for recovering jobs.</p> <p>-z,--zookeeperNamespace <zookeeperNamespace>: specifies the namespace of ZooKeeper.</p> <p>-yat,--yarnapplicationType <arg>: Defines the Yarn application type.</p> <p>-yD <arg>: Dynamic parameter configuration.</p> <p>-yd,--yarn detached: Starts Yarn in the detached mode.</p> <p>-yh,--yarnhelp: Obtains the Yarn help.</p> <p>-yid,--yarnapplicationId <arg>: Binds a job to a Yarn session.</p> <p>-yj,--yarnjar <arg>: Sets the path to Flink jar file.</p> <p>-yjm,--yarnjobManagerMemory <arg>: Sets the JobManager memory (MB).</p> <p>-ynm,--yarnname <arg>: Customizes a name for the application on Yarn.</p> <p>-yq,--yarnquery: Queries available Yarn resources (memory and CPUs).</p>	<p>Submit a Flink job.</p> <ol style="list-style-type: none"> 1. The -y* parameter is used in the yarn-cluster mode. 2. If the parameter is not -y*, you need to run the yarn-session command to start the Flink cluster before running this command to submit a task.

Command	Description	Description
	<p>-yqu,--yarnqueue <arg>: Specifies a Yarn queue.</p> <p>-ys,--yarnslots: Sets the number of slots for each TaskManager.</p> <p>-yt,--yarnship <arg>: Specifies the path of the file to be sent.</p> <p>-ytm,--yarntaskManagerMemory <arg>: Sets the TaskManager memory (MB).</p> <p>-yz,--yarnzookeeperNamespace <arg>: Specifies the namespace of ZooKeeper. The value must be the same as the value of yarn-session.sh -z.</p> <p>-h: Gets help information.</p>	
flink info	<p>-c,--class <classname>: Specifies a class as the entry for running programs.</p> <p>-p,--parallelism <parallelism>: Specifies the DOP for running programs.</p> <p>-h: Gets help information.</p>	Display the execution plan (JSON) of the running program.
flink list	<p>-a,--all: displays all jobs.</p> <p>-m,--jobmanager <host:port>: specifies the JobManager.</p> <p>-r,--running: displays only jobs in the running state.</p> <p>-s,--scheduled: displays only jobs in the scheduled state.</p> <p>-z,--zookeeperNamespace <zookeeperNamespace>: specifies the namespace of ZooKeeper.</p> <p>-yid,--yarnapplicationId <arg>: binds a job to a Yarn session.</p> <p>-h: gets help information.</p>	Query running programs in the cluster.

Command	Description	Description
flink stop	<p>-d,--drain: sends MAX_WATERMARK before the savepoint is triggered and the job is stopped.</p> <p>-p,--savepointPath <savepointPath>: path for storing savepoints. The default value is state.savepoints.dir.</p> <p>-m,--jobmanager <host:port>: specifies the JobManager.</p> <p>-z,--zookeeperNamespace <zookeeperNamespace>: specifies the namespace of ZooKeeper.</p> <p>-yid,--yarnapplicationId <arg>: binds a job to a Yarn session.</p> <p>-h: gets help information.</p>	<p>Forcibly stop a running job (only streaming jobs are supported).</p> <p>StoppableFunction needs to be implemented on the source side in service code).</p>
flink cancel	<p>-m,--jobmanager <host:port>: specifies the JobManager.</p> <p>-s,--withSavepoint <targetDirectory>: triggers a savepoint when a job is canceled. The default directory is state.savepoints.dir.</p> <p>-z,--zookeeperNamespace <zookeeperNamespace>: specifies the namespace of ZooKeeper.</p> <p>-yid,--yarnapplicationId <arg>: binds a job to a Yarn session.</p> <p>-h: gets help information.</p>	<p>Cancel a running job.</p>
flink savepoint	<p>-d,--dispose <arg>: specifies a directory for storing the savepoint.</p> <p>-m,--jobmanager <host:port>: specifies the JobManager.</p> <p>-z,--zookeeperNamespace <zookeeperNamespace>: specifies the namespace of ZooKeeper.</p> <p>-yid,--yarnapplicationId <arg>: binds a job to a Yarn session.</p> <p>-h: gets help information.</p>	<p>Trigger a savepoint.</p>

Command	Description	Description
<pre>source Client installation directory/ bigdata_en v</pre>	None	<p>Import client environment variables.</p> <p>Restriction: If the user uses a custom script (for example, A.sh) and runs this command in the script, variables cannot be imported to the A.sh script. If variables need to be imported to the custom script A.sh, the user needs to use the secondary calling method.</p> <p>For example, first call the B.sh script in the A.sh script, and then run this command in the B.sh script. Parameters can be imported to the A.sh script but cannot be imported to the B.sh script.</p>
<pre>start-scala- shell.sh</pre>	local remote <host> <port> yarn: running mode	Start the scala shell.
<pre>sh generate_ke ystore.sh</pre>	-	<p>Run the generate_keystore.sh script to generate security cookie, flink.keystore, and flink.truststore. You need to enter a user-defined password that does not contain number signs (#).</p>

----End

6.9 Reference

6.9.1 Example of Issuing a Certificate

Generate the `generate_keystore.sh` script based on the sample code and save the script to the `bin` directory on the Flink client.

```
#!/bin/bash

KEYTOOL=${JAVA_HOME}/bin/keytool
KEystorePATH="$FLINK_HOME/conf/"
CA_ALIAS="ca"
CA_KEystore_NAME="ca.keystore"
CA_DNAME="CN=Flink_CA"
CA_KEYALG="RSA"
CLIENT_CONF_YAML="$FLINK_HOME/conf/flink-conf.yaml"
KEYTABPRINCEPAL=""

function getConf()
{
    if [ $# -ne 2 ]; then
        echo "invalid parameters for getConf"
        exit 1
    fi

    confName="$1"
    if [ -z "$confName" ]; then
        echo "conf name is empty."
        exit 2
    fi

    configFile=$FLINK_HOME/conf/client.properties
    if [ ! -f $configFile ]; then
        echo "$configFile" is not exist."
        exit 3
    fi

    defaultValue="$2"
    cnt=$(grep $1 $configFile | wc -l)
    if [ $cnt -gt 1 ]; then
        echo "$confName" has multi values in "$configFile"
        exit 4
    elif [ $cnt -lt 1 ]; then
        echo $defaultValue
    else
        line=$(grep $1 $configFile)
        confValue=$(echo "${line#*=}")
        echo "$confValue"
    fi
}

function createSelfSignedCA()
{
    #variable from user input
    keystorePath=$1
    storepassValue=$2
    keypassValue=$3

    #generate ca keystore
    rm -rf $keystorePath/$CA_KEystore_NAME
    $KEYTOOL -genkeypair -alias $CA_ALIAS -keystore $keystorePath/$CA_KEystore_NAME -dname
    $CA_DNAME -storepass $storepassValue -keypass $keypassValue -validity 3650 -keyalg $CA_KEYALG -
    keysize 3072 -ext bc=ca:true
    if [ $? -ne 0 ]; then
        echo "generate ca.keystore failed."
        exit 1
    fi

    #generate ca.cer
    rm -rf "$keystorePath/ca.cer"
    $KEYTOOL -keystore "$keystorePath/$CA_KEystore_NAME" -storepass "$storepassValue" -alias
```

```

$CA_ALIAS -validity 3650 -exportcert > "$keystorePath/ca.cer"
  if [ $? -ne 0 ]; then
    echo "generate ca.cer failed."
    exit 1
  fi

  #generate ca.truststore
  rm -rf "$keystorePath/flink.truststore"
  $KEYTOOL -importcert -keystore "$keystorePath/flink.truststore" -alias $CA_ALIAS -storepass
"$storepassValue" -noprompt -file "$keystorePath/ca.cer"
  if [ $? -ne 0 ]; then
    echo "generate ca.truststore failed."
    exit 1
  fi
}

function generateKeystore()
{
  #get path/pass from input
  keystorePath=$1
  storepassValue=$2
  keypassValue=$3

  #get value from conf
  aliasValue=$(getConf "flink.keystore.rsa.alias" "flink")
  validityValue=$(getConf "flink.keystore.rsa.validity" "3650")
  keyalgValue=$(getConf "flink.keystore.rsa.keyalg" "RSA")
  dnameValue=$(getConf "flink.keystore.rsa.dname" "CN=flink.com")
  SANValue=$(getConf "flink.keystore.rsa.ext" "ip:127.0.0.1")
  SANValue=$(echo "$SANValue" | xargs)
  SANValue="ip:$(echo "$SANValue" | sed 's/,/,ip:/g')"

  #generate keystore
  rm -rf $keystorePath/flink.keystore
  $KEYTOOL -genkeypair -alias $aliasValue -keystore $keystorePath/flink.keystore -dname $dnameValue -
ext SAN=$SANValue -storepass $storepassValue -keypass $keypassValue -keyalg $keyalgValue -keysize
3072 -validity 3650
  if [ $? -ne 0 ]; then
    echo "generate flink.keystore failed."
    exit 1
  fi

  #generate cer
  rm -rf $keystorePath/flink.csr
  $KEYTOOL -certreq -keystore $keystorePath/flink.keystore -storepass $storepassValue -alias $aliasValue -
file $keystorePath/flink.csr
  if [ $? -ne 0 ]; then
    echo "generate flink.csr failed."
    exit 1
  fi

  #generate flink.cer
  rm -rf $keystorePath/flink.cer
  $KEYTOOL -gencert -keystore $keystorePath/ca.keystore -storepass $storepassValue -alias $CA_ALIAS -
ext SAN=$SANValue -infile $keystorePath/flink.csr -outfile $keystorePath/flink.cer -validity 3650
  if [ $? -ne 0 ]; then
    echo "generate flink.cer failed."
    exit 1
  fi

  #import cer into keystore
  $KEYTOOL -importcert -keystore $keystorePath/flink.keystore -storepass $storepassValue -file
$keystorePath/ca.cer -alias $CA_ALIAS -noprompt
  if [ $? -ne 0 ]; then
    echo "importcert ca."
    exit 1
  fi

  $KEYTOOL -importcert -keystore $keystorePath/flink.keystore -storepass $storepassValue -file

```

```

$keystorePath/flink.cer -alias $aliasValue -noprompt;
    if [ $? -ne 0 ]; then
        echo "generate flink.truststore failed."
        exit 1
    fi
}

function configureFlinkConf()
{
    # set config
    if [ -f "$CLIENT_CONF_YAML" ]; then
        SSL_ENCRYPT_ENABLED=$(grep "security.ssl.encrypt.enabled" "$CLIENT_CONF_YAML" | awk '{print
$2}')
        if [ "$SSL_ENCRYPT_ENABLED" = "false" ];then

            sed -i s/"security.ssl.key-password:.*"/"security.ssl.key-password:"\ "${keyPass}"/g
"$CLIENT_CONF_YAML"
            if [ $? -ne 0 ]; then
                echo "set security.ssl.key-password failed."
                return 1
            fi

            sed -i s/"security.ssl.keystore-password:.*"/"security.ssl.keystore-password:"\ "${storePass}"/g
"$CLIENT_CONF_YAML"
            if [ $? -ne 0 ]; then
                echo "set security.ssl.keystore-password failed."
                return 1
            fi

            sed -i s/"security.ssl.truststore-password:.*"/"security.ssl.truststore-password:"\ "${storePass}"/g
"$CLIENT_CONF_YAML"
            if [ $? -ne 0 ]; then
                echo "set security.ssl.keystore-password failed."
                return 1
            fi

            echo "security.ssl.encrypt.enabled is false, set security.ssl.key-password security.ssl.keystore-
password security.ssl.truststore-password success."
        else
            echo "security.ssl.encrypt.enabled is true, please enter security.ssl.key-password security.ssl.keystore-
password security.ssl.truststore-password encrypted value in flink-conf.yaml."
        fi

        keystoreFilePath="${keystorePath}/flink.keystore
sed -i 's#"security.ssl.keystore:.*#"security.ssl.keystore:"\ "${keystoreFilePath}"#g'
"$CLIENT_CONF_YAML"
        if [ $? -ne 0 ]; then
            echo "set security.ssl.keystore failed."
            return 1
        fi

        truststoreFilePath="${keystorePath}/flink.truststore"
sed -i 's#"security.ssl.truststore:.*#"security.ssl.truststore:"\ "${truststoreFilePath}"#g'
"$CLIENT_CONF_YAML"
        if [ $? -ne 0 ]; then
            echo "set security.ssl.truststore failed."
            return 1
        fi

        command -v sha256sum >/dev/null
        if [ $? -ne 0 ];then
            echo "sha256sum is not exist, it will produce security.cookie with date +%F-%H-%M-%s-%N."
            cookie=$(date +%F-%H-%M-%s-%N)
        else
            cookie=$(echo "${KEYTABPRINCEPAL}" | sha256sum | awk '{print $1}')
        fi

        sed -i s/"security.cookie:.*"/"security.cookie:"\ "${cookie}"/g "$CLIENT_CONF_YAML"

```

```
        if [ $? -ne 0 ]; then
            echo "set security.cookie failed."
            return 1
        fi
    fi
    return 0;
}

main()
{
    #check environment variable is set or not
    if [ -z ${FLINK_HOME+x} ]; then
        echo "errro: environment variables are not set."
        exit 1
    fi
    stty -echo
    read -rp "Enter password:" password
    stty echo
    echo

    KEYTABPRINCEPAL=$(grep "security.kerberos.login.principal" "$CLIENT_CONF_YAML" | awk '{print $2}')
    if [ -z "$KEYTABPRINCEPAL" ];then
        echo "please config security.kerberos.login.principal info first."
        exit 1
    fi

    #get input
    keystorePath="$KEYSTOREPATH"
    storePass="$password"
    keyPass="$password"

    #generate self signed CA
    createSelfSignedCA "$keystorePath" "$storePass" "$keyPass"
    if [ $? -ne 0 ]; then
        echo "create self signed ca failed."
        exit 1
    fi

    #generate keystore
    generateKeystore "$keystorePath" "$storePass" "$keyPass"
    if [ $? -ne 0 ]; then
        echo "create keystore failed."
        exit 1
    fi

    echo "generate keystore/truststore success."

    # set flink config
    configureFlinkConf "$keystorePath" "$storePass" "$keyPass"
    if [ $? -ne 0 ]; then
        echo "configure Flink failed."
        exit 1
    fi

    return 0;
}

#the start main
main "$@"

exit 0
```

 NOTE

Run the **sh generate_keystore.sh** *<password>* command. *<password>* is user-defined.

- If *<password>* contains the special character \$, use the following method to avoid the password being escaped: **sh generate_keystore.sh 'Bigdata_2013'**.
- The password cannot contain #.
- Before using the **generate_keystore.sh** script, run the **source bigdata_env** command in the client directory.
- When the **generate_keystore.sh** script is used, the absolute paths of **security.ssl.keystore** and **security.ssl.truststore** are automatically filled in **flink-conf.yaml**. Therefore, you need to manually change the paths to relative paths as required. Example:
 - Change **/opt/client/Flink/flink/conf//flink.keystore** to **security.ssl.keystore: ssl/flink.keystore**.
 - Change **/opt/client/Flink/flink/conf//flink.truststore** to **security.ssl.truststore: ssl/flink.truststore**.
 - Create the **ssl** folder in any directory on the Flink client. For example, create the **ssl** folder in the **/opt/client/Flink/flink/conf/** directory and save the **flink.keystore** and **flink.truststore** files to the **ssl** folder.
 - When you run the **yarn-session** or **flink run -m yarn-cluster** command, run the **yarn-session.sh -t ssl -d** or **flink run -m yarn-cluster -yt ssl -d WordCount.jar** command in the same directory as the **ssl** folder.

6.10 Flink Restart Policy

Overview

Flink supports different restart policies to control whether and how to restart a job when a fault occurs. If no restart policy is specified, the cluster uses the default restart policy. You can also specify a restart policy when submitting a job. For details about how to configure such a policy on the job development page of MRS 3.1.0 or later, see [Creating a Job](#).

The restart policy can be specified by configuring the **restart-strategy** parameter in the Flink configuration file *Client installation directory*/**Flink/flink/conf/flink-conf.yaml** or can be dynamically specified in the application code. The configuration takes effect globally. Restart policies include **failure-rate** and the following two default policies:

- **No restart:** If CheckPoint is not enabled, this policy is used by default.
- **Fixed-delay:** If CheckPoint is enabled but no restart policy is configured, this policy is used by default.

No restart Policy

When a fault occurs, the job fails and does not attempt to restart.

Configure the parameter as follows:

```
restart-strategy: none
```

fixed-delay Policy

When a fault occurs, the job attempts to restart for a fixed number of times. If the number of attempts exceeds the times you specified, the job fails. The restart policy waits for a fixed period of time between two consecutive restart attempts.

In the following example, a job fails if the job attempts to restart for three times at an interval of 10 seconds. Configure the parameters as follows:

```
restart-strategy: fixed-delay
restart-strategy.fixed-delay.attempts: 3
restart-strategy.fixed-delay.delay: 10 s
```

failure-rate Policy

When a job fails, the job restarts directly. If the failure rate exceeds the value you configured, the job is considered as failed. The restart policy waits for a fixed period of time between two consecutive restart attempts.

In the following example, a job is considered as failed if the job attempts to restart for three times at an interval of 10 minutes. Configure the parameters as follows:

```
restart-strategy: failure-rate
restart-strategy.failure-rate.max-failures-per-interval: 3
restart-strategy.failure-rate.failure-rate-interval: 10 min
restart-strategy.failure-rate.delay: 10 s
```

Choosing a Restart Policy

- If you do not want to retry a failed job, select the **No restart** policy.
- To retry a failed job, select the **failure-rate** policy. If the fixed-delay policy is used, the number of job failures may reach the maximum number of retries due to hardware faults such as network and memory faults. As a result, the job fails.

To prevent repeated restarts when the failure-rate policy is used, configure parameters as follows:

```
restart-strategy: failure-rate
restart-strategy.failure-rate.max-failures-per-interval: 3
restart-strategy.failure-rate.failure-rate-interval: 10 min
restart-strategy.failure-rate.delay: 10 s
```

7 Using Flume

7.1 Using Flume from Scratch

Scenario

You can use Flume to import collected log information to Kafka.

Prerequisites

- A streaming cluster that contains components such as Flume and Kafka and has Kerberos authentication enabled has been created.
- The streaming cluster can properly communicate with the node where logs are generated.

Using the Flume Client (Versions Earlier Than MRS 3.x)

 NOTE

You do not need to perform [Step 2](#) to [Step 6](#) for a normal cluster.

Step 1 Install the Flume client.

Install the Flume client in a directory, for example, `/opt/Flumeclient`, on the node where logs are generated by referring to [Installing the Flume Client on Clusters of Versions Earlier Than MRS 3.x](#). The Flume client installation directories in the following steps are only examples. Change them to the actual installation directories.

Step 2 Copy the configuration file of the authentication server from the Master1 node to the *Flume client installation directory*/`fusioninsight-flume-Flume component version number/conf` directory on the node where the Flume client is installed.

``${BIGDATA_HOME}/MRS_Current/1_X_KerberosClient/etc/kdc.conf` is used as the full file path.

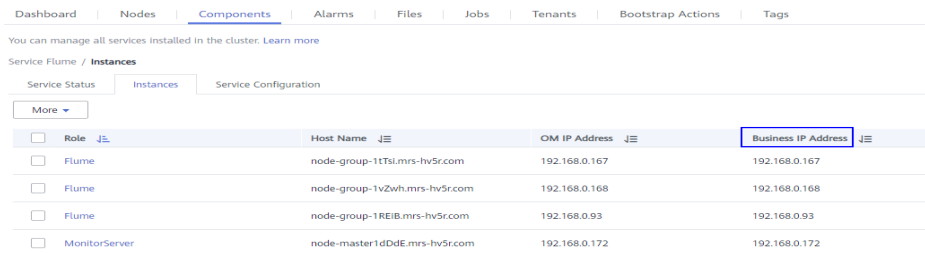
In the preceding paths, **X** indicates a random number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 3 Check the service IP address of any node where the Flume role is deployed.

Log in to the cluster details page, choose *Name of the desired cluster* > **Components** > **Flume** > **Instances**, and check the service IP address of any node where the Flume role is deployed.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)



Role	Host Name	OM IP Address	Business IP Address
Flume	node-group-1Tsl.mrs-hv5r.com	192.168.0.167	192.168.0.167
Flume	node-group-1vZvh.mrs-hv5r.com	192.168.0.168	192.168.0.168
Flume	node-group-1REIB.mrs-hv5r.com	192.168.0.93	192.168.0.93
MonitorServer	node-master1dDdE.mrs-hv5r.com	192.168.0.172	192.168.0.172

Step 4 Copy the user authentication file from this node to the *Flume client installation directory/fusioninsight-flume-Flume component version number/conf* directory on the Flume client node.

$\${BIGDATA_HOME}/MRS_XXX/install/FusionInsight-Flume-Flume\ component\ version\ number/flume/conf/flume.keytab$ is used as the full file path.

In the preceding paths, **XXX** indicates the product version number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 5 Copy the **jaas.conf** file from this node to the **conf** directory on the Flume client node.

$\${BIGDATA_HOME}/MRS_Current/1_X_Flume/etc/jaas.conf$ is used as the full file path.

In the preceding path, **X** indicates a random number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 6 Log in to the Flume client node and go to the client installation directory. Run the following command to modify the file:

```
vi conf/jaas.conf
```

Change the full path of the user authentication file defined by **keyTab** to the **Flume client installation directory/fusioninsight-flume-Flume component version number/conf** saved in **Step 4**, and save the modification and exit.

Step 7 Run the following command to modify the **flume-env.sh** configuration file of the Flume client:

```
vi Flume client installation directory/fusioninsight-flume-Flume component version number/conf/flume-env.sh
```

Add the following information after **-XX:+UseCMSCompactAtFullCollection**:

```
-Djava.security.krb5.conf=Flume client installation directory/fusioninsight-flume-1.9.0/conf/kdc.conf -  
Djava.security.auth.login.config=Flume client installation directory/fusioninsight-flume-1.9.0/conf/jaas.conf -  
Dzookeeper.request.timeout=120000
```

Example: "**-XX:+UseCMSCompactAtFullCollection -
Djava.security.krb5.conf=/opt/FlumeClient/fusioninsight-flume-Flume
component version number/conf/kdc.conf -
Djava.security.auth.login.config=/opt/FlumeClient/fusioninsight-flume-Flume
component version number/conf/jaas.conf -
Dzookeeper.request.timeout=120000**"

Change *Flume client installation directory* to the actual installation directory. Then save and exit.

Step 8 Run the following command to restart the Flume client:

```
cd Flume client installation directory/fusioninsight-flume-Flume component  
version number/bin
```

```
./flume-manage.sh restart
```

Example:

```
cd /opt/FlumeClient/fusioninsight-flume-Flume component version number/bin
```

```
./flume-manage.sh restart
```

Step 9 Run the following command to configure and save jobs in the Flume client configuration file **properties.properties** based on service requirements.

```
vi Flume client installation directory/fusioninsight-flume-Flume component  
version number/conf/properties.properties
```

The following uses SpoolDir Source+File Channel+Kafka Sink as an example:

```
#####  
#####  
client.sources = static_log_source  
client.channels = static_log_channel  
client.sinks = kafka_sink  
#####  
#####  
#LOG_TO_HDFS_ONLINE_1  
  
client.sources.static_log_source.type = spooldir  
client.sources.static_log_source.spoolDir = Monitoring directory  
client.sources.static_log_source.fileSuffix = .COMPLETED  
client.sources.static_log_source.ignorePattern = ^$  
client.sources.static_log_source.trackerDir = Metadata storage path during transmission  
client.sources.static_log_source.maxBlobLength = 16384  
client.sources.static_log_source.batchSize = 51200  
client.sources.static_log_source.inputCharset = UTF-8  
client.sources.static_log_source.deserializer = LINE  
client.sources.static_log_source.selector.type = replicating  
client.sources.static_log_source.fileHeaderKey = file  
client.sources.static_log_source.fileHeader = false  
client.sources.static_log_source.basenameHeader = true  
client.sources.static_log_source.basenameHeaderKey = basename  
client.sources.static_log_source.deletePolicy = never  
  
client.channels.static_log_channel.type = file  
client.channels.static_log_channel.dataDirs = Data cache path. Multiple paths, separated by commas (,), can  
be configured to improve performance.  
client.channels.static_log_channel.checkpointDir = Checkpoint storage path  
client.channels.static_log_channel.maxFileSize = 2146435071  
client.channels.static_log_channel.capacity = 1000000
```

```
client.channels.static_log_channel.transactionCapacity = 612000
client.channels.static_log_channel.minimumRequiredSpace = 524288000

client.sinks.kafka_sink.type = org.apache.flume.sink.kafka.KafkaSink
client.sinks.kafka_sink.kafka.topic = Topic to which data is written, for example, flume_test
client.sinks.kafka_sink.kafka.bootstrap.servers = XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number
client.sinks.kafka_sink.flumeBatchSize = 1000
client.sinks.kafka_sink.kafka.producer.type = sync
client.sinks.kafka_sink.kafka.security.protocol = SASL_PLAINTEXT
client.sinks.kafka_sink.kafka.kerberos.domain.name = Kafka domain name. This parameter is mandatory for a security cluster, for example, hadoop.xxx.com.
client.sinks.kafka_sink.requiredAcks = 0

client.sources.static_log_source.channels = static_log_channel
client.sinks.kafka_sink.channel = static_log_channel
```

 NOTE

- **client.sinks.kafka_sink.kafka.topic:** Topic to which data is written. If the topic does not exist in Kafka, it is automatically created by default.
- **client.sinks.kafka_sink.kafka.bootstrap.servers:** List of Kafka Brokers, which are separated by commas (.). By default, the port is **21007** for a security cluster and **9092** for a normal cluster.
- **client.sinks.kafka_sink.kafka.security.protocol:** The value is **SASL_PLAINTEXT** for a security cluster and **PLAINTEXT** for a normal cluster.
- **client.sinks.kafka_sink.kafka.kerberos.domain.name:**
You do not need to set this parameter for a normal cluster. For a security cluster, the value of this parameter is the value of **kerberos.domain.name** in the Kafka cluster.
Obtain the value by checking **\${BIGDATA_HOME}/MRS_Current/1_X_Broker/etc/server.properties** on the node where the broker instance resides.
In the preceding paths, **X** indicates a random number. Change it based on site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 10 After the parameters are set and saved, the Flume client automatically loads the content configured in **properties.properties**. When new log files are generated by `spoolDir`, the files are sent to Kafka producers and can be consumed by Kafka consumers.

----End

Using the Flume Client (MRS 3.x or Later)

 NOTE

You do not need to perform [Step 2](#) to [Step 6](#) for a normal cluster.

Step 1 Install the Flume client.

Install the Flume client in a directory, for example, `/opt/Flumeclient`, on the node where logs are generated by referring to [Installing the Flume Client on MRS 3.x or Later Clusters](#). The Flume client installation directories in the following steps are only examples. Change them to the actual installation directories.

Step 2 Copy the configuration file of the authentication server from the Master1 node to the `Flume client installation directory/fusioninsight-flume-Flume component version number/conf` directory on the node where the Flume client is installed.

The full file path is `${BIGDATA_HOME}/FusionInsight_BASE_XXX/1_X_KerberosClient/etc/kdc.conf`. In the preceding path, **XXX** indicates the

product version number. **X** indicates a random number. Replace them based on site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 3 Check the service IP address of any node where the Flume role is deployed.

Log in to FusionInsight Manager, click **Cluster**, click **Services**, and click Flume. On the displayed page, click **Instance**. Check the service IP address of any node where the Flume role is deployed.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 4 Copy the user authentication file from this node to the *Flume client installation directory*/**fusioninsight-flume-Flume component version number/conf** directory on the Flume client node.

The full file path is **`${BIGDATA_HOME}/FusionInsight_Porter_XXX/install/FusionInsight-Flume-Flume component version number/flume/conf/flume.keytab`**.

In the preceding paths, **XXX** indicates the product version number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 5 Copy the **jaas.conf** file from this node to the **conf** directory on the Flume client node.

The full file path is **`${BIGDATA_HOME}/FusionInsight_Current/1_X_Flume/etc/jaas.conf`**.

In the preceding path, **X** indicates a random number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 6 Log in to the Flume client node and go to the client installation directory. Run the following command to modify the file:

```
vi conf/jaas.conf
```

Change the full path of the user authentication file defined by **keyTab** to the **Flume client installation directory**/**fusioninsight-flume-Flume component version number/conf** saved in [Step 4](#), and save the modification and exit.

Step 7 Run the following command to modify the **flume-env.sh** configuration file of the Flume client:

```
vi Flume client installation directory/fusioninsight-flume-Flume component version number/conf/flume-env.sh
```

Add the following information after **-XX:+UseCMSCompactAtFullCollection**:

```
-Djava.security.krb5.conf=Flume client installation directory/fusioninsight-flume-1.9.0/conf/kdc.conf -  
Djava.security.auth.login.config=Flume client installation directory/fusioninsight-flume-1.9.0/conf/jaas.conf -  
Dzookeeper.request.timeout=120000
```

Example: **"-XX:+UseCMSCompactAtFullCollection -
Djava.security.krb5.conf=/opt/FlumeClient/fusioninsight-flume-Flume**

```
component version number/conf/kdc.conf -  
Djava.security.auth.login.config=/opt/FlumeClient/fusioninsight-flume-Flume  
component version number/conf/jaas.conf -  
Dzookeeper.request.timeout=120000"
```

Change *Flume client installation directory* to the actual installation directory. Then save and exit.

Step 8 Run the following command to restart the Flume client:

```
cd Flume client installation directory/fusioninsight-flume-Flume component  
version number/bin
```

```
./flume-manage.sh restart
```

Example:

```
cd /opt/FlumeClient/fusioninsight-flume-Flume component version number/bin  
./flume-manage.sh restart
```

Step 9 Configure jobs based on actual service scenarios.

- Some parameters, for MRS 3.x or later, can be configured on Manager.
- Set the parameters in the **properties.properties** file. The following uses SpoolDir Source+File Channel+Kafka Sink as an example.

Run the following command on the node where the Flume client is installed. Configure and save jobs in the Flume client configuration file **properties.properties** based on actual service requirements.

```
vi Flume client installation directory/fusioninsight-flume-Flume component  
version number/conf/properties.properties
```

```
#####  
#####  
client.sources = static_log_source  
client.channels = static_log_channel  
client.sinks = kafka_sink  
#####  
#####  
#LOG_TO_HDFS_ONLINE_1  
  
client.sources.static_log_source.type = spooldir  
client.sources.static_log_source.spoolDir = Monitoring directory  
client.sources.static_log_source.fileSuffix = .COMPLETED  
client.sources.static_log_source.ignorePattern = ^$  
client.sources.static_log_source.trackerDir = Metadata storage path during transmission  
client.sources.static_log_source.maxBlobLength = 16384  
client.sources.static_log_source.batchSize = 51200  
client.sources.static_log_source.inputCharset = UTF-8  
client.sources.static_log_source.deserializer = LINE  
client.sources.static_log_source.selector.type = replicating  
client.sources.static_log_source.fileHeaderKey = file  
client.sources.static_log_source.fileHeader = false  
client.sources.static_log_source.basenameHeader = true  
client.sources.static_log_source.basenameHeaderKey = basename  
client.sources.static_log_source.deletePolicy = never  
  
client.channels.static_log_channel.type = file  
client.channels.static_log_channel.dataDirs = Data cache path. Multiple paths, separated by commas  
(,), can be configured to improve performance.  
client.channels.static_log_channel.checkpointDir = Checkpoint storage path  
client.channels.static_log_channel.maxFileSize = 2146435071  
client.channels.static_log_channel.capacity = 1000000  
client.channels.static_log_channel.transactionCapacity = 612000
```

```
client.channels.static_log_channel.minimumRequiredSpace = 524288000

client.sinks.kafka_sink.type = org.apache.flume.sink.kafka.KafkaSink
client.sinks.kafka_sink.kafka.topic = Topic to which data is written, for example, flume_test
client.sinks.kafka_sink.kafka.bootstrap.servers = XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number
client.sinks.kafka_sink.flumeBatchSize = 1000
client.sinks.kafka_sink.kafka.producer.type = sync
client.sinks.kafka_sink.kafka.security.protocol = SASL_PLAINTEXT
client.sinks.kafka_sink.kafka.kerberos.domain.name = Kafka domain name. This parameter is mandatory for a security cluster, for example, hadoop.xxx.com.
client.sinks.kafka_sink.requiredAcks = 0

client.sources.static_log_source.channels = static_log_channel
client.sinks.kafka_sink.channel = static_log_channel
```

NOTE

- **client.sinks.kafka_sink.kafka.topic:** Topic to which data is written. If the topic does not exist in Kafka, it is automatically created by default.
- **client.sinks.kafka_sink.kafka.bootstrap.servers:** List of Kafka Brokers, which are separated by commas (.). By default, the port is **21007** for a security cluster and **9092** for a normal cluster.
- **client.sinks.kafka_sink.kafka.security.protocol:** The value is **SASL_PLAINTEXT** for a security cluster and **PLAINTEXT** for a normal cluster.
- **client.sinks.kafka_sink.kafka.kerberos.domain.name:**
You do not need to set this parameter for a normal cluster. For a security cluster, the value of this parameter is the value of **kerberos.domain.name** in the Kafka cluster.
Obtain the value by checking **`\${BIGDATA_HOME}/MRS_Current/1_X_Broker/etc/server.properties** on the node where the broker instance resides.
In the preceding paths, **X** indicates a random number. Change it based on site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 10 After the parameters are set and saved, the Flume client automatically loads the content configured in **properties.properties**. When new log files are generated by `spoolDir`, the files are sent to Kafka producers and can be consumed by Kafka consumers.

----End

7.2 Overview

Flume is a distributed, reliable, and highly available system for aggregating massive logs, which can efficiently collect, aggregate, and move massive log data from different data sources and store the data in a centralized data storage system. Various data senders can be customized in the system to collect data. Additionally, Flume provides simple data processes capabilities and writes data to data receivers (which is customizable).

Flume consists of the client and server, both of which are FlumeAgents. The server corresponds to the FlumeServer instance and is directly deployed in a cluster. The client can be deployed inside or outside the cluster. The client-side and service-side FlumeAgents work independently and provide the same functions.

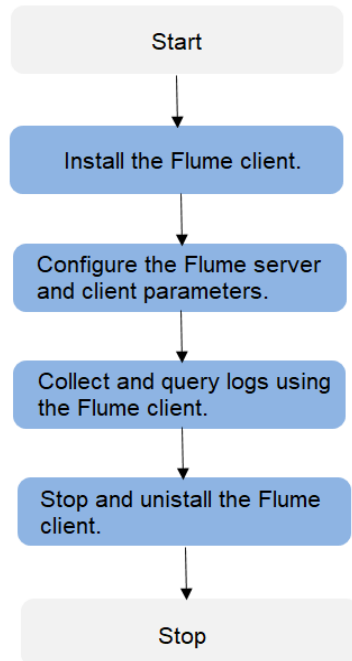
The client-side FlumeAgent needs to be independently installed. Data can be directly imported to components such as HDFS and Kafka. Additionally, the client-side and service-side FlumeAgents can also work together to provide services.

Process

The process for collecting logs using Flume is as follows:

1. Installing the flume client
2. Configuring the Flume server and client parameters
3. Collecting and querying logs using the Flume client
4. Stopping and uninstalling the Flume client

Figure 7-1 Log collection process



Flume Client

A Flume client consists of the source, channel, and sink. The source sends the data to the channel, and then the sink transmits the data from the channel to the external device. [Table 7-1](#) describes Flume modules.

Table 7-1 Module description

Name	Description
Source	<p>A source receives or generates data and sends the data to one or multiple channels. The source can work in either data-driven or polling mode.</p> <p>Typical sources include:</p> <ul style="list-style-type: none"> • Sources that are integrated with the system and receives data, such as Syslog and Netcat • Sources that automatically generate event data, such as Exec and SEQ • IPC sources that are used for communication between agents, such as Avro <p>A Source must associate with at least one channel.</p>
Channel	<p>A channel is used to buffer data between a source and a sink. After the sink transmits the data to the next channel or the destination, the cache is deleted automatically.</p> <p>The persistency of the channels varies with the channel types:</p> <ul style="list-style-type: none"> • Memory channel: non-persistency • File channel: persistency implemented based on write-ahead logging (WAL) • JDBC channel: persistency implemented based on the embedded database <p>Channels support the transaction feature to ensure simple sequential operations. A channel can work with sources and sinks of any quantity.</p>
Sink	<p>Sink is responsible for sending data to the next hop or final destination and removing the data from the channel after successfully sending the data.</p> <p>Typical sinks include:</p> <ul style="list-style-type: none"> • Sinks that send storage data to the final destination, such as HDFS and Kafka • Sinks that are consumed automatically, such as Null Sink • IPC sinks that are used for communication between agents, such as Avro <p>A sink must associate with at least one channel.</p>

A Flume client can have multiple sources, channels, and sinks. A source can send data to multiple channels, and then multiple sinks send the data out of the client.

Multiple Flume clients can be cascaded. That is, a sink can send data to the source of another client.

Supplementary Information

1. Flume provides the following reliability measures:
 - The transaction mechanism is implemented between sources and channels, and between channels and sinks.
 - The sink processor supports the failover and load balancing (load_balance) mechanisms.

The following is an example of the load balancing (load_balance) configuration:

```
server.sinkgroups=g1
server.sinkgroups.g1.sinks=k1 k2
server.sinkgroups.g1.processor.type=load_balance
server.sinkgroups.g1.processor.backoff=true
server.sinkgroups.g1.processor.selector=random
```

2. The following are precautions for the aggregation and cascading of multiple Flume clients:
 - Avro or Thrift protocol can be used for cascading.
 - When the aggregation end contains multiple nodes, evenly distribute the clients to these nodes. Do not connect all the clients to a single node.
3. The Flume client can contain multiple independent data flows. That is, multiple sources, channels, and sinks can be configured in the **properties.properties** configuration file. These components can be linked to form multiple flows.

For example, to configure two data flows in a configuration, run the following commands:

```
server.sources = source1 source2
server.sinks = sink1 sink2
server.channels = channel1 channel2

#dataflow1
server.sources.source1.channels = channel1
server.sinks.sink1.channel = channel1

#dataflow2
server.sources.source2.channels = channel2
server.sinks.sink2.channel = channel2
```

7.3 Installing the Flume Client

7.3.1 Installing the Flume Client on Clusters of Versions Earlier Than MRS 3.x

Scenario

To use Flume to collect logs, you must install the Flume client on a log host. You can create an ECS and install the Flume client on it.

This section applies to MRS 3.x or earlier clusters.

Prerequisites

- A streaming cluster with the Flume component has been created.

- The log host is in the same VPC and subnet with the MRS cluster.
- You have obtained the username and password for logging in to the log host.

Procedure

Step 1 Create an ECS that meets the requirements.

Step 2 Go to the cluster details page and choose **Components**.

NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 3 Click **Download Client**.

1. In **Client Type**, select **All client files**.
2. In **Download to**, select **Remote host**.
3. Set **Host IP Address** to the IP address of the ECS, **Host Port** to **22**, and **Save Path** to **/tmp**.
 - If the default port **22** for logging in to an ECS through SSH has been changed, set **Host Port** to a new port.
 - The value of **Save Path** contains a maximum of 256 characters.
4. Set **Login User** to **root**.

If another user is used, ensure that the user has permissions to read, write, and execute the save path.

5. 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.
6. Click **OK** to generate a client file.

If the following information is displayed, the client package is saved.

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 3** to download the client again.

Failed to connect to the server. Please check the network connection or parameter settings.

Step 4 Choose **Flume > Instance**. Query the **Business IP Address** of any Flume instance and any two MonitorServer instances.

Role	Host Name	OM IP Address	Business IP Address
Flume	node-str-coreGVaQ		
Flume	node-str-coreHEel		
MonitorServer	node-master1LHaB		
MonitorServer	node-master2mecC		

Step 5 Log in to the ECS using VNC. See section "Login Using VNC" in the *Elastic Cloud Service User Guide* (**Instances** > **Logging In to a Linux ECS** > **Login Using VNC**).

All images support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation. You are advised to change the password upon the first login.

Step 6 On the ECS, switch to user **root** and copy the installation package to the **/opt** directory.

```
sudo su - root
```

```
cp /tmp/MRS_Flume_Client.tar /opt
```

Step 7 Run the following command in the **/opt** directory to decompress the package and obtain the verification file and the configuration package of the client:

```
tar -xvf MRS_Flume_Client.tar
```

Step 8 Run the following command to verify the configuration package of the client:

```
sha256sum -c MRS_Flume_ClientConfig.tar.sha256
```

If the following information is displayed, the file package is successfully verified:

```
MRS_Flume_ClientConfig.tar: OK
```

Step 9 Run the following command to decompress **MRS_Flume_ClientConfig.tar**:

```
tar -xvf MRS_Flume_ClientConfig.tar
```

Step 10 Run the following command to install the client running environment to a new directory, for example, **/opt/Flumeenv**. A directory is automatically generated during the client installation.

```
sh /opt/MRS_Flume_ClientConfig/install.sh /opt/Flumeenv
```

If the following information is displayed, the client running environment is successfully installed:

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

Step 11 Run the following command to configure environment variables:

```
source /opt/Flumeenv/bigdata_env
```

Step 12 Run the following commands to decompress the Flume client package:

```
cd /opt/MRS_Flume_ClientConfig/Flume
```

```
tar -xvf FusionInsight-Flume-1.6.0.tar.gz
```

Step 13 Run the following command to check whether the password of the current user has expired:

```
chage -l root
```

If the value of **Password expires** is earlier than the current time, the password has expired. Run the **chage -M -1 root** command to validate the password.

Step 14 Run the following command to install the Flume client to a new directory, for example, **/opt/FlumeClient**. A directory is automatically generated during the client installation.

```
sh /opt/MRS_Flume_ClientConfig/Flume/install.sh -d /opt/FlumeClient -f  
service IP address of the MonitorServer instance -c path of the Flume  
configuration file -l /var/log/ -e service IP address of Flume -n name of the Flume  
client
```

The parameters are described as follows:

- **-d**: indicates the installation path of the Flume client.
- (Optional) **-f**: indicates the service IP addresses of the two MonitorServer instances, separated by a comma (,). If the IP addresses are not configured, the Flume client will not send alarm information to MonitorServer, and the client information will not be displayed on MRS Manager.
- (Optional) **-c**: indicates the **properties.properties** configuration file that the Flume client loads after installation. If this parameter is not specified, the **fusioninsight-flume-1.6.0/conf/properties.properties** file in the client installation directory is used by default. The configuration file of the client is empty. You can modify the configuration file as required and the Flume client will load it automatically.
- (Optional) **-l**: indicates the log directory. The default value is **/var/log/Bigdata**.
- (Optional) **-e**: indicates the service IP address of the Flume instance. It is used to receive the monitoring indicators reported by the client.
- (Optional) **-n**: indicates the name of the Flume client.
- IBM JDK does not support **-Xloggc**. You must change **-Xloggc** to **Xverbosegclog** in **flume/conf/flume-env.sh**. For 32-bit JDK, the value of **-Xmx** must not exceed 3.25 GB.
- In **flume/conf/flume-env.sh**, the default value of **-Xmx** is 4 GB. If the client memory is too small, you can change it to 512 MB or even 1 GB.

For example, run **sh install.sh -d /opt/FlumeClient**.

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

```
install flume client successfully.
```

```
----End
```

7.3.2 Installing the Flume Client on MRS 3.x or Later Clusters

Scenario

To use Flume to collect logs, you must install the Flume client on a log host. You can create an ECS and install the Flume client on it.

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

Prerequisites


- A cluster with the Flume component has been created.
- The log host is in the same VPC and subnet with the MRS cluster.
- You have obtained the username and password for logging in to the log host.
- The installation directory is automatically created if it does not exist. If it exists, the directory must be left blank. The directory path cannot contain any space.

Procedure

Step 1 Obtain the software package.

Log in to the FusionInsight Manager. Choose **Cluster** > *Name of the target cluster* > **Services** > **Flume**. On the Flume service page that is displayed, choose **More** > **Download Client** in the upper right corner and set **Select Client Type** to **Complete Client** to download the Flume service client file.

The file name of the client is **FusionInsight_Cluster_<Cluster ID>_Flume_Client.tar**. This section takes the client file **FusionInsight_Cluster_1_Flume_Client.tar** as an example.



Step 2 Upload the software package.

Upload the software package to a directory, for example, **/opt/client**, on the node where the Flume client is to be installed as user **user**.

NOTE

user is the user who installs and runs the Flume client.

Step 3 Decompress the software package.

Log in to the node where the Flume service client is to be installed as user **user**. Go to the directory where the installation package is installed, for example, **/opt/client**, and run the following command to decompress the installation package to the current directory:

```
cd /opt/client
```

```
tar -xvf FusionInsight_Cluster_1_Flume_Client.tar
```

Step 4 Verify the software package.

Run the **sha256sum -c** command to verify the decompressed file. If **OK** is returned, the verification is successful. Example:

```
sha256sum -c FusionInsight_Cluster_1_Flume_ClientConfig.tar.sha256
```

```
FusionInsight_Cluster_1_Flume_ClientConfig.tar: OK
```

Step 5 Decompress the package.

```
tar -xvf FusionInsight_Cluster_1_Flume_ClientConfig.tar
```

Step 6 Run the following command in the Flume client installation directory to install the client to a specified directory (for example, **opt/FlumeClient**): After the client is installed successfully, the installation is complete.

```
cd /opt/client/FusionInsight_Cluster_1_Flume_ClientConfig/Flume/FlumeClient
```

```
./install.sh -d /opt/FlumeClient -f MonitorServerService IP address or host name  
of the role -c User service configuration filePath for storing properties.properties -s  
CPU threshold -l /var/log/Bigdata -e FlumeServer service IP address or host name  
-n Flume
```

 NOTE

- **-d**: Flume client installation path
- (Optional) **-f**: IP addresses or host names of two MonitorServer roles. The IP addresses or host names are separated by commas (.). If this parameter is not configured, the Flume client does not send alarm information to MonitorServer and information about the client cannot be viewed on the FusionInsight Manager GUI.
- (Optional) **-c**: Service configuration file, which needs to be generated on the configuration tool page of the Flume server based on your service requirements. Upload the file to any directory on the node where the client is to be installed. If this parameter is not specified during the installation, you can upload the generated service configuration file **properties.properties** to the **/opt/FlumeClient/fusioninsight-flume-1.9.0/conf** directory after the installation.
- (Optional) **-s**: cgroup threshold. The value is an integer ranging from 1 to 100 x *N*. *N* indicates the number of CPU cores. The default threshold is **-1**, indicating that the processes added to the cgroup are not restricted by the CPU usage.
- (Optional) **-l**: Log path. The default value is **/var/log/Bigdata**. The user **user** must have the write permission on the directory. When the client is installed for the first time, a subdirectory named **flume-client** is generated. After the installation, subdirectories named **flume-client-*n*** will be generated in sequence. The letter *n* indicates a sequence number, which starts from 1 in ascending order. In the **/conf/** directory of the Flume client installation directory, modify the **ENV_VARS** file and search for the **FLUME_LOG_DIR** attribute to view the client log path.
- (Optional) **-e**: Service IP address or host name of FlumeServer, which is used to receive statistics for the monitoring indicator reported by the client.
- (Optional) **-n**: Name of the Flume client. You can choose **Cluster > Name of the desired cluster > Service > Flume > Flume Management** on FusionInsight Manager to view the client name on the corresponding node.
- If the following error message is displayed, run the **export JAVA_HOME=*JDK path*** command.
`JAVA_HOME is null in current user,please install the JDK and set the JAVA_HOME`
- IBM JDK does not support **-Xloggc**. You must change **-Xloggc** to **-Xverbosegclog** in **flume/conf/flume-env.sh**. For 32-bit JDK, the value of **-Xmx** must not exceed 3.25 GB.
- When installing a cross-platform client in a cluster, go to the **/opt/client/FusionInsight_Cluster_1_Flume_ClientConfig/Flume/FusionInsight-Flume-1.9.0.tar.gz** directory to install the Flume client.

----End

7.4 Viewing Flume Client Logs

Scenario

You can view logs to locate faults.

Prerequisites

The Flume client has been installed.

Procedure

Step 1 Go to the Flume client log directory (**/var/log/Bigdata** by default).

Step 2 Run the following command to view the log file:

ls -lR flume-client-*

A log file is shown as follows:

```
flume-client-1/flume:
total 7672
-rw-----, 1 root root    0 Sep  8 19:43 Flume-audit.log
-rw-----, 1 root root 1562037 Sep 11 06:05 FlumeClient.2017-09-11_04-05-09.[1].log.zip
-rw-----, 1 root root 6127274 Sep 11 14:47 FlumeClient.log
-rw-----, 1 root root  2935 Sep  8 22:20 flume-root-20170908202009-pid72456-gc.log.0.current
-rw-----, 1 root root  2935 Sep  8 22:27 flume-root-20170908202634-pid78789-gc.log.0.current
-rw-----, 1 root root  4382 Sep  8 22:47 flume-root-20170908203137-pid84925-gc.log.0.current
-rw-----, 1 root root  4390 Sep  8 23:46 flume-root-20170908204918-pid103920-gc.log.0.current
-rw-----, 1 root root  3196 Sep  9 10:12 flume-root-20170908215351-pid44372-gc.log.0.current
-rw-----, 1 root root  2935 Sep  9 10:13 flume-root-20170909101233-pid55119-gc.log.0.current
-rw-----, 1 root root  6441 Sep  9 11:10 flume-root-20170909101631-pid59301-gc.log.0.current
-rw-----, 1 root root    0 Sep  9 11:10 flume-root-20170909111009-pid119477-gc.log.0.current
-rw-----, 1 root root  92896 Sep 11 13:24 flume-root-20170909111126-pid120689-gc.log.0.current
-rw-----, 1 root root  5588 Sep 11 14:46 flume-root-20170911132445-pid42259-gc.log.0.current
-rw-----, 1 root root  2576 Sep 11 13:24 prestartDetail.log
-rw-----, 1 root root  3303 Sep 11 13:24 startDetail.log
-rw-----, 1 root root  1253 Sep 11 13:24 stopDetail.log

flume-client-1/monitor:
total 8
-rw-----, 1 root root  141 Sep  8 19:43 flumeMonitorChecker.log
-rw-----, 1 root root  2946 Sep 11 13:24 flumeMonitor.log
```

In the log file, **FlumeClient.log** is the run log of the Flume client.

----End

7.5 Stopping or Uninstalling the Flume Client

Scenario

You can stop and start the Flume client or uninstall the Flume client when the Flume data ingestion channel is not required.

Procedure

- Stop the Flume client of the Flume role.
Assume that the Flume client installation path is **/opt/FlumeClient**. Run the following command to stop the Flume client:

```
cd /opt/FlumeClient/fusioninsight-flume-Flume component version number/bin  
./flume-manage.sh stop
```

If the following information is displayed after the command execution, the Flume client is successfully stopped.

```
Stop Flume PID=120689 successful..
```

NOTE

The Flume client will be automatically restarted after being stopped. If you do not need automatic restart, run the following command:

```
./flume-manage.sh stop force
```

If you want to restart the Flume client, run the following command:

```
./flume-manage.sh start force
```


- Uninstall the Flume client of the Flume role.
Assume that the Flume client installation path is `/opt/FlumeClient`. Run the following command to uninstall the Flume client:

```
cd /opt/FlumeClient/fusioninsight-flume-Flume component version number/inst  
./uninstall.sh
```

7.6 Using the Encryption Tool of the Flume Client

Scenario

You can use the encryption tool provided by the Flume client to encrypt some parameter values in the configuration file.

Prerequisites

The Flume client has been installed.

Procedure

Step 1 Log in to the Flume client node and go to the client installation directory, for example, `/opt/FlumeClient`.

Step 2 Run the following command to switch the directory:

```
cd fusioninsight-flume-Flume component version number/bin
```

Step 3 Run the following command to encrypt information:

```
./genPwFile.sh
```

Input the information that you want to encrypt twice.

Step 4 Run the following command to query the encrypted information:

```
cat password.property
```

NOTE

If the encryption parameter is used for the Flume server, you need to perform encryption on the corresponding Flume server node. You need to run the encryption script as user `omm` for encryption.

- For versions earlier than MRS 3.x, the encryption path is `/opt/Bigdata/MRS_XXX/install/FusionInsight-Flume-Flume component version number/flume/bin/genPwFile.sh`.
- For MRS 3.x or later, the encryption path is `/opt/Bigdata/FusionInsight_Porter_XXX/install/FusionInsight-Flume-Flume component version number/flume/bin/genPwFile.sh`. *XXX* indicates the product version number.

----End

7.7 Flume Service Configuration Guide

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

This configuration guide describes how to configure common Flume services.

 **NOTE**

- Parameters in bold in the following tables are mandatory.
- The value of **BatchSize** of the Sink must be less than that of **transactionCapacity** of the Channel.
- Only some parameters of Source, Channel, and Sink are displayed on the Flume configuration tool page. For details, see the following configurations.
- The Customer Source, Customer Channel, and Customer Sink displayed on the Flume configuration tool page need to be configured based on self-developed code. The following common configurations are not displayed.

Common Source Configurations

- **Avro Source**

An Avro source listens to the Avro port, receives data from the external Avro client, and places data into configured channels. Common configurations are as follows:

Table 7-2 Common configurations of an Avro source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured.
type	avro	Specifies the type of the avro source, which must be avro .
bind	-	Specifies the listening host name/IP address.
port	-	Specifies the bound listening port. Ensure that this port is not occupied.
threads	-	Specifies the maximum number of source threads.
compression-type	none	Specifies the message compression format, which can be set to none or deflate . none indicates that data is not compressed, while deflate indicates that data is compressed.

Parameter	Default Value	Description
compression-level	6	Specifies the data compression level, which ranges from 1 to 9 . The larger the value is, the higher the compression rate is.
ssl	false	Specifies whether to use SSL encryption. If this parameter is set to true , the values of keystore and keystore-password must be specified.
truststore-type	JKS	Specifies the Java trust store type, which can be set to JKS or PKCS12 . NOTE Different passwords are used to protect the key store and private key of JKS , while the same password is used to protect the key store and private key of PKCS12 .
truststore	-	Specifies the Java trust store file.
truststore-password	-	Specifies the Java trust store password.
keystore-type	JKS	Specifies the keystore type set after SSL is enabled, which can be set to JKS or PKCS12 . NOTE Different passwords are used to protect the key store and private key of JKS , while the same password is used to protect the key store and private key of PKCS12 .
keystore	-	Specifies the keystore file path set after SSL is enabled. This parameter is mandatory if SSL is enabled.

Parameter	Default Value	Description
keystore-password	-	Specifies the keystore password set after SSL is enabled. This parameter is mandatory if SSL is enabled.
trust-all-certs	false	Specifies whether to disable the check for the SSL server certificate. If this parameter is set to true , the SSL server certificate of the remote source is not checked. You are not advised to perform this operation during the production.
exclude-protocols	SSLv3	Specifies the excluded protocols. The entered protocols must be separated by spaces. The default value is SSLv3 .
ipFilter	false	Specifies whether to enable the IP address filtering.
ipFilter.rules	-	Specifies the rules of <i>N</i> network ipFilters . Host names or IP addresses must be separated by commas (.). If this parameter is set to true , there are two configuration rules: allow and forbidden. The configuration format is as follows: ipFilterRules=allow:ip:127.*, allow:name:localhost, deny:ip:*

- **SpoolDir Source**

SpoolDir Source monitors and transmits new files that have been added to directories in real-time mode. Common configurations are as follows:

Table 7-3 Common configurations of a Spooling Directory source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured.
type	spooldir	Specifies the type of the spooling source, which must be set to spooldir .
spoolDir	-	Specifies the monitoring directory of the Spooldir source. A Flume running user must have the read, write, and execution permissions on the directory.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the source is restarted. Unit: second
fileSuffix	.COMPLETED	Specifies the suffix added after file transmission is complete.
deletePolicy	never	Specifies the source file deletion policy after file transmission is complete. The value can be either never or immediate . never indicates that the source file is not deleted after file transmission is complete, while immediate indicates that the source file is immediately deleted after file transmission is complete.
ignorePattern	^\$	Specifies the regular expression of a file to be ignored. The default value is ^\$, indicating that spaces are ignored.
includePattern	^.*\$	Specifies the regular expression that contains a file. This parameter can be used together with ignorePattern . If a file meets both ignorePattern and includePattern , the file is ignored. In addition, when a file starts with a period (.), the file will not be filtered.
trackerDir	.flumespool	Specifies the metadata storage path during data transmission.
batchSize	1000	Specifies the number of events written to the channel in batches.
decodeErrorPolicy	FAIL	Specifies the code error policy. NOTE If a code error occurs in the file, set decodeErrorPolicy to REPLACE or IGNORE . Flume will skip the code error and continue to collect subsequent logs.

Parameter	Default Value	Description
deserializer	LINE	Specifies the file parser. The value can be either LINE or BufferedLine . <ul style="list-style-type: none"> When the value is set to LINE, characters read from the file are transcoded one by one. When the value is set to BufferedLine, one line or multiple lines of characters read from the file are transcoded in batches, which delivers better performance.
deserializer.max LineLength	2048	Specifies the maximum length for resolution by line.
deserializer.max BatchLine	1	Specifies the maximum number of lines for resolution by line. If multiple lines are set, maxLineLength must be set to a corresponding multiplier. NOTE When configuring the Interceptor, take the multi-line combination into consideration to avoid data loss. If the Interceptor cannot process combined lines, set this parameter to 1 .
selector.type	replicating	Specifies the selector type. The value can be either replicating or multiplexing . replicating indicates that data is replicated and then transferred to each channel so that each channel receives the same data, while multiplexing indicates that a channel is selected based on the value of the header in the event and each channel has different data.
interceptors	-	Specifies the interceptor. Multiple interceptors are separated by spaces.
inputCharset	UTF-8	Specifies the encoding format of a read file. The encoding format must be the same as that of the data source file that has been read. Otherwise, an error may occur during character parsing.
fileHeader	false	Specifies whether to add the file name (including the file path) to the event header.

Parameter	Default Value	Description
fileHeaderKey	-	Specifies that the data storage structure in header is set in the <key,value> mode. Parameters fileHeaderKey and fileHeader must be used together. Following is an example if fileHeader is set to true: Define fileHeaderKey as file . When the / root/a.txt file is read, fileHeaderKey exists in the header in the file=/root/a.txt format.
basenameHeader	false	Specifies whether to add the file name (excluding the file path) to the event header.
basenameHeaderKey	-	Specifies that the data storage structure in header is set in the <key,value> mode. Parameters basenameHeaderKey and basenameHeader must be used together. Following is an example if basenameHeader is set to true: Define basenameHeaderKey as file . When the a.txt file is read, fileHeaderKey exists in the header in the file=a.txt format.
pollDelay	500	Specifies the delay for polling new files in the monitoring directory. Unit: milliseconds
recursiveDirectorySearch	false	Specifies whether to monitor new files in the subdirectory of the configured directory.
consumeOrder	oldest	Specifies the consumption order of files in a directory. If this parameter is set to oldest or youngest , the sequence of files to be read is determined by the last modification time of files in the monitored directory. If there are a large number of files in the directory, it takes a long time to search for oldest or youngest files. If this parameter is set to random , an earlier created file may not be read for a long time. If this parameter is set to oldest or youngest , it takes a long time to find the latest and the earliest file. The options are as follows: random , youngest , and oldest .

Parameter	Default Value	Description
maxBackoff	4000	Specifies the maximum time to wait between consecutive attempts to write to a channel if the channel is full. If the time exceeds the threshold, an exception is thrown. The corresponding source starts to write at a smaller time value. Each time the source attempts, the digital exponent increases until the current specified value is reached. If data cannot be written, the data write fails. Unit: second
emptyFileEvent	true	Specifies whether to collect empty file information and send it to the sink end. The default value is true , indicating that empty file information is sent to the sink end. This parameter is valid only for HDFS Sink. Taking HDFS Sink as an example, if this parameter is set to true and an empty file exists in the spoolDir directory, an empty file with the same name will be created in the hdfs.path directory of HDFS.

 **NOTE**

SpoolDir Source ignores the last line feed character of each event when data is reading by row. Therefore, Flume does not calculate the data volume counters used by the last line feed character.

- **Kafka Source**

A Kafka source consumes data from Kafka topics. Multiple sources can consume data of the same topic, and the sources consume different partitions of the topic. Common configurations are as follows:

Table 7-4 Common configurations of a Kafka source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured.
type	org.apache.flume.source.kafka.KafkaSource	Specifies the type of the Kafka source, which must be set to org.apache.flume.source.kafka.KafkaSource .

Parameter	Default Value	Description
kafka.bootstrap.servers	-	Specifies the bootstrap address port list of Kafka. If Kafka has been installed in the cluster and the configuration has been synchronized to the server, you do not need to set this parameter on the server. The default value is the list of all brokers in the Kafka cluster. This parameter must be configured on the client. Use commas (,) to separate multiple values of <i>IP address:Port number</i> . The rules for matching ports and security protocols must be as follows: port 21007 matches the security mode (SASL_PLAINTEXT), and port 9092 matches the common mode (PLAINTEXT).
kafka.topics	-	Specifies the list of subscribed Kafka topics, which are separated by commas (,).
kafka.topics.regex	-	Specifies the subscribed topics that comply with regular expressions. kafka.topics.regex has a higher priority than kafka.topics and will overwrite kafka.topics .
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the source is restarted. Unit: second
nodatotime	0 (Disabled)	Specifies the alarm threshold. An alarm is triggered when the duration that Kafka does not release data to subscribers exceeds the threshold. Unit: second This parameter can be configured in the properties.properties file.
batchSize	1000	Specifies the number of events written to the channel in batches.
batchDurationMillis	1000	Specifies the maximum duration of topic data consumption at a time, expressed in milliseconds.
keepTopicInHeader	false	Specifies whether to save topics in the event header. If the parameter value is true , topics configured in Kafka Sink become invalid.

Parameter	Default Value	Description
setTopicHeader	true	If this parameter is set to true , the topic name defined in topicHeader is stored in the header.
topicHeader	topic	When setTopicHeader is set to true , this parameter specifies the name of the topic received by the storage device. If the property is used with that of Kafka Sink topicHeader , be careful not to send messages to the same topic cyclically.
useFlumeEventFormat	false	By default, an event is transferred from a Kafka topic to the body of the event in the form of bytes. If this parameter is set to true , the Avro binary format of Flume is used to read events. When used together with the parseAsFlumeEvent parameter with the same name in KafkaSink or KafkaChannel, any set header generated from the data source is retained.
keepPartitionHeader	false	Specifies whether to save partition IDs in the event header. If the parameter value is true , Kafka Sink writes data to the corresponding partition.
kafka.consumer.group.id	flume	Specifies the Kafka consumer group ID. Sources or proxies having the same ID are in the same consumer group.
kafka.security.protocol	SASL_PLAINTEXT	Specifies the Kafka security protocol. The parameter value must be set to PLAINTEXT in a common cluster. The rules for matching ports and security protocols must be as follows: port 21007 matches the security mode (SASL_PLAINTEXT), and port 9092 matches the common mode (PLAINTEXT).

Parameter	Default Value	Description
Other Kafka Consumer Properties	-	Specifies other Kafka configurations. This parameter can be set to any consumption configuration supported by Kafka, and the .kafka prefix must be added to the configuration.

- **Taildir Source**

A Taildir source monitors file changes in a directory and automatically reads the file content. In addition, it can transmit data in real time. Common configurations are as follows:

Table 7-5 Common configurations of a Taildir source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured.
type	TAILDIR	Specifies the type of the taildir source, which must be set to TAILDIR.
filegroups	-	Specifies the group name of a collection file directory. Group names are separated by spaces.
filegroups.<filegroupName>.parentDir	-	Specifies the parent directory. The value must be an absolute path.
filegroups.<filegroupName>.filePattern	-	Specifies the relative file path of the file group's parent directory. Directories can be included and regular expressions are supported. It must be used together with parentDir .
positionFile	-	Specifies the metadata storage path during data transmission.
headers.<filegroupName>.<headerKey>	-	Specifies the key-value of an event when data of a group is being collected.
byteOffsetHeader	false	Specifies whether each event header contains the event location information in the source file. If the parameter value is true, the location information is saved in the byteoffset variable.

Parameter	Default Value	Description
maxBatchCount	Long.MAX_VALUE	Specifies the maximum number of batches that can be consecutively read from a file. If the monitored directory reads multiple files consecutively and one of the files is written at a rapid rate, other files may fail to be processed. This is because the file that is written at a high speed will be in an infinite read loop. In this case, set this parameter to a smaller value.
skipToEnd	false	Specifies whether Flume can locate the latest location of a file and read the latest data after restart. If the parameter value is true, Flume locates and reads the latest file data after restart.
idleTimeout	120000	Specifies the idle duration during file reading, expressed in milliseconds. If file content is not changed in the preset time duration, close the file. If data is written to this file after the file is closed, open the file and read data.
writePosInterval	3000	Specifies the interval for writing metadata to a file, expressed in milliseconds.
batchSize	1000	Specifies the number of events written to the channel in batches.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the source is restarted. Unit: second
fileHeader	false	Specifies whether to add the file name (including the file path) to the event header.
fileHeaderKey	file	Specifies that the data storage structure in header is set in the <key,value> mode. Parameters fileHeaderKey and fileHeader must be used together. Following is an example if fileHeader is set to true: Define fileHeaderKey as file . When the /root/a.txt file is read, fileHeaderKey exists in the header in the file=/root/a.txt format.

- **Http Source**

An HTTP source receives data from an external HTTP client and sends the data to the configured channels. Common configurations are as follows:

Table 7-6 Common configurations of an HTTP source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured.
type	http	Specifies the type of the http source, which must be set to http.
bind	-	Specifies the listening host name/IP address.
port	-	Specifies the bound listening port. Ensure that this port is not occupied.
handler	org.apache.flume.source.http.JSONHandler	Specifies the message parsing method of an HTTP request. Two formats are supported: JSON (org.apache.flume.source.http.JSONHandler) and BLOB (org.apache.flume.sink.solr.morphline.BlobHandler).
handler.*	-	Specifies handler parameters.
exclude-protocols	SSLv3	Specifies the excluded protocols. The entered protocols must be separated by spaces. The default value is SSLv3 .
include-cipher-suites	-	Specifies the included protocols. The entered protocols must be separated by spaces. If this parameter is left empty, all protocols are supported by default.
enableSSL	false	Specifies whether SSL is enabled in HTTP. If this parameter is set to true , the values of keystore and keystore-password must be specified.
keystore-type	JKS	Specifies the keystore type, which can be JKS or PKCS12 .

Parameter	Default Value	Description
keystore	-	Specifies the keystore path set after SSL is enabled in HTTP.
keystorePassword	-	Specifies the keystore password set after SSL is enabled in HTTP.

- **Thrift Source**

Thrift Source monitors the thrift port, receives data from the external Thrift clients, and puts the data into the configured channel. Common configurations are as follows:

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured.
type	thrift	Specifies the type of the thrift source, which must be set to thrift .
bind	-	Specifies the listening host name/IP address.
port	-	Specifies the bound listening port. Ensure that this port is not occupied.
threads	-	Specifies the maximum number of worker threads that can be run.
kerberos	false	Specifies whether Kerberos authentication is enabled.
agent-keytab	-	Specifies the address of the keytab file used by the server. The machine-machine account must be used. You are advised to use flume/conf/flume_server.keytab in the Flume service installation directory.
agent-principal	-	Specifies the principal of the security user used by the server. The principal must be a machine-machine account. You are advised to use the default user of Flume: <code>flume_server/hadoop.<system domain name>@<system domain name></code> NOTE <code>flume_server/hadoop.<system domain name></code> is the username. All letters in the system domain name contained in the username are lowercase letters. For example, Local Domain is set to 9427068F-6EFA-4833-B43E-60CB641E5B6C.COM , and the username is flume_server/hadoop.9427068f-6efa-4833-b43e-60cb641e5b6c.com .

Parameter	Default Value	Description
compression-type	none	Specifies the message compression format, which can be set to none or deflate . none indicates that data is not compressed, while deflate indicates that data is compressed.
ssl	false	Specifies whether to use SSL encryption. If this parameter is set to true , the values of keystore and keystore-password must be specified.
keystore-type	JKS	Specifies the keystore type set after SSL is enabled.
keystore	-	Specifies the keystore file path set after SSL is enabled. This parameter is mandatory if SSL is enabled.
keystore-password	-	Specifies the keystore password set after SSL is enabled. This parameter is mandatory if SSL is enabled.

Common Channel Configurations

- **Memory Channel**

A memory channel uses memory as the cache. Events are stored in memory queues. Common configurations are as follows:

Table 7-7 Common configurations of a memory channel

Parameter	Default Value	Description
type	-	Specifies the type of the memory channel, which must be set to memory .
capacity	10000	Specifies the maximum number of events cached in a channel.

Parameter	Default Value	Description
transactionCapacity	1000	Specifies the maximum number of events accessed each time. NOTE <ul style="list-style-type: none"> The parameter value must be greater than the batchSize of the source and sink. The value of transactionCapacity must be less than or equal to that of capacity.
channelFullCount	10	Specifies the channel full count. When the count reaches the threshold, an alarm is reported.
keep-alive	3	Specifies the waiting time of the Put and Take threads when the transaction or channel cache is full. Unit: second
byteCapacity	80% of the maximum JVM memory	Specifies the total bytes of all event bodies in a channel. The default value is the 80% of the maximum JVM memory (indicated by -Xmx). Unit: bytes
byteCapacityBufferPercentage	20	Specifies the percentage of bytes in a channel (%).

- File Channel**

A file channel uses local disks as the cache. Events are stored in the folder specified by **dataDirs**. Common configurations are as follows:

Table 7-8 Common configurations of a file channel

Parameter	Default Value	Description
type	-	Specifies the type of the file channel, which must be set to file .

Parameter	Default Value	Description
checkpointDir	\${BIGDATA_DATA_HOME}/ hadoop/data1~N/flume/ checkpoint NOTE This path is changed with the custom data path.	Specifies the checkpoint storage directory.
dataDirs	\${BIGDATA_DATA_HOME}/ hadoop/data1~N/flume/data NOTE This path is changed with the custom data path.	Specifies the data cache directory. Multiple directories can be configured to improve performance. The directories are separated by commas (.).
maxFileSize	2146435071	Specifies the maximum size of a single cache file, expressed in bytes.
minimumRequiredSpace	524288000	Specifies the minimum idle space in the cache, expressed in bytes.
capacity	1000000	Specifies the maximum number of events cached in a channel.
transactionCapacity	10000	Specifies the maximum number of events accessed each time. NOTE <ul style="list-style-type: none"> The parameter value must be greater than the batchSize of the source and sink. The value of transactionCapacity must be less than or equal to that of capacity.
channelFullCount	10	Specifies the channel full count. When the count reaches the threshold, an alarm is reported.

Parameter	Default Value	Description
useDualCheckpoints	false	Specifies the backup checkpoint. If this parameter is set to true , the backupCheckpointDir parameter value must be set.
backupCheckpointDir	-	Specifies the path of the backup checkpoint.
checkpointInterval	30000	Specifies the check interval, expressed in seconds.
keep-alive	3	Specifies the waiting time of the Put and Take threads when the transaction or channel cache is full. Unit: second
use-log-replay-v1	false	Specifies whether to enable the old reply logic.
use-fast-replay	false	Specifies whether to enable the queue reply.
checkpointOnClose	true	Specifies that whether a checkpoint is created when a channel is disabled.

- **Memory File Channel**

A memory file channel uses both memory and local disks as its cache and supports message persistence. It provides similar performance as a memory channel and better performance than a file channel. This channel is currently experimental and not recommended for use in production. The following table describes common configuration items: Common configurations are as follows:

Table 7-9 Common configurations of a memory file channel

Parameter	Default Value	Description
type	org.apache.flume.channel.MemoryFileChannel	Specifies the type of the memory file channel, which must be set to org.apache.flume.channel.MemoryFileChannel .

Parameter	Default Value	Description
capacity	50000	Specifies the maximum number of events cached in a channel.
transactionCapacity	5000	Specifies the maximum number of events processed by a transaction. NOTE <ul style="list-style-type: none"> The parameter value must be greater than the batchSize of the source and sink. The value of transactionCapacity must be less than or equal to that of capacity.
subqueueByteCapacity	20971520	Specifies the maximum size of events that can be stored in a subqueue, expressed in bytes. A memory file channel uses both queues and subqueues to cache data. Events are stored in a subqueue, and subqueues are stored in a queue. subqueueCapacity and subqueueInterval determine the size of events that can be stored in a subqueue. subqueueCapacity specifies the capacity of a subqueue, and subqueueInterval specifies the duration that a subqueue can store events. Events in a subqueue are sent to the destination only after the subqueue reaches the upper limit of subqueueCapacity or subqueueInterval . NOTE The value of subqueueByteCapacity must be greater than the number of events specified by batchSize .
subqueueInterval	2000	Specifies the maximum duration that a subqueue can store events, expressed in milliseconds.
keep-alive	3	Specifies the waiting time of the Put and Take threads when the transaction or channel cache is full. Unit: second
dataDir	-	Specifies the cache directory for local files.
byteCapacity	80% of the maximum JVM memory	Specifies the channel cache capacity. Unit: bytes

Parameter	Default Value	Description
compression-type	None	Specifies the message compression format, which can be set to none or deflate . none indicates that data is not compressed, while deflate indicates that data is compressed.
channelfullcount	10	Specifies the channel full count. When the count reaches the threshold, an alarm is reported.

The following is a configuration example of a memory file channel:

```
server.channels.c1.type = org.apache.flume.channel.MemoryFileChannel
server.channels.c1.dataDir = /opt/flume/mfdata
server.channels.c1.subqueueByteCapacity = 20971520
server.channels.c1.subqueueInterval=2000
server.channels.c1.capacity = 500000
server.channels.c1.transactionCapacity = 40000
```

- **Kafka Channel**

A Kafka channel uses a Kafka cluster as the cache. Kafka provides high availability and multiple copies to prevent data from being immediately consumed by sinks when Flume or Kafka Broker crashes.

Table 7-10 Common configurations of a Kafka channel

Parameter	Default Value	Description
type	-	Specifies the type of the Kafka channel, which must be set to org.apache.flume.channel.kafka.KafkaChannel .

Parameter	Default Value	Description
kafka.bootstrap.servers	-	<p>Specifies the bootstrap address port list of Kafka.</p> <p>If Kafka has been installed in the cluster and the configuration has been synchronized to the server, you do not need to set this parameter on the server. The default value is the list of all brokers in the Kafka cluster. This parameter must be configured on the client. Use commas (,) to separate multiple values of <i>IP address:Port number</i>. The rules for matching ports and security protocols must be as follows: port 21007 matches the security mode (SASL_PLAINTEXT), and port 9092 matches the common mode (PLAINTEXT).</p>
kafka.topic	flume-channel	Specifies the Kafka topic used by the channel to cache data.
kafka.consumer.group.id	flume	Specifies the data group ID obtained from Kafka. This parameter cannot be left blank.
parseAsFlumeEvent	true	Specifies whether data is parsed into Flume events.
migrateZookeeperOffsets	true	Specifies whether to search for offsets in ZooKeeper and submit them to Kafka when there is no offset in Kafka.

Parameter	Default Value	Description
kafka.consumer.auto.offset.reset	latest	Specifies where to consume if there is no offset record, which can be set to earliest , latest , or none . earliest indicates that the offset is reset to the initial point, latest indicates that the offset is set to the latest position, and none indicates that an exception is thrown if there is no offset.
kafka.producer.security.protocol	SASL_PLAINTEXT	Specifies the Kafka producer security protocol. The rules for matching ports and security protocols must be as follows: port 21007 matches the security mode (SASL_PLAINTEXT), and port 9092 matches the common mode (PLAINTEXT). NOTE If the parameter is not displayed, click + in the lower left corner of the dialog box to display all parameters.
kafka.consumer.security.protocol	SASL_PLAINTEXT	Specifies the Kafka consumer security protocol. The rules for matching ports and security protocols must be as follows: port 21007 matches the security mode (SASL_PLAINTEXT), and port 9092 matches the common mode (PLAINTEXT).
pollTimeout	500	Specifies the maximum timeout interval for the consumer to invoke the poll function. Unit: milliseconds

Parameter	Default Value	Description
ignoreLongMessage	false	Specifies whether to discard oversized messages.
messageMaxLength	1000012	Specifies the maximum length of a message written by Flume to Kafka.

Common Sink Configurations

- **HDFS Sink**

An HDFS sink writes data into HDFS. Common configurations are as follows:

Table 7-11 Common configurations of an HDFS sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink.
type	hdfs	Specifies the type of the hdfs sink, which must be set to hdfs .
hdfs.path	-	Specifies the data storage path in HDFS. The value must start with hdfs://hacluster/ .
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the sink is restarted. Unit: second
hdfs.inUseSuffix	.tmp	Specifies the suffix of the HDFS file to which data is being written.
hdfs.rollInterval	30	Specifies the interval for file rolling, expressed in seconds.
hdfs.rollSize	1024	Specifies the size for file rolling, expressed in bytes.
hdfs.rollCount	10	Specifies the number of events for file rolling. NOTE Parameters rollInterval , rollSize , and rollCount can be configured at the same time. The parameter meeting the requirements takes precedence for compression.

Parameter	Default Value	Description
hdfs.idleTimeout	0	Specifies the timeout interval for closing idle files automatically, expressed in seconds.
hdfs.batchSize	1000	Specifies the number of events written into HDFS in batches.
hdfs.kerberosPrincipal	-	Specifies the Kerberos principal of HDFS authentication. This parameter is mandatory in a secure mode, but not required in a common mode.
hdfs.kerberosKeytab	-	Specifies the Kerberos keytab of HDFS authentication. This parameter is not required in a common mode, but in a secure mode, the Flume running user must have the permission to access keyTab path in the jaas.cof file.
hdfs.fileCloseByEndEvent	true	Specifies whether to close the HDFS file when the last event of the source file is received.
hdfs.batchCallTimeout	-	<p>Specifies the timeout control duration when events are written into HDFS in batches. Unit: milliseconds</p> <p>If this parameter is not specified, the timeout duration is controlled when each event is written into HDFS. When the value of hdfs.batchSize is greater than 0, configure this parameter to improve the performance of writing data into HDFS.</p> <p>NOTE The value of hdfs.batchCallTimeout depends on hdfs.batchSize. A greater hdfs.batchSize requires a larger hdfs.batchCallTimeout. If the value of hdfs.batchCallTimeout is too small, writing events to HDFS may fail.</p>
serializer.appendNewline	true	Specifies whether to add a line feed character (\n) after an event is written to HDFS. If a line feed character is added, the data volume counters used by the line feed character will not be calculated by HDFS sinks.

Parameter	Default Value	Description
hdfs.filePrefix	over_ % {base name}	Specifies the file name prefix after data is written to HDFS.
hdfs.fileSuffix	-	Specifies the file name suffix after data is written to HDFS.
hdfs.inUsePrefix	-	Specifies the prefix of the HDFS file to which data is being written.
hdfs.fileType	DataStream	Specifies the HDFS file format, which can be set to SequenceFile , DataStream , or CompressedStream . NOTE If the parameter is set to SequenceFile or DataStream , output files are not compressed, and the codeC parameter cannot be configured. However, if the parameter is set to CompressedStream , the output files are compressed, and the codeC parameter must be configured together.
hdfs.codeC	-	Specifies the file compression format, which can be set to gzip , bzip2 , lzo , lzop , or snappy .
hdfs.maxOpenFiles	5000	Specifies the maximum number of HDFS files that can be opened. If the number of opened files reaches this value, the earliest opened files are closed.
hdfs.writeFormat	Writable	Specifies the file write format, which can be set to Writable or Text .
hdfs.callTimeout	10000	Specifies the timeout control duration each time events are written into HDFS, expressed in milliseconds.
hdfs.threadsPoolSize	-	Specifies the number of threads used by each HDFS sink for HDFS I/O operations.
hdfs.rollTimerPoolSize	-	Specifies the number of threads used by each HDFS sink to schedule the scheduled file rolling.
hdfs.round	false	Specifies whether to round off the timestamp value. If this parameter is set to true, all time-based escape sequences (except %t) are affected.

Parameter	Default Value	Description
hdfs.roundUnit	second	Specifies the unit of the timestamp value that has been rounded off, which can be set to second , minute , or hour .
hdfs.useLocalTimeStamp	true	Specifies whether to enable the local timestamp. The recommended parameter value is true .
hdfs.closeTries	0	Specifies the maximum attempts for the hdfs sink to stop renaming a file. If the parameter is set to the default value 0 , the sink does not stop renaming the file until the file is successfully renamed.
hdfs.retryInterval	180	Specifies the interval of request for closing the HDFS file, expressed in seconds. NOTE For each closing request, there are multiple RPCs working on the NameNode back and forth, which may make the NameNode overloaded if the parameter value is too small. Also, when the parameter is set to 0 , the Sink will not attempt to close the file, but opens the file or uses .tmp as the file name extension, if the first closing attempt fails.
hdfs.failcount	10	Specifies the number of times that data fails to be written to HDFS. If the number of times that the sink fails to write data to HDFS exceeds the parameter value, an alarm indicating abnormal data transmission is reported.

- **Avro Sink**

An Avro sink converts events into Avro events and sends them to the monitoring ports of the hosts. Common configurations are as follows:

Table 7-12 Common configurations of an Avro sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink.
type	-	Specifies the type of the avro sink, which must be set to avro .

Parameter	Default Value	Description
hostname	-	Specifies the bound host name or IP address.
port	-	Specifies the bound listening port. Ensure that this port is not occupied.
batch-size	1000	Specifies the number of events sent in a batch.
client.type	DEFAULT	<p>Specifies the client instance type. Set this parameter based on the communication protocol used by the configured model. The options are as follows:</p> <ul style="list-style-type: none"> • DEFAULT: The client instance of the AvroRPC type is returned. • OTHER: NULL is returned. • THRIFT: The client instance of the Thrift RPC type is returned. • DEFAULT_LOADBALANCING: The client instance of the LoadBalancing RPC type is returned. • DEFAULT_FAILOVER: The client instance of the Failover RPC type is returned.
ssl	false	Specifies whether to use SSL encryption. If this parameter is set to true , the values of keystore and keystore-password must be specified.

Parameter	Default Value	Description
truststore-type	JKS	Specifies the Java trust store type, which can be set to JKS or PKCS12 . NOTE Different passwords are used to protect the key store and private key of JKS , while the same password is used to protect the key store and private key of PKCS12 .
truststore	-	Specifies the Java trust store file.
truststore-password	-	Specifies the Java trust store password.
keystore-type	JKS	Specifies the keystore type set after SSL is enabled.
keystore	-	Specifies the keystore file path set after SSL is enabled. This parameter is mandatory if SSL is enabled.
keystore-password	-	Specifies the keystore password after SSL is enabled. This parameter is mandatory if SSL is enabled.
connect-timeout	20000	Specifies the timeout for the first connection, expressed in milliseconds.
request-timeout	20000	Specifies the maximum timeout for a request after the first request, expressed in milliseconds.

Parameter	Default Value	Description
reset-connection-interval	0	Specifies the interval between a connection failure and a second connection, expressed in seconds. If the parameter is set to 0 , the system continuously attempts to perform a connection.
compression-type	none	Specifies the compression type of the batch data, which can be set to none or deflate . none indicates that data is not compressed, while deflate indicates that data is compressed. This parameter value must be the same as that of the AvroSource compression-type.
compression-level	6	Specifies the compression level of batch data, which can be set to 1 to 9 . A larger value indicates a higher compression rate.
exclude-protocols	SSLv3	Specifies the excluded protocols. The entered protocols must be separated by spaces. The default value is SSLv3 .

- **HBase Sink**

An HBase sink writes data into HBase. Common configurations are as follows:

Table 7-13 Common configurations of an HBase sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink.

Parameter	Default Value	Description
type	-	Specifies the type of the HBase sink, which must be set to hbase .
table	-	Specifies the HBase table name.
columnFamily	-	Specifies the HBase column family.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the sink is restarted. Unit: second
batchSize	1000	Specifies the number of events written into HBase in batches.
kerberosPrincipal	-	Specifies the Kerberos principal of HBase authentication. This parameter is mandatory in a secure mode, but not required in a common mode.
kerberosKeytab	-	Specifies the Kerberos keytab of HBase authentication. This parameter is not required in a common mode, but in a secure mode, the Flume running user must have the permission to access keyTab path in the jaas.cof file.
coalesceIncrements	true	Specifies whether to perform multiple operations on the same hbase cell in a same processing batch. Setting this parameter to true improves performance.

- **Kafka Sink**

A Kafka sink writes data into Kafka. Common configurations are as follows:

Table 7-14 Common configurations of a Kafka sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink.
type	-	Specifies the type of the kafka sink, which must be set to org.apache.flume.sink.kafka.KafkaSink .

Parameter	Default Value	Description
kafka.bootstrap.servers	-	Specifies the bootstrap address port list of Kafka. If Kafka has been installed in the cluster and the configuration has been synchronized to the server, you do not need to set this parameter on the server. The default value is the list of all brokers in the Kafka cluster. The client must be configured with this parameter. If there are multiple values, use commas (,) to separate the values. The rules for matching ports and security protocols must be as follows: port 21007 matches the security mode (SASL_PLAINTEXT), and port 9092 matches the common mode (PLAINTEXT).
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the sink is restarted. Unit: second
kafka.producer.acks	1	Successful write is determined by the number of received acknowledgement messages about replicas. The value 0 indicates that no confirm message needs to be received, the value 1 indicates that the system is only waiting for only the acknowledgement information from a leader, and the value -1 indicates that the system is waiting for the acknowledgement messages of all replicas. If this parameter is set to -1 , data loss can be avoided in some leader failure scenarios.
kafka.topic	-	Specifies the topic to which data is written. This parameter is mandatory.
flumeBatchSize	1000	Specifies the number of events written into Kafka in batches.
kafka.security.protocol	SASL_PLAINTEXT	Specifies the Kafka security protocol. The parameter value must be set to PLAINTEXT in a common cluster. The rules for matching ports and security protocols must be as follows: port 21007 matches the security mode (SASL_PLAINTEXT), and port 9092 matches the common mode (PLAINTEXT).

Parameter	Default Value	Description
ignoreLongMessage	false	Specifies whether to discard oversized messages.
messageMaxLength	1000012	Specifies the maximum length of a message written by Flume to Kafka.
defaultPartitionId	-	Specifies the Kafka partition ID to which the events of a channel is transferred. The partitionIdHeader value overwrites this parameter value. By default, if this parameter is left blank, events will be distributed by the Kafka Producer's partitioner (by a specified key or a partitioner customized by kafka.partitionner.class).
partitionIdHeader	-	When you set this parameter, the sink will take the value of the field named using the value of this property from the event header and send the message to the specified partition of the topic. If the value does not have a valid partition, EventDeliveryException is thrown. If the header value already exists, this setting overwrites the defaultPartitionId parameter.
Other Kafka Producer Properties	-	Specifies other Kafka configurations. This parameter can be set to any production configuration supported by Kafka, and the .kafka prefix must be added to the configuration.

- **Thrift Sink**

A Thrift sink converts events to Thrift events and sends them to the monitoring port of the configured host. Common configurations are as follows:

Table 7-15 Common configurations of a Thrift sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink.
type	thrift	Specifies the type of the thrift sink, which must be set to thrift .

Parameter	Default Value	Description
hostname	-	Specifies the bound host name or IP address.
port	-	Specifies the bound listening port. Ensure that this port is not occupied.
batch-size	1000	Specifies the number of events sent in a batch.
connect-timeout	20000	Specifies the timeout for the first connection, expressed in milliseconds.
request-timeout	20000	Specifies the maximum timeout for a request after the first request, expressed in milliseconds.
kerberos	false	Specifies whether Kerberos authentication is enabled.
client-keytab	-	Specifies the path of the client keytab file. The Flume running user must have the access permission on the authentication file.
client-principal	-	Specifies the principal of the security user used by the client.
server-principal	-	Specifies the principal of the security user used by the server.
compression-type	none	Specifies the compression type of data sent by Flume, which can be set to none or deflate . none indicates that data is not compressed, while deflate indicates that data is compressed.

Parameter	Default Value	Description
maxConnections	5	Specifies the maximum size of the connection pool for Flume to send data.
ssl	false	Specifies whether to use SSL encryption.
truststore-type	JKS	Specifies the Java trust store type.
truststore	-	Specifies the Java trust store file.
truststore-password	-	Specifies the Java trust store password.
reset-connection-interval	0	Specifies the interval between a connection failure and a second connection, expressed in seconds. If the parameter is set to 0 , the system continuously attempts to perform a connection.

Precautions

- What are the reliability measures of Flume?
 - Use the transaction mechanisms between Source and Channel as well as between Channel and Sink.
 - Configure the failover and load_balance mechanisms for Sink Processor. The following shows a load balancing example.


```
server.sinkgroups=g1
server.sinkgroups.g1.sinks=k1 k2
server.sinkgroups.g1.processor.type=load_balance
server.sinkgroups.g1.processor.backoff=true
server.sinkgroups.g1.processor.selector=random
```
- What are the precautions for the aggregation and cascading of multiple Flume agents?
 - Avro or Thrift protocol can be used for cascading.
 - When the aggregation end contains multiple nodes, evenly distribute the agents and do not aggregate all agents on a single node.

7.8 Flume Configuration Parameter Description

For versions earlier than MRS 3.x, configure Flume parameters in the **properties.properties** file.

For MRS 3.x or later, some parameters can be configured on Manager.

Overview

This section describes how to configure the sources, channels, and sinks of Flume, and modify the configuration items of each module.

For MRS 3.x or later, log in to FusionInsight Manager and choose **Cluster > Services > Flume**. On the displayed page, click the **Configuration Tool** tab, select and drag the source, channel, and sink to be used to the GUI on the right, and double-click them to configure corresponding parameters. Parameters such as **channels** and **type** are configured only in the client configuration file **properties.properties**, the path of which is *Flume client installation directory/fusioninsight-flume-Flume version/conf/properties.properties*.

NOTE

You must input encrypted information for some configurations. For details on how to encrypt information, see [Using the Encryption Tool of the Flume Client](#).

Common Source Configurations

- **Avro Source**

An Avro source listens to the Avro port, receives data from the external Avro client, and places data into configured channels. [Table 7-16](#) lists common configurations.

Table 7-16 Common configurations of an Avro source

Parameter	Default Value	Description
channels	-	<p>Specifies the channel connected to the source. Multiple channels can be configured. Use spaces to separate them.</p> <p>In a single proxy process, sources and sinks are connected through channels. A source instance corresponds to multiple channels, but a sink instance corresponds only to one channel.</p> <p>The format is as follows:</p> <pre><Agent >.sources.<Source>.channels = <channel1> <channel2> <channel3>...</pre> <pre><Agent >.sinks.<Sink>.channels = <channel1></pre> <p>This parameter can be configured only in the properties.properties file.</p>

Parameter	Default Value	Description
type	avro	Specifies the type, which is set to avro . The type of each source is a fixed value. This parameter can be configured only in the properties.properties file.
bind	-	Specifies the host name or IP address associated with the source.
port	-	Specifies the bound port number.
ssl	false	Specifies whether to use SSL encryption. <ul style="list-style-type: none"> • true • false
truststore-type	JKS	Specifies the Java trust store type. Set this parameter to JKS or other truststore types supported by Java.
truststore	-	Specifies the Java trust store file.
truststore-password	-	Specifies the Java trust store password.
keystore-type	JKS	Specifies the key storage type. Set this parameter to JKS or other truststore types supported by Java.
keystore	-	Specifies the key storage file.
keystore-password	-	Specifies the key storage password.

- **SpoolDir Source**

A SpoolDir source monitors and transmits new files that have been added to directories in quasi-real-time mode. Common configurations are as follows:

Table 7-17 Common configurations of a SpoolDir source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured. This parameter can be configured only in the properties.properties file.

Parameter	Default Value	Description
type	spooldir	Type, which is set to spooldir . This parameter can be configured only in the properties.properties file.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the source is restarted. Unit: second
spoolDir	-	Specifies the monitoring directory.
fileSuffix	.COMPLETED	Specifies the suffix added after file transmission is complete.
deletePolicy	never	Specifies the source file deletion policy after file transmission is complete. The value can be either never or immediate .
ignorePattern	^\$	Specifies the regular expression of a file to be ignored.
trackerDir	.flumespool	Specifies the metadata storage path during data transmission.
batchSize	1000	Specifies the source transmission granularity.
decodeErrorPolicy	FAIL	<p>Specifies the code error policy. This parameter can be configured only in the properties.properties file.</p> <p>The value can be FAIL, REPLACE, or IGNORE.</p> <p>FAIL: Generate an exception and fail the parsing.</p> <p>REPLACE: Replace the characters that cannot be identified with other characters, such as U+FFFD.</p> <p>IGNORE: Discard character strings that cannot be parsed.</p> <p>NOTE If a code error occurs in the file, set decodeErrorPolicy to REPLACE or IGNORE. Flume will skip the code error and continue to collect subsequent logs.</p>

Parameter	Default Value	Description
deserializer	LINE	Specifies the file parser. The value can be either LINE or BufferedLine . <ul style="list-style-type: none"> When the value is set to LINE, characters read from the file are transcoded one by one. When the value is set to BufferedLine, one line or multiple lines of characters read from the file are transcoded in batches, which delivers better performance.
deserializer.maxLineLength	2048	Specifies the maximum length for resolution by line, ranging from 0 to 2,147,483,647.
deserializer.maxBatchLine	1	Specifies the maximum number of lines for resolution by line. If multiple lines are set, maxLineLength must be set to a corresponding multiplier. For example, if maxBatchLine is set to 2 , maxLineLength is set to 4096 (2048 x 2).
selector.type	replicating	Specifies the selector type. The value can be either replicating or multiplexing . <ul style="list-style-type: none"> replicating indicates that the same content is sent to each channel. multiplexing indicates that the content is sent only to certain channels according to the distribution rule.
interceptors	-	Specifies the interceptor. This parameter can be configured only in the properties.properties file.

 **NOTE**

The Spooling source ignores the last line feed character of each event when data is read by line. Therefore, Flume does not calculate the data volume counters used by the last line feed character.

- **Kafka Source**

A Kafka source consumes data from Kafka topics. Multiple sources can consume data of the same topic, and the sources consume different partitions of the topic. Common configurations are as follows:

Table 7-18 Common configurations of a Kafka source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured. This parameter can be configured only in the properties.properties file.
type	org.apache.flume.source.kafka.KafkaSource	Specifies the type, which is set to org.apache.flume.source.kafka.KafkaSource . This parameter can be configured only in the properties.properties file.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the source is restarted. Unit: second
nodatotime	0 (Disabled)	Specifies the alarm threshold. An alarm is triggered when the duration that Kafka does not release data to subscribers exceeds the threshold. Unit: second
batchSize	1000	Specifies the number of events written into a channel at a time.
batchDurationMillis	1000	Specifies the maximum duration of topic data consumption at a time, expressed in milliseconds.
keepTopicInHeader	false	Specifies whether to save topics in the event header. If topics are saved, topics configured in Kafka sinks become invalid. <ul style="list-style-type: none"> • true • false This parameter can be configured only in the properties.properties file.

Parameter	Default Value	Description
keepPartitionIn-Header	false	Specifies whether to save partition IDs in the event header. If partition IDs are saved, Kafka sinks write data to the corresponding partitions. <ul style="list-style-type: none"> • true • false This parameter can be set only in the properties.properties file.
kafka.bootstrap.servers	-	Specifies the list of Broker addresses, which are separated by commas.
kafka.consumer.group.id	-	Specifies the Kafka consumer group ID.
kafka.topics	-	Specifies the list of subscribed Kafka topics, which are separated by commas (,).
kafka.topics.regex	-	Specifies the subscribed topics that comply with regular expressions. kafka.topics.regex has a higher priority than kafka.topics and will overwrite kafka.topics .
kafka.security.protocol	SASL_PLAINTEXT	Specifies the security protocol of Kafka. The value must be set to PLAINTEXT for clusters in which Kerberos authentication is disabled.
kafka.kerberos.domain.name	-	Specifies the value of default_realm of Kerberos in the Kafka cluster, which should be configured only for security clusters. This parameter can be set only in the properties.properties file.
Other Kafka Consumer Properties	-	Specifies other Kafka configurations. This parameter can be set to any consumption configuration supported by Kafka, and the .kafka prefix must be added to the configuration. This parameter can be set only in the properties.properties file.

- **Taildir Source**

A Taildir source monitors file changes in a directory and automatically reads the file content. In addition, it can transmit data in real time. [Table 7-19](#) lists common configurations.

Table 7-19 Common configurations of a Taildir source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured. This parameter can be set only in the properties.properties file.
type	taildir	Specifies the type, which is set to taildir . This parameter can be set only in the properties.properties file.
filegroups	-	Specifies the group name of a collection file directory. Group names are separated by spaces.
filegroups.<filegroup Name>.parentDir	-	Specifies the parent directory. The value must be an absolute path. This parameter can be set only in the properties.properties file.
filegroups.<filegroup Name>.filePattern	-	Specifies the relative file path of the file group's parent directory. Directories can be included and regular expressions are supported. It must be used together with parentDir . This parameter can be set only in the properties.properties file.
positionFile	-	Specifies the metadata storage path during data transmission.
headers.<filegroup Name>.<headerKey>	-	Specifies the key-value of an event when data of a group is being collected. This parameter can be set only in the properties.properties file.
byteOffsetHeader	false	Specifies whether each event header should contain the location information about the event in the source file. The location information is saved in the byteoffset variable.

Parameter	Default Value	Description
skipToEnd	false	Specifies whether Flume can locate the latest location of a file and read the latest data after restart.
idleTimeout	120000	Specifies the idle duration during file reading, expressed in milliseconds. If the file data is not changed in this idle period, the source closes the file. If data is written into this file after it is closed, the source opens the file and reads data.
writePosInterval	3000	Specifies the interval for writing metadata to a file, expressed in milliseconds.
batchSize	1000	Specifies the number of events written to the channel in batches.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the source is restarted. Unit: second

- **Http Source**

An HTTP source receives data from an external HTTP client and sends the data to the configured channels. [Table 7-20](#) lists common configurations.

Table 7-20 Common configurations of an HTTP source

Parameter	Default Value	Description
channels	-	Specifies the channel connected to the source. Multiple channels can be configured. This parameter can be set only in the properties.properties file.
type	http	Specifies the type, which is set to http . This parameter can be set only in the properties.properties file.
bind	-	Specifies the name or IP address of the bound host.
port	-	Specifies the bound port.

Parameter	Default Value	Description
handler	org.apache.flume.source.http.JSONHandler	Specifies the message parsing method of an HTTP request. The following methods are supported: <ul style="list-style-type: none"> org.apache.flume.source.http.JSONHandler: JSON org.apache.flume.sink.solr.morphline.BlobHandler: BLOB
handler.*	-	Specifies handler parameters.
enableSSL	false	Specifies whether SSL is enabled in HTTP.
keystore	-	Specifies the keystore path set after SSL is enabled in HTTP.
keystorePassword	-	Specifies the keystore password set after SSL is enabled in HTTP.

Common Channel Configurations

- **Memory Channel**

A memory channel uses memory as the cache. Events are stored in memory queues. [Table 7-21](#) lists common configurations.

Table 7-21 Common configurations of a memory channel

Parameter	Default Value	Description
type	-	Specifies the type, which is set to memory . This parameter can be set only in the properties.properties file.
capacity	10000	Specifies the maximum number of events cached in a channel.
transactionCapacity	1000	Specifies the maximum number of events accessed each time.
channelFullcount	10	Specifies the channel full count. When the count reaches the threshold, an alarm is reported.

- **File Channel**

A file channel uses local disks as the cache. Events are stored in the folder specified by **dataDirs**. [Table 7-22](#) lists common configurations.

Table 7-22 Common configurations of a file channel

Parameter	Default Value	Description
type	-	Specifies the type, which is set to file . This parameter can be set only in the properties.properties file.
checkpointDir	\${BIGDATA_DATA_HOME}/flume/checkpoint	Specifies the checkpoint storage directory.
dataDirs	\${BIGDATA_DATA_HOME}/flume/data	Specifies the data cache directory. Multiple directories can be configured to improve performance. The directories are separated by commas (,).
maxFileSize	2146435071	Specifies the maximum size of a single cache file, expressed in bytes.
minimumRequired-Space	524288000	Specifies the minimum idle space in the cache, expressed in bytes.
capacity	1000000	Specifies the maximum number of events cached in a channel.
transactionCapacity	10000	Specifies the maximum number of events accessed each time.
channelFullcount	10	Specifies the channel full count. When the count reaches the threshold, an alarm is reported.

- **Kafka Channel**

A Kafka channel uses a Kafka cluster as the cache. Kafka provides high availability and multiple copies to prevent data from being immediately consumed by sinks when Flume or Kafka Broker crashes. [Table 10 Common configurations of a Kafka channel](#) lists common configurations.

Table 7-23 Common configurations of a Kafka channel

Parameter	Default Value	Description
type	-	Specifies the type, which is set to org.apache.flume.channel.kafka.KafkaChannel . This parameter can be set only in the properties.properties file.

Parameter	Default Value	Description
kafka.bootstrap.servers	-	Specifies the list of Brokers in the Kafka cluster.
kafka.topic	flume-channel	Specifies the Kafka topic used by the channel to cache data.
kafka.consumer.group.id	flume	Specifies the Kafka consumer group ID.
parseAsFlumeEvent	true	Specifies whether data is parsed into Flume events.
migrateZookeeper-Offsets	true	Specifies whether to search for offsets in ZooKeeper and submit them to Kafka when there is no offset in Kafka.
kafka.consumer.auto.offset.reset	latest	Consumes data from the specified location when there is no offset.
kafka.producer.security.protocol	SASL_PLAINTEXT	Specifies the Kafka producer security protocol.
kafka.consumer.security.protocol	SASL_PLAINTEXT	Specifies the Kafka consumer security protocol.

Common Sink Configurations

- **HDFS Sink**

An HDFS sink writes data into HDFS. [Table 7-24](#) lists common configurations.

Table 7-24 Common configurations of an HDFS sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink. This parameter can be set only in the properties.properties file.
type	hdfs	Specifies the type, which is set to hdfs . This parameter can be set only in the properties.properties file.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the sink is restarted. Unit: second
hdfs.path	-	Specifies the HDFS path.

Parameter	Default Value	Description
hdfs.inUseSuffix	.tmp	Specifies the suffix of the HDFS file to which data is being written.
hdfs.rollInterval	30	Specifies the interval for file rolling, expressed in seconds.
hdfs.rollSize	1024	Specifies the size for file rolling, expressed in bytes.
hdfs.rollCount	10	Specifies the number of events for file rolling.
hdfs.idleTimeout	0	Specifies the timeout interval for closing idle files automatically, expressed in seconds.
hdfs.batchSize	1000	Specifies the number of events written into HDFS at a time.
hdfs.kerberosPrincipal	-	Specifies the Kerberos username for HDFS authentication. This parameter is not required for a cluster in which Kerberos authentication is disabled.
hdfs.kerberosKeytab	-	Specifies the Kerberos keytab of HDFS authentication. This parameter is not required for a cluster in which Kerberos authentication is disabled.
hdfs.fileCloseByEvent	true	Specifies whether to close the file when the last event is received.
hdfs.batchCallTimeout	-	<p>Specifies the timeout control duration each time events are written into HDFS, expressed in milliseconds.</p> <p>If this parameter is not specified, the timeout duration is controlled when each event is written into HDFS. When the value of hdfs.batchSize is greater than 0, configure this parameter to improve the performance of writing data into HDFS.</p> <p>NOTE The value of hdfs.batchCallTimeout depends on hdfs.batchSize. A greater hdfs.batchSize requires a larger hdfs.batchCallTimeout. If the value of hdfs.batchCallTimeout is too small, writing events to HDFS may fail.</p>

Parameter	Default Value	Description
serializer.appendNewLine	true	Specifies whether to add a line feed character (\n) after an event is written to HDFS. If a line feed character is added, the data volume counters used by the line feed character will not be calculated by HDFS sinks.

- **Avro Sink**

An Avro sink converts events into Avro events and sends them to the monitoring ports of the hosts. [Table 7-25](#) lists common configurations.

Table 7-25 Common configurations of an Avro sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink. This parameter can be set only in the properties.properties file.
type	-	Specifies the type, which is set to avro . This parameter can be set only in the properties.properties file.
hostname	-	Specifies the name or IP address of the bound host.
port	-	Specifies the monitoring port.
batch-size	1000	Specifies the number of events sent in a batch.
ssl	false	Specifies whether to use SSL encryption.
truststore-type	JKS	Specifies the Java trust store type.
truststore	-	Specifies the Java trust store file.
truststore-password	-	Specifies the Java trust store password.
keystore-type	JKS	Specifies the key storage type.
keystore	-	Specifies the key storage file.
keystore-password	-	Specifies the key storage password.

- **HBase Sink**

An HBase sink writes data into HBase. [Table 7-26](#) lists common configurations.

Table 7-26 Common configurations of an HBase sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink. This parameter can be set only in the properties.properties file.
type	-	Specifies the type, which is set to hbase . This parameter can be set only in the properties.properties file.
table	-	Specifies the HBase table name.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the sink is restarted. Unit: second
columnFamily	-	Specifies the HBase column family.
batchSize	1000	Specifies the number of events written into HBase at a time.
kerberosPrincipal	-	Specifies the Kerberos username for HBase authentication. This parameter is not required for a cluster in which Kerberos authentication is disabled.
kerberosKeytab	-	Specifies the Kerberos keytab of HBase authentication. This parameter is not required for a cluster in which Kerberos authentication is disabled.

- **Kafka Sink**

A Kafka sink writes data into Kafka. [Table 7-27](#) lists common configurations.

Table 7-27 Common configurations of a Kafka sink

Parameter	Default Value	Description
channel	-	Specifies the channel connected to the sink. This parameter can be set only in the properties.properties file.
type	-	Specifies the type, which is set to org.apache.flume.sink.kafka.Kafka Sink . This parameter can be set only in the properties.properties file.

Parameter	Default Value	Description
kafka.bootstrap.servers	-	Specifies the list of Kafka Brokers, which are separated by commas.
monTime	0 (Disabled)	Specifies the thread monitoring threshold. When the update time exceeds the threshold, the sink is restarted. Unit: second
kafka.topic	default-flume-topic	Specifies the topic where data is written.
flumeBatchSize	1000	Specifies the number of events written into Kafka at a time.
kafka.security.protocol	SASL_PLAINTEXT	Specifies the security protocol of Kafka. The value must be set to PLAINTEXT for clusters in which Kerberos authentication is disabled.
kafka.kerberos.domain.name	-	Specifies the Kafka domain name. This parameter is mandatory for a security cluster. This parameter can be set only in the properties.properties file.
Other Kafka Producer Properties	-	Specifies other Kafka configurations. This parameter can be set to any production configuration supported by Kafka, and the .kafka prefix must be added to the configuration. This parameter can be set only in the properties.properties file.

7.9 Using Environment Variables in the properties.properties File

Scenario

This section describes how to use environment variables in the **properties.properties** configuration file.

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

Prerequisites

The Flume service is running properly and the Flume client has been installed.

Procedure

Step 1 Log in to the node where the Flume client is installed as user **root**.

Step 2 Switch to the following directory:

```
cd Flume client installation directory/fusioninsight-flume-Flume component version/conf
```

Step 3 Add environment variables to the **flume-env.sh** file in the directory.

- Format:

```
export Variable name=Variable value
```

- Example:

```
JAVA_OPTS="-Xms2G -Xmx4G -XX:CMSFullGCsBeforeCompaction=1 -XX:+UseConcMarkSweepGC -  
XX:+CMSParallelRemarkEnabled -XX:+UseCMSCompactAtFullCollection -  
DpropertiesImplementation=org.apache.flume.node.EnvVarResolverProperties"  
export TAILDIR_PATH=/tmp/flumetest/201907/20190703/1/*.log.*
```

Step 4 Restart the Flume instance process.

1. Log in to FusionInsight Manager.
2. Choose **Cluster** > **Services** > **Flume**. On the page that is displayed, click the **Instance** tab, select all Flume instances, and choose **More** > **Restart Instance**. In the displayed **Verify Identity** dialog box, enter the password, and click **OK**.

NOTICE

Do not restart the Flume service on FusionInsight Manager after **flume-env.sh** takes effect on the server. Otherwise, the user-defined environment variables will be lost. You only need to restart the corresponding instances on FusionInsight Manager.

Step 5 In the *Flume client installation directory*/fusioninsight-flume-*Flume component version number*/conf/**properties.properties** configuration file, reference variables in the **$\${Variable name}$** format. The following is an example:

```
client.sources.s1.type = TAILDIR  
client.sources.s1.filegroups = f1  
client.sources.s1.filegroups.f1 =  $\${TAILDIR\_PATH}$   
client.sources.s1.positionFile = /tmp/flumetest/201907/20190703/1/taildir_position.json  
client.sources.s1.channels = c1
```

NOTICE

- Ensure that **flume-env.sh** takes effect before you go to [Step 5](#) to configure the **properties.properties** file.
- If you configure file on the local host, upload the file on FusionInsight Manager by performing the following steps. The user-defined environment variables may be lost if the operations are not performed in the correct sequence.
 1. Log in to FusionInsight Manager.
 2. Choose **Cluster > Services > Flume**. On the page that is displayed, click the **Configurations** tab, select the Flume instance, and click **Upload File** next to **flume.config.file** to upload the **properties.properties** file.

----End

7.10 Non-Encrypted Transmission

7.10.1 Configuring Non-encrypted Transmission

Scenario

This section describes how to configure Flume server and client parameters after the cluster and the Flume service are installed to ensure proper running of the service.

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

NOTE

By default, the cluster network environment is secure and the SSL authentication is not enabled during the data transmission process. For details about how to use the encryption mode, see [Configuring the Encrypted Transmission](#).

Prerequisites

- The Flume client has been installed.
- The cluster and Flume service have been installed.
- The network environment of the cluster is secure.

Procedure

Step 1 Configure the client parameters of the Flume role.

1. Use the Flume configuration tool on FusionInsight Manager to configure the Flume role client parameters and generate a configuration file.
 - a. Log in to FusionInsight Manager. Choose **Cluster > Services > Flume > Configuration Tool**.

Figure 7-2 Clicking Configuration Tool



- b. Set **Agent Name** to **client**. Select and drag the source, channel, and sink to be used to the GUI on the right, and connect them. For example, use SpoolDir Source, File Channel, and Avro Sink, as shown in [Figure 7-3](#).

Figure 7-3 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-28](#) based on the actual environment.

NOTE

- If the client parameters of the Flume role have been configured, you can obtain the existing client parameter configuration file from *client installation directory/fusioninsight-flume-1.9.0/conf/properties.properties* to ensure that the configuration is in concordance with the previous. Log in to FusionInsight Manager, choose **Cluster > Services > Flume > Configuration > Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.

Table 7-28 Parameters to be modified for the Flume role client


Parameter	Description	Example Value
ssl	<p>Specifies whether to enable the SSL authentication. (You are advised to enable this function to ensure security.)</p> <p>Only Sources of the Avro type have this configuration item.</p> <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the function is not enabled. 	false

- d. Click **Export** to save the **properties.properties** configuration file to the local server.
2. Upload the **properties.properties** file to **flume/conf/** under the installation directory of the Flume client.

Step 2 Configure the server parameters of the Flume role and upload the configuration file to the cluster.


1. Use the Flume configuration tool on the FusionInsight Manager portal to configure the server parameters and generate the configuration file.
 - a. Log in to FusionInsight Manager. Choose **Cluster > Services > Flume > Configuration Tool**.

Figure 7-4 Clicking Configuration Tool



- b. Set **Agent Name** to **server**. Select and drag the source, channel, and sink to be used to the GUI on the right, and connect them. For example, use Avro Source, File Channel, and HDFS Sink, as shown in [Figure 7-5](#).

Figure 7-5 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-29](#) based on the actual environment.

 NOTE

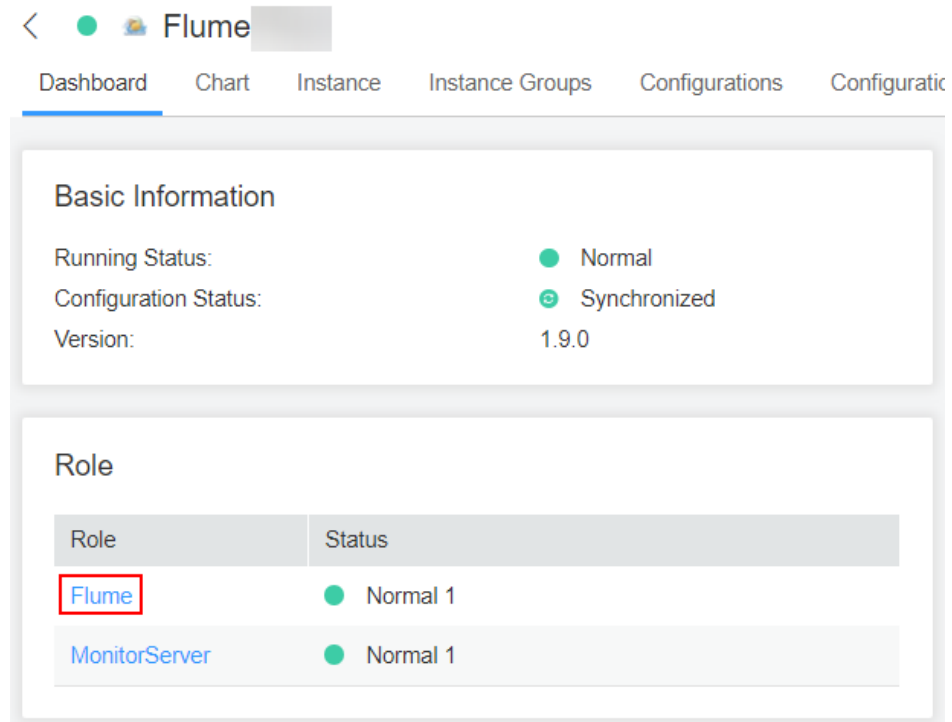
- If the server parameters of the Flume role have been configured, you can choose **Cluster > Services > Flume > Instance** on FusionInsight Manager. Then select the corresponding Flume role instance and click the **Download** button behind the **flume.config.file** parameter on the **Instance Configurations** page to obtain the existing server parameter configuration file. Choose **Cluster > Service > Flume > Configurations > Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.
- A unique checkpoint directory needs to be configured for each File Channel.

Table 7-29 Parameters to be modified for the Flume role server

Parameter	Description	Example Value
ssl	<p>Specifies whether to enable the SSL authentication. (You are advised to enable this function to ensure security.)</p> <p>Only Sources of the Avro type have this configuration item.</p> <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the function is not enabled. 	false

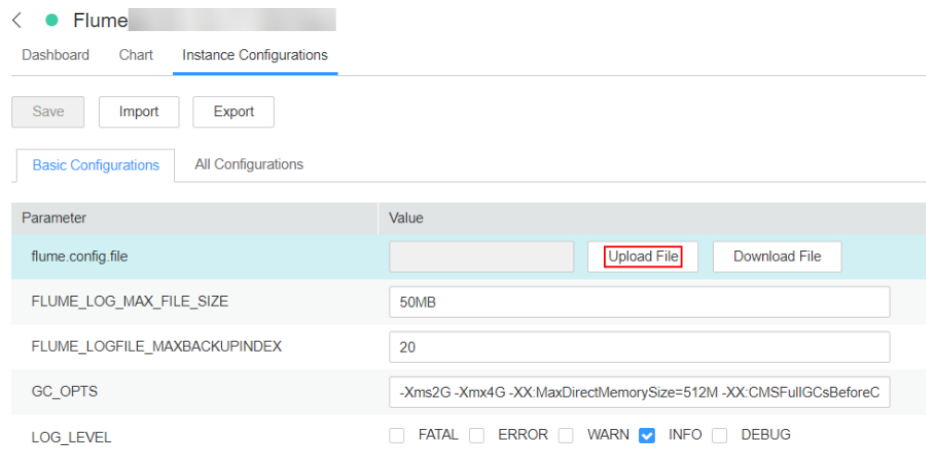
- d. Click **Export** to save the **properties.properties** configuration file to the local server.
2. Log in to FusionInsight Manager and choose **Cluster > Services > Flume**. On the **Instances** tab page, click **Flume**.

Figure 7-6 Clicking the Flume role



3. Select the Flume role of the node where the configuration file is to be uploaded, choose **Instance Configurations > Import** beside the **flume.config.file**, and select the **properties.properties** file.

Figure 7-7 Uploading a file



NOTE

- An independent server configuration file can be uploaded to each Flume instance.
 - This step is required for updating the configuration file. Modifying the configuration file on the background is an improper operation because the modification will be overwritten after configuration synchronization.
4. Click **Save**, and then click **OK**.

5. Click **Finish**.
- End

7.10.2 Typical Scenario: Collecting Local Static Logs and Uploading Them to Kafka

Scenario

This section describes how to use the Flume server to collect static logs from a local host and save them to the topic list (test1) of Kafka.

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

NOTE

By default, the cluster network environment is secure and the SSL authentication is not enabled during the data transmission process. For details about how to use the encryption mode, see [Configuring the Encrypted Transmission](#). The configuration applies to scenarios where only the Flume is configured, for example, SpoolDir Source+Memory Channel+Kafka Sink.

Prerequisites

- The cluster has been installed, including the Kafka and Flume services.
- The network environment of the cluster is secure.
- The MRS cluster administrator has understood service requirements and prepared Kafka administrator `flume_kafka`.

Procedure


Step 1 Set Flume parameters.

Use the Flume configuration tool on Manager to configure the Flume role server parameters and generate a configuration file.

1. Log in to FusionInsight Manager. Choose **Cluster > Services > Flume > Configuration Tool**.
2. Set **Agent Name** to `client`. Select and drag the source, channel, and sink to be used to the GUI on the right, and connect them.

Use SpoolDir Source, Memory Channel, and Kafka Sink.

Figure 7-8 Example for the Flume configuration tool



3. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-30](#) based on the actual environment.

 **NOTE**

- If you want to continue using the **properties.properties** file by modifying it, log in to FusionInsight Manager, choose **Cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab, click **Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.

Table 7-30 Parameters to be modified for the Flume role server

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
spoolDir	Specifies the directory where the file to be collected resides. This parameter cannot be left blank. The directory needs to exist and have the write, read, and execute permissions on the flume running user.	/srv/BigData/hadoop/data1/zb
trackerDir	Specifies the path for storing the metadata of files collected by Flume.	/srv/BigData/hadoop/data1/tracker
batchSize	Specifies the number of events that Flume sends in a batch (number of data pieces). A larger value indicates higher performance and lower timeliness.	61200
kafka.topics	Specifies the list of subscribed Kafka topics, which are separated by commas (.). This parameter cannot be left blank.	test1
kafka.bootstrap.servers	Specifies the bootstrap IP address and port list of Kafka. The default value is all Kafkabrokers in the Kafka cluster.	192.168.101.10:21007

4. Click **Export** to save the **properties.properties** configuration file to the local server.

Step 2 Upload the configuration file.

Upload the file exported in [Step 1.4](#) to the `flume/conf` directory of the cluster. For details, see [Step 2.2](#).

Step 3 Verify log transmission.

1. Log in to the Kafka client.

```
cd Kafka client installation directory/Kafka/kafka
```

```
kinit flume_kafka (Enter the password.)
```

2. Read data from a Kafka topic.

```
bin/kafka-console-consumer.sh --topic topic name --bootstrap-server Kafka service IP address of the node where the role instance is located:21007 --consumer.config config/consumer.properties --from-beginning
```

The system displays the contents of the file to be collected.

```
[root@host1 kafka]# bin/kafka-console-consumer.sh --topic test1 --bootstrap-server 192.168.101.10:21007 --consumer.config config/consumer.properties --from-beginning
Welcome to flume
```

----End

7.10.3 Typical Scenario: Collecting Local Static Logs and Uploading Them to HDFS

Scenario

This section describes how to use the Flume server to collect static logs from a local host and save them to the `/flume/test` directory on HDFS.

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

NOTE

By default, the cluster network environment is secure and the SSL authentication is not enabled during the data transmission process. For details about how to use the encryption mode, see [Configuring the Encrypted Transmission](#). The configuration applies to scenarios where only the Flume is configured, for example, Spooldir Source+Memory Channel+HDFS Sink.

Prerequisites

- The cluster has been installed, including the HDFS and Flume services.
- The network environment of the cluster is secure.
- User `flume_hdfs` has been created, and the HDFS directory and data used for log verification have been authorized to the user.

Procedure

Step 1 On FusionInsight Manager, choose **System > Permission > User**, select user `flume_hdfs`, and choose **More > Download Authentication Credential** to download the Kerberos certificate file of user `flume_hdfs` and save it to the local host.

Step 2 Set Flume parameters.

Use Flume on FusionInsight Manager to configure the Flume role server parameters and generate a configuration file.

1. Log in to FusionInsight Manager. Choose **Cluster > Services > Flume > Configuration Tool**.
2. Set **Agent Name** to **client**. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.

Use SpoolDir Source, Memory Channel, and HDFS Sink.

Figure 7-9 Example for the Flume configuration tool



3. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-31](#) based on the actual environment.

NOTE

- If you want to continue using the **properties.properties** file by modifying it, log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab, click **Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.

Table 7-31 Parameters to be modified for the Flume role server

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
spoolDir	Specifies the directory where the file to be collected resides. This parameter cannot be left blank. The directory needs to exist and have the write, read, and execute permissions on the flume running user.	/srv/BigData/hadoop/data1/zb
trackerDir	Specifies the path for storing the metadata of files collected by Flume.	/srv/BigData/hadoop/data1/tracker
batchSize	Specifies the number of events that Flume sends in a batch.	61200

Parameter	Description	Example Value
hdfs.path	Specifies the HDFS data write directory. This parameter cannot be left blank.	hdfs://hacluster/flume/test
hdfs.filePrefix	Specifies the file name prefix after data is written to HDFS.	TMP_
hdfs.batchSize	Specifies the maximum number of events that can be written to HDFS once.	61200
hdfs.kerberosPrincipal	Specifies the Kerberos authentication user, which is mandatory in security versions. This configuration is required only in security clusters.	flume_hdfs
hdfs.kerberosKeytab	Specifies the keytab file path for Kerberos authentication, which is mandatory in security versions. This configuration is required only in security clusters.	/opt/test/conf/user.keytab NOTE Obtain the user.keytab file from the Kerberos certificate file of the user flume_hdfs . In addition, ensure that the user who installs and runs the Flume client has the read and write permissions on the user.keytab file.
hdfs.useLocalTimeStamp	Specifies whether to use the local time. Possible values are true and false .	true

4. Click **Export** to save the **properties.properties** configuration file to the local.


Step 3 Upload the configuration file.

Upload the file exported in [Step 2.4](#) to the **flume/conf** directory of the cluster. For details, see [Step 2.2](#).

Step 4 Verify log transmission.

1. Log in to FusionInsight Manager as a user who has the management permission on HDFS. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Services > HDFS**. On the page that is displayed, click the **NameNode(Node name,Active)** link next to **NameNode WebUI** to go to the HDFS web UI. On the displayed page, choose **Utilities > Browse the file system**.
2. Check whether the data is generated in the **/flume/test** directory on the HDFS.

Figure 7-10 Checking HDFS directories and files



----End

7.10.4 Typical Scenario: Collecting Local Dynamic Logs and Uploading Them to HDFS

Scenario

This section describes how to use the Flume server to collect dynamic logs from a local host and save them to the **/flume/test** directory on HDFS.

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

NOTE

By default, the cluster network environment is secure and the SSL authentication is not enabled during the data transmission process. For details about how to use the encryption mode, see [Configuring the Encrypted Transmission](#). The configuration applies to scenarios where only the Flume is configured, for example, Taildir Source+Memory Channel+HDFS Sink.

Prerequisites

- The cluster has been installed, including the HDFS and Flume services.
- The network environment of the cluster is secure.
- You have created user **flume_hdfs** and authorized the HDFS directory and data to be operated during log verification.

Procedure

Step 1 On FusionInsight Manager, choose **System > User** and choose **More > Download Authentication Credential** to download the Kerberos certificate file of user **flume_hdfs** and save it to the local host.


Step 2 Set Flume parameters.

Use the Flume configuration tool on FusionInsight Manager to configure the Flume role server parameters and generate a configuration file.

1. Log in to FusionInsight Manager and choose **Cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab.
2. Set **Agent Name** to **client**. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.

Use Taildir Source, Memory Channel, and HDFS Sink.

Figure 7-11 Example for the Flume configuration tool



3. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-32](#) based on the actual environment.

NOTE

- If you want to continue using the **properties.propertites** file by modifying it, log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab, click **Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.

Table 7-32 Parameters to be modified for the Flume role server

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
filegroups	Specifies the file group list name. This parameter cannot be left blank. The value contains the following two parts: <ul style="list-style-type: none"> - Name: name of the file group list. - filegroups: absolute path of dynamic log files. 	-
positionFile	Specifies the location where the collected file information (file name and location from which the file collected) is saved. This parameter cannot be left blank. The file does not need to be created manually, but the Flume running user needs to have the write permission on its upper-level directory.	/home/omm/flume/positionfile
batchSize	Specifies the number of events that Flume sends in a batch.	61200

Parameter	Description	Example Value
hdfs.path	Specifies the HDFS data write directory. This parameter cannot be left blank.	hdfs://hacluster/flume/test
hdfs.filePrefix	Specifies the file name prefix after data is written to HDFS.	TMP_
hdfs.batchSize	Specifies the maximum number of events that can be written to HDFS once.	61200
hdfs.kerberosPrincipal	Specifies the Kerberos authentication user, which is mandatory in security versions. This configuration is required only in security clusters.	flume_hdfs
hdfs.kerberosKeytab	Specifies the keytab file path for Kerberos authentication, which is mandatory in security versions. This configuration is required only in security clusters.	/opt/test/conf/user.keytab NOTE Obtain the user.keytab file from the Kerberos certificate file of the user flume_hdfs . In addition, ensure that the user who installs and runs the Flume client has the read and write permissions on the user.keytab file.
hdfs.useLocalTimeStamp	Specifies whether to use the local time. Possible values are true and false .	true

4. Click **Export** to save the **properties.properties** configuration file to the local.


Step 3 Upload the configuration file.

Upload the file exported in [Step 2.4](#) to the **flume/conf** directory of the cluster. For details, see [Step 2.2](#).

Step 4 Verify log transmission.

1. Log in to FusionInsight Manager as a user who has the management permission on HDFS. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Services > HDFS**. On the page that is displayed, click the **NameNode(Node name,Active)** link next to **NameNode WebUI** to go to the HDFS web UI. On the displayed page, choose **Utilities > Browse the file system**.
2. Check whether the data is generated in the **/flume/test** directory on the HDFS.

Figure 7-12 Checking HDFS directories and files



----End

7.10.5 Typical Scenario: Collecting Logs from Kafka and Uploading Them to HDFS

Scenario

This section describes how to use the Flume server to collect logs from the topic list (test1) of Kafka and save them to the `/flume/test` directory on HDFS.

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

NOTE

By default, the cluster network environment is secure and the SSL authentication is not enabled during the data transmission process. For details about how to use the encryption mode, see [Configuring the Encrypted Transmission](#). The configuration applies to scenarios where only the Flume is configured, for example, Kafka Source+Memory Channel+HDFS Sink.

Prerequisites

- The cluster has been installed, including the HDFS, Kafka, and Flume services.
- The network environment of the cluster is secure.
- You have created user `flume_hdfs` and authorized the HDFS directory and data to be operated during log verification.

Procedure

Step 1 On FusionInsight Manager, choose **System > User** and choose **More > Download Authentication Credential** to download the Kerberos certificate file of user `flume_hdfs` and save it to the local host.


Step 2 Configure the server parameters of the Flume role.

Use the Flume configuration tool on FusionInsight Manager to configure the Flume role server parameters and generate a configuration file.

1. Log in to FusionInsight Manager and choose **Cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab.
2. Set **Agent Name** to `client`. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.

For example, use Kafka Source, Memory Channel, and HDFS Sink.

Figure 7-13 Example for the Flume configuration tool



3. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-33](#) based on the actual environment.

NOTE

- If you want to continue using the **properties.propertites** file by modifying it, log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab, click **Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.

Table 7-33 Parameters to be modified for the Flume role server

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
kafka.topics	Specifies the subscribed Kafka topic list, in which topics are separated by commas (.). This parameter cannot be left blank.	test1
kafka.consumer.group.id	Specifies the data group ID obtained from Kafka. This parameter cannot be left blank.	flume
kafka.bootstrap.servers	Specifies the bootstrap IP address and port list of Kafka. The default value is all Kafka lists in a Kafka cluster. If Kafka has been installed in the cluster and its configurations have been synchronized, this parameter can be left blank.	192.168.101.10:9092
batchSize	Specifies the number of events that Flume sends in a batch (number of data pieces).	61200

Parameter	Description	Example Value
hdfs.path	Specifies the HDFS data write directory. This parameter cannot be left blank.	hdfs://hacluster/flume/test
hdfs.filePrefix	Specifies the file name prefix after data is written to HDFS.	TMP_
hdfs.batchSize	Specifies the maximum number of events that can be written to HDFS once.	61200
hdfs.kerberosPrincipal	Specifies the Kerberos authentication user, which is mandatory in security versions. This configuration is required only in security clusters.	flume_hdfs
hdfs.kerberosKeytab	Specifies the keytab file path for Kerberos authentication, which is mandatory in security versions. This configuration is required only in security clusters.	/opt/test/conf/user.keytab NOTE Obtain the user.keytab file from the Kerberos certificate file of the user flume_hdfs . In addition, ensure that the user who installs and runs the Flume client has the read and write permissions on the user.keytab file.
hdfs.useLocalTimeStamp	Specifies whether to use the local time. Possible values are true and false .	true

4. Click **Export** to save the **properties.properties** configuration file to the local.


Step 3 Upload the configuration file.

Upload the file exported in [Step 2.4](#) to the **flume/conf** directory of the cluster. For details, see [Step 2.2](#).

Step 4 Verify log transmission.

1. Log in to FusionInsight Manager as a user who has the management permission on HDFS. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Services > HDFS**. On the page that is displayed, click the **NameNode(Node name,Active)** link next to **NameNode WebUI** to go to the HDFS web UI. On the displayed page, choose **Utilities > Browse the file system**.
2. Check whether the data is generated in the **/flume/test** directory on the HDFS.

Figure 7-14 Checking HDFS directories and files



----End

7.10.6 Typical Scenario: Collecting Logs from Kafka and Uploading Them to HDFS Through the Flume Client

Scenario

This section describes how to use the Flume client to collect logs from the topic list (test1) of the Kafka client and save them to the `/flume/test` directory on HDFS.

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

NOTE

By default, the cluster network environment is secure and the SSL authentication is not enabled during the data transmission process. For details about how to use the encryption mode, see [Configuring the Encrypted Transmission](#).

Prerequisites

- The Flume client has been installed.
- The cluster has been installed, including the HDFS, Kafka, and Flume services.
- You have created user `flume_hdfs` and authorized the HDFS directory and data to be operated during log verification.
- The network environment of the cluster is secure.


Procedure

Step 1 On FusionInsight Manager, choose **System > User** and choose **More > Download Authentication Credential** to download the Kerberos certificate file of user `flume_hdfs` and save it to the local host.

Step 2 Configure the client parameters of the Flume role.

1. Use the Flume configuration tool on FusionInsight Manager to configure the Flume role server parameters and generate a configuration file.
 - a. Log in to FusionInsight Manager and choose **Cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab.
 - b. Set **Agent Name** to `client`. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.
For example, use Kafka Source, File Channel, and HDFS Sink, as shown in [Figure 7-15](#).

Figure 7-15 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-34](#) based on the actual environment.

NOTE

- If you want to continue using the **properties.properties** file by modifying it, log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab, click **Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.

Table 7-34 Parameters to be modified for the Flume role client

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
kafka.topics	Specifies the subscribed Kafka topic list, in which topics are separated by commas (,). This parameter cannot be left blank.	test1
kafka.consumer.group.id	Specifies the data group ID obtained from Kafka. This parameter cannot be left blank.	flume

Parameter	Description	Example Value
kafka.bootstrap.servers	Specifies the bootstrap IP address and port list of Kafka. The default value is all Kafka lists in a Kafka cluster. If Kafka has been installed in the cluster and its configurations have been synchronized, this parameter can be left blank.	192.168.101.10:21007
batchSize	Specifies the number of events that Flume sends in a batch (number of data pieces).	61200
dataDirs	Specifies the directory for storing buffer data. The run directory is used by default. Configuring multiple directories on disks can improve transmission efficiency. Use commas (,) to separate multiple directories. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/data directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	/srv/BigData/hadoop/data1/flume/data

Parameter	Description	Example Value
checkpointDir	Specifies the directory for storing the checkpoint information, which is under the run directory by default. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/checkpoint directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	/srv/BigData/hadoop/data1/flume/checkpoint
transactionCapacity	Specifies the transaction size, that is, the number of events in a transaction that can be processed by the current Channel. The size cannot be smaller than the batchSize of Source. Setting the same size as batchSize is recommended.	61200
hdfs.path	Specifies the HDFS data write directory. This parameter cannot be left blank.	hdfs://hacluster/flume/test
hdfs.filePrefix	Specifies the file name prefix after data is written to HDFS.	TMP_
hdfs.batchSize	Specifies the maximum number of events that can be written to HDFS once.	61200

Parameter	Description	Example Value
hdfs.kerberosPrincipal	Specifies the Kerberos authentication user, which is mandatory in security versions. This configuration is required only in security clusters.	flume_hdfs
hdfs.kerberosKeytab	Specifies the keytab file path for Kerberos authentication, which is mandatory in security versions. This configuration is required only in security clusters.	/opt/test/conf/ user.keytab NOTE Obtain the user.keytab file from the Kerberos certificate file of the user flume_hdfs . In addition, ensure that the user who installs and runs the Flume client has the read and write permissions on the user.keytab file.
hdfs.useLocalTimeStamp	Specifies whether to use the local time. Possible values are true and false .	true

- d. Click **Export** to save the **properties.properties** configuration file to the local.
2. Upload the **properties.properties** file to **flume/conf/** in the Flume client installation directory.
3. To connect the Flume client to the HDFS, you need to add the following configuration:
 - a. Download the Kerberos certificate of account **flume_hdfs** and obtain the **krb5.conf** configuration file. Upload the configuration file to the **fusioninsight-flume-1.9.0/conf/** directory on the node where the client is installed.
 - b. In **fusioninsight-flume-1.9.0/conf/**, create the **jaas.conf** configuration file.

vi jaas.conf

```
KafkaClient {
com.sun.security.auth.module.Krb5LoginModule required
useKeyTab=true
keyTab="/opt/test/conf/user.keytab"
principal="flume_hdfs@<System domain name>"
useTicketCache=false
storeKey=true
debug=true;
};
```


Values of **keyTab** and **principal** vary depending on the actual situation.

- c. Obtain configuration files **core-site.xml** and **hdfs-site.xml** from **/opt/FusionInsight_Cluster_<Cluster ID>_Flume_ClientConfig/Flume/config** and upload them to **fusioninsight-flume-1.9.0/conf/**.
4. Go to **fusioninsight-flume-1.9.0/bin** in the installation directory of the client node and run the following command to restart the Flume process:
./flume-manage.sh restart

Step 3 Verify log transmission.

1. Log in to FusionInsight Manager as a user who has the management permission on HDFS. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Services > HDFS**. On the page that is displayed, click the **NameNode(Node name,Active)** link next to **NameNode WebUI** to go to the HDFS web UI. On the displayed page, choose **Utilities > Browse the file system**.
2. Check whether the data is generated in the **/flume/test** directory on the HDFS.

Figure 7-16 Checking HDFS directories and files



----End

7.10.7 Typical Scenario: Collecting Local Static Logs and Uploading Them to HBase

Scenario

This section describes how to use the Flume client to collect static logs from a local host and save them to the **flume_test** HBase table. In this scenario, multi-level agents are cascaded.

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

NOTE

By default, the cluster network environment is secure and the SSL authentication is not enabled during the data transmission process. For details about how to use the encryption mode, see [Configuring the Encrypted Transmission](#). The configuration applies to scenarios where only the server is configured, for example, SpoolDir Source+File Channel+HBase Sink.

Prerequisites

- The cluster has been installed, including the HBase and Flume services.
- The Flume client has been installed.
- The network environment of the cluster is secure.
- An HBase table has been created by running the **create 'flume_test', 'cf'** command.

- The MRS cluster administrator has understood service requirements and prepared HBase administrator **flume_hbase**.

Procedure

Step 1 On FusionInsight Manager, choose **System > User** and choose **More > Download Authentication Credential** to download the Kerberos certificate file of user **flume_hbase** and save it to the local host.

Step 2 Configure the client parameters of the Flume role.

1. Use the Flume configuration tool on FusionInsight Manager to configure the Flume role client parameters and generate a configuration file.
 - a. Log in to FusionInsight Manager and choose **Cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab.
 - b. Set **Agent Name** to **client**. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.
Use SpoolDir Source, File Channel, and Avro Sink.

Figure 7-17 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to **Table 7-35** based on the actual environment.

NOTE

- If you want to continue using the **properties.propertites** file by modifying it, log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab, click **Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.

Table 7-35 Parameters to be modified for the Flume role client

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test

Parameter	Description	Example Value
spoolDir	Specifies the directory where the file to be collected resides. This parameter cannot be left blank. The directory needs to exist and have the write, read, and execute permissions on the flume running user.	/srv/BigData/hadoop/data1/zb
trackerDir	Specifies the path for storing the metadata of files collected by Flume.	/srv/BigData/hadoop/data1/tracker
batchSize	Specifies the number of events that Flume sends in a batch (number of data pieces). A larger value indicates higher performance and lower timeliness.	61200
dataDirs	Specifies the directory for storing buffer data. The run directory is used by default. Configuring multiple directories on disks can improve transmission efficiency. Use commas (,) to separate multiple directories. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/data directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	/srv/BigData/hadoop/data1/flume/data

Parameter	Description	Example Value
checkpointDir	Specifies the directory for storing the checkpoint information, which is under the run directory by default. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/checkpoint directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	/srv/BigData/hadoop/data1/flume/checkpoint
transactionCapacity	Specifies the transaction size, that is, the number of events in a transaction that can be processed by the current Channel. The size cannot be smaller than the batchSize of Source. Setting the same size as batchSize is recommended.	61200
hostname	Specifies the name or IP address of the host whose data is to be sent. This parameter cannot be left blank. Name or IP address must be configured to be the name or IP address that the Avro source associated with it.	192.168.108.11

Parameter	Description	Example Value
port	Specifies the port that sends the data. This parameter cannot be left blank. It must be consistent with the port that is monitored by the connected Avro Source.	21154
ssl	Specifies whether to enable the SSL authentication. (You are advised to enable this function to ensure security.) Only Sources of the Avro type have this configuration item. <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the client authentication function is not enabled. 	false

- d. Click **Export** to save the **properties.properties** configuration file to the local.
2. Upload the **properties.properties** file to **flume/conf/** under the installation directory of the Flume client.

Step 3 Configure the server parameters of the Flume role and upload the configuration file to the cluster.

1. Use the Flume configuration tool on the FusionInsight Manager portal to configure the server parameters and generate the configuration file.
 - a. Log in to FusionInsight Manager and choose **Cluster > Services**. On the page that is displayed, choose **Flume**. On the displayed page, click the **Configuration Tool** tab.
 - b. Set **Agent Name** to **server**. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.
For example, use Avro Source, File Channel, and HBase Sink.

Figure 7-18 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by referring to [Table 7-36](#) based on the actual environment.

NOTE

- If the server parameters of the Flume role have been configured, you can choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume** > **Instance** on FusionInsight Manager. Then select the corresponding Flume role instance and click the **Download** button behind the **flume.config.file** parameter on the **Instance Configurations** page to obtain the existing server parameter configuration file. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume** > **Configuration Tool** > **Import**, import the file, and modify the configuration items related to non-encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.
- A unique checkpoint directory needs to be configured for each File Channel.

Table 7-36 Parameters to be modified for the Flume role server

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
bind	Specifies the IP address to which Avro Source is bound. This parameter cannot be left blank. It must be configured as the IP address that the server configuration file will upload.	192.168.108.11
port	Specifies the ID of the port that the Avro Source monitors. This parameter cannot be left blank. It must be configured as an unused port.	21154

Parameter	Description	Example Value
ssl	<p>Specifies whether to enable the SSL authentication. (You are advised to enable this function to ensure security.)</p> <p>Only Sources of the Avro type have this configuration item.</p> <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the client authentication function is not enabled. 	false
dataDirs	<p>Specifies the directory for storing buffer data. The run directory is used by default. Configuring multiple directories on disks can improve transmission efficiency. Use commas (,) to separate multiple directories. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/data directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.</p>	/srv/BigData/hadoop/data1/flumeserver/data
checkpointDir	<p>Specifies the directory for storing the checkpoint information, which is under the run directory by default. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/checkpoint directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.</p>	/srv/BigData/hadoop/data1/flumeserver/checkpoint
transactionCapacity	<p>Specifies the transaction size, that is, the number of events in a transaction that can be processed by the current Channel. The size cannot be smaller than the batchSize of Source. Setting the same size as batchSize is recommended.</p>	61200

Parameter	Description	Example Value
table	Specifies the HBase table name. This parameter cannot be left blank.	flume_test
columnFamily	Specifies the HBase column family name. This parameter cannot be left blank.	cf
batchSize	Specifies the maximum number of events written to HBase by Flume in a batch.	61200
kerberosPrincipal	Specifies the Kerberos authentication user, which is mandatory in security versions. This configuration is required only in security clusters.	flume_hbase
kerberosKeytab	Specifies the file path for Kerberos authentication, which is mandatory in security versions. This configuration is required only in security clusters.	/opt/test/conf/user.keytab NOTE Obtain the user.keytab file from the Kerberos certificate file of the user flume_hbase . In addition, ensure that the user who installs and runs the Flume client has the read and write permissions on the user.keytab file.

- d. Click **Export** to save the **properties.properties** configuration file to the local.
2. Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > Flume**. On the displayed page, click the **Flume** role on the **Instance** tab page.
3. Select the Flume role of the node where the configuration file is to be uploaded, choose **Instance Configurations > Import** beside the **flume.config.file**, and select the **properties.properties** file.

 **NOTE**

- An independent server configuration file can be uploaded to each Flume instance.
 - This step is required for updating the configuration file. Modifying the configuration file on the background is an improper operation because the modification will be overwritten after configuration synchronization.
4. Click **Save**, and then click **OK**.
 5. Click **Finish**.

Step 4 Verify log transmission.

1. Go to the directory where the HBase client is installed.
cd /Client installation directory/ HBase/hbase
kinit flume_hbase (Enter the password.)
2. Run the **hbase shell** command to access the HBase client.
3. Run the **scan 'flume_test'** statement. Logs are written in the HBase column family by line.

```
hbase(main):001:0> scan 'flume_test'
ROW                                COLUMN
+CELL

2017-09-18 16:05:36,394 INFO [hconnection-0x415a3f6a-shared--pool2-t1] ipc.AbstractRpcClient:
RPC Server Kerberos principal name for service=ClientService is hbase/hadoop.<system domain
name>@<system domain name>
default4021ff4a-9339-4151-a4d0-00f20807e76d          column=cf:pCol,
timestamp=1505721909388, value=Welcome to
flume
incRow                                           column=cf:iCol, timestamp=1505721909461, value=
\x00\x00\x00\x00\x00\x00\x00\x00\x01
2 row(s) in 0.3660 seconds
```

----End

7.11 Encrypted Transmission

7.11.1 Configuring the Encrypted Transmission

Scenario

This section describes how to configure the server and client parameters of the Flume service (including the Flume and MonitorServer roles) after the cluster is installed to ensure proper running of the service.

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

Prerequisites

The cluster and Flume service have been installed.

Procedure

- Step 1** Generate the certificate trust lists of the server and client of the Flume role respectively.

1. Remotely log in to the node using ECM where the Flume server is to be installed as user **omm**. Go to the **`\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin** directory.

```
cd `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin
```

 **NOTE**

The version 8.1.0.1 is used as an example. Replace it with the actual version number.

2. Run the following command to generate and export the server and client certificates of the Flume role:

```
sh geneJKS.sh -f xxx -g xxx
```

The generated certificate is saved in the ``${BIGDATA_HOME}`/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf` path .

- `flume_sChat.jks` is the certificate library of the Flume role server. `flume_sChat.crt` is the exported file of the `flume_sChat.jks` certificate. `-f` indicates the password of the certificate and certificate library.
- `flume_cChat.jks` is the certificate library of the Flume role client. `flume_cChat.crt` is the exported file of the `flume_cChat.jks` certificate. `-g` indicates the password of the certificate and certificate library.
- `flume_sChatt.jks` and `flume_cChatt.jks` are the SSL certificate trust lists of the Flume server and client, respectively.

 NOTE

All user-defined passwords involved in this section must meet the following strength requirements:

- The password must contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- The password must contain 8 to 64 characters.
- It is recommended that the user-defined passwords be changed periodically (for example, every three months), and certificates and trust lists be generated again to ensure security.

Step 2 Configure the server parameters of the Flume role and upload the configuration file to the cluster.

1. Remotely log in to any node where the Flume role is located as user `omm` using ECM. Run the following command to go to the ``${BIGDATA_HOME}`/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin` directory:

```
cd `${BIGDATA_HOME}`/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin
```

2. Run the following command to generate and obtain Flume server keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. It is the password of the `flume_sChat.jks` certificate library.


```
./genPwFile.sh
```

```
cat password.property
```

3. Use the Flume configuration tool on the FusionInsight Manager portal to configure the server parameters and generate the configuration file.
 - a. Log in to FusionInsight Manager. Choose **Services > Flume > Configuration Tool**.
 - b. Set **Agent Name** to `server`. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.

For example, use Avro Source, File Channel, and HDFS Sink, as shown in [Figure 7-19](#).

Figure 7-19 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by seeing [Table 7-37](#) based on the actual environment.

NOTE

- If the server parameters of the Flume role have been configured, you can choose **Services > Flume > Instance** on FusionInsight Manager. Then select the corresponding Flume role instance and click the **Download** button behind the **flume.config.file** parameter on the **Instance Configurations** page to obtain the existing server parameter configuration file. Choose **Services > Flume > Import** to change the relevant configuration items of encrypted transmission after the file is imported.
 - It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.
- d. Click **Export** to save the **properties.properties** configuration file to the local.

Table 7-37 Parameters to be modified of the Flume role server

Parameter	Description	Example Value
ssl	<p>Specifies whether to enable the SSL authentication. (You are advised to enable this function to ensure security.)</p> <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the client authentication function is not enabled. 	true
keystore	Indicates the server certificate.	<code>\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_sChat.jks</code>

Parameter	Description	Example Value
keystore-password	Specifies the password of the key library, which is the password required to obtain the keystore information. Enter the value of password obtained in Step 2.2 .	-
truststore	Indicates the SSL certificate trust list of the server.	\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_sChatt.jks
truststore-password	Specifies the trust list password, which is the password required to obtain the truststore information. Enter the value of password obtained in Step 2.2 .	-

- Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume**. On the displayed page, click the **Flume** role under **Role**.
- Select the Flume role of the node where the configuration file is to be uploaded, choose **Instance Configurations** > **Import** beside the **flume.config.file**, and select the **properties.properties** file.

 **NOTE**

- An independent server configuration file can be uploaded to each Flume instance.
 - This step is required for updating the configuration file. Modifying the configuration file on the background is an improper operation because the modification will be overwritten after configuration synchronization.
- Click **Save**, and then click **OK**. Click **Finish**.

Step 3 Set the client parameters of the Flume role.

- Run the following commands to copy the generated client certificate (**flume_cChat.jks**) and client trust list (**flume_cChatt.jks**) to the client directory, for example, **/opt/flume-client/fusionInsight-flume-1.9.0/conf/**. (The Flume client must have been installed.) **10.196.26.1** is the service plane IP address of the node where the client resides.

```
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_cChat.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
```

```
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_cChatt.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
```

 NOTE

When copying the client certificate, you need to enter the password of user **user** of the host (for example, **10.196.26.1**) where the client resides.

2. Log in to the node where the Flume client is decompressed as user **user**. Run the following command to go to the client directory **opt/flume-client/fusionInsight-flume-1.9.0/bin**.

cd opt/flume-client/fusionInsight-flume-1.9.0/bin

3. Run the following command to generate and obtain Flume client keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. The password is the same as the password of the certificate whose alias is *flumechatclient* and the password of the *flume_cChat.jks* certificate library.

./genPwFile.sh

cat password.property

 NOTE

If the following error message is displayed, run the export **JAVA_HOME=JDK path** command.

JAVA_HOME is null in current user,please install the JDK and set the JAVA_HOME

4. Run the **echo \$SCC_PROFILE_DIR** command to check whether the **SCC_PROFILE_DIR** environment variable is empty.
 - If yes, run the **source .sccfile** command.
 - If no, go to [Step 3.5](#).
5. Use the Flume configuration tool on FusionInsight Manager to configure the Flume role client parameters and generate a configuration file.
 - a. Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > Flume > Configuration Tool**.
 - b. Set **Agent Name** to **client**. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.

For example, use SpoolDir Source, File Channel, and Avro Sink, as shown in [Figure 7-20](#).

Figure 7-20 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by seeing [Table 7-38](#) based on the actual environment.

 NOTE

- If the client parameters of the Flume role have been configured, you can obtain the existing client parameter configuration file from *client installation directory/fusioninsight-flume-1.9.0/conf/properties.properties* to ensure that the configuration is in concordance with the previous. Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Flume > Configuration Tool > Import**, import the file, and modify the configuration items related to encrypted transmission.
 - It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.
 - A unique checkpoint directory needs to be configured for each File Channel.
- d. Click **Export** to save the **properties.properties** configuration file to the local.

Table 7-38 Parameters to be modified of the Flume role client

Parameter	Description	Example Value
ssl	Indicates whether to enable the SSL authentication. (You are advised to enable this function to ensure security.) <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the client authentication function is not enabled. 	true
keystore	Specified the client certificate.	/opt/flume-client/fusionInsight-flume-1.9.0/conf/flume_cChat.jks
keystore-password	Specifies the password of the key library, which is the password required to obtain the keystore information. Enter the value of password obtained in Step 3.3 .	-

Parameter	Description	Example Value
truststore	Indicates the SSL certificate trust list of the client.	/opt/flume-client/fusionInsight-flume-1.9.0/conf/flume_cChat.jks
truststore-password	Specifies the trust list password, which is the password required to obtain the truststore information. Enter the value of password obtained in Step 3.3 .	-

6. Upload the **properties.properties** file to **flume/conf/** under the installation directory of the Flume client.

Step 4 Generate the certificate and trust list of the server and client of the MonitorServer role respectively.

1. Log in to the host using ECM with the MonitorServer role assigned as user **omm**.

Go to the **`\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin** directory.

```
cd `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin
```

2. Run the following command to generate and export the server and client certificates of the MonitorServer role:

```
sh geneJKS.sh -m xxx -n xxx
```

The generated certificate is saved in the **`\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf** path. Where:

- **ms_sChat.jks** is the certificate library of the MonitorServer role server. **ms_sChat.crt** is the exported file of the **ms_sChat.jks** certificate. **-m** indicates the password of the certificate and certificate library.
- **ms_cChat.jks** is the certificate library of the MonitorServer role client. **ms_cChat.crt** is the exported file of the **ms_cChat.jks** certificate. **-n** indicates the password of the certificate and certificate library.
- **ms_sChatt.jks** and **ms_cChatt.jks** are the SSL certificate trust lists of the MonitorServer server and client, respectively.

Step 5 Set the server parameters of the MonitorServer role.

1. Run the following command to generate and obtain MonitorServer server keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. The password is the same as the password of the certificate whose alias is *mschatserver* and the password of the *ms_sChat.jks* certificate library.

./genPwFile.sh

cat password.property

- Run the following command to open the `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/service/application.properties` file: Modify related parameters based on the description in [Table 7-39](#), save the modification, and exit.

vi `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/service/application.properties`

Table 7-39 Parameters to be modified of the MonitorServer role server

Parameter	Description	Example Value
<code>ssl_need_kspas swd_decrypt_k ey</code>	Specifies whether to enable the user-defined key encryption and decryption function. (You are advised to enable this function to ensure security.) <ul style="list-style-type: none"> true indicates that the function is enabled. false indicates that the client authentication function is not enabled. 	true
<code>ssl_server_enab le</code>	Indicates whether to enable the SSL authentication. (You are advised to enable this function to ensure security.) <ul style="list-style-type: none"> true indicates that the function is enabled. false indicates that the client authentication function is not enabled. 	true
<code>ssl_server_key_ store</code>	Set this parameter based on the specific storage location.	<code>\${BIGDATA_HOME}/ FusionInsight_Porter_8.1.0.1 /install/FusionInsight- Flume-1.9.0/flume/conf/ ms_sChat.jks</code>
<code>ssl_server_trust _key_store</code>	Set this parameter based on the specific storage location.	<code>\${BIGDATA_HOME}/ FusionInsight_Porter_8.1.0.1 /install/FusionInsight- Flume-1.9.0/flume/conf/ ms_sChatt.jks</code>

Parameter	Description	Example Value
ssl_server_key_store_password	Indicates the client certificate password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the certificate). Enter the value of password obtained in Step 5.1 .	-
ssl_server_trust_key_store_password	Specifies the trustkeystore password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the trust list). Enter the value of password obtained in Step 5.1 .	-
ssl_need_client_auth	Indicates whether to enable the client authentication. (You are advised to enable this function to ensure security.) <ul style="list-style-type: none"> - true indicates that the function is enabled. - false indicates that the client authentication function is not enabled. 	true

- Restart the MonitorServer instance. Choose **Services > Flume > Instance > MonitorServer**, select the MonitorServer instance, and choose **More > Restart Instance**. Enter the cluster administrator password and click **OK**. After the restart is complete, click **Finish**.

Step 6 Set the client parameters of the MonitorServer role.

- Run the following commands to copy the generated client certificate (**ms_cChat.jks**) and client trust list (**ms_cChatt.jks**) to the **/opt/flume-client/fusionInsight-flume-1.9.0/conf/** client directory. **10.196.26.1** is the service plane IP address of the node where the client resides.

```
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/ms_cChat.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
```

```
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/ms_cChatt.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
```


2. Log in to the node where the Flume client is located as **user**. Run the following command to go to the client directory **/opt/flume-client/fusionInsight-flume-1.9.0/bin**.
cd /opt/flume-client/fusionInsight-flume-1.9.0/bin
3. Run the following command to generate and obtain MonitorServer client keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. The password is the same as the password of the certificate whose alias is *mschatclient* and the password of the *ms_cChat.jks* certificate library.
./genPwFile.sh
cat password.property
4. Run the following command to open the **/opt/flume-client/fusionInsight-flume-1.9.0/conf/service/application.properties** file. (**/opt/flume-client/fusionInsight-flume-1.9.0** is the directory where the client software is installed.) Modify related parameters based on the description in [Table 7-40](#), save the modification, and exit.
vi /opt/flume-client/fusionInsight-flume-1.9.0/flume/conf/service/application.properties

Table 7-40 Parameters to be modified of the MonitorServer role client

Parameter	Description	Example Value
ssl_need_kspas swd_decrypt_key	Indicates whether to enable the user-defined key encryption and decryption function. (You are advised to enable this function to ensure security.) – true indicates that the function is enabled. – false indicates that the client authentication function is not enabled.	true
ssl_client_enable	Indicates whether to enable the SSL authentication. (You are advised to enable this function to ensure security.) – true indicates that the function is enabled. – false indicates that the client authentication function is not enabled.	true

Parameter	Description	Example Value
ssl_client_key_store	Set this parameter based on the specific storage location.	\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/ms_cChat.jks
ssl_client_trust_key_store	Set this parameter based on the specific storage location.	\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/ms_cChatt.jks
ssl_client_key_store_password	Specifies the keystore password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the certificate). Enter the value of password obtained in Step 6.3 .	-
ssl_client_trust_key_store_password	Specifies the trustkeystore password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the trust list). Enter the value of password obtained in Step 6.3 .	-
ssl_need_client_auth	Indicates whether to enable the client authentication. (You are advised to enable this function to ensure security.) <ul style="list-style-type: none"> - true indicates that the function is enabled. - false indicates that the client authentication function is not enabled. 	true

----End

7.11.2 Typical Scenario: Collecting Local Static Logs and Uploading Them to HDFS

Scenario

This section describes how to use Flume to collect static logs from a local host and save them to the `/flume/test` directory on HDFS.

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

Prerequisites

- The cluster, HDFS and Flume services, and Flume client have been installed.
- User `flume_hdfs` has been created, and the HDFS directory and data used for log verification have been authorized to the user.

Procedure

Step 1 Generate the certificate trust lists of the server and client of the Flume role respectively.

1. Log in to the node where the Flume server is located as user `omm`. Go to the ``${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin` directory.

```
cd `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin
```

2. Run the following command to generate and export the server and client certificates of the Flume role:

```
sh geneJKS.sh -f Password -g Password
```

The generated certificate is saved in the ``${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf` path .

- `flume_sChat.jks` is the certificate library of the Flume role server. `flume_sChat.crt` is the exported file of the `flume_sChat.jks` certificate. `-f` indicates the password of the certificate and certificate library.
- `flume_cChat.jks` is the certificate library of the Flume role client. `flume_cChat.crt` is the exported file of the `flume_cChat.jks` certificate. `-g` indicates the password of the certificate and certificate library.
- `flume_sChatt.jks` and `flume_cChatt.jks` are the SSL certificate trust lists of the Flume server and client, respectively.

NOTE

All user-defined passwords involved in this section must meet the following requirements:

- Contain at least four types of the following: uppercase letters, lowercase letters, digits, and special characters.
- Contain at least eight characters and a maximum of 64 characters.
- It is recommended that the user-defined passwords be changed periodically (for example, every three months), and certificates and trust lists be generated again to ensure security.

Step 2 On FusionInsight Manager, choose **System > User** and choose **More > Download Authentication Credential** to download the Kerberos certificate file of user **flume_hdfs** and save it to the local host.

Step 3 Configure the server parameters of the Flume role and upload the configuration file to the cluster.

1. Log in to any node where the Flume role is located as user **omm**. Run the following command to go to the `/${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin` directory:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin
```


2. Run the following command to generate and obtain Flume server keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. It is the password of the **flume_sChat.jks** certificate library.

```
./genPwFile.sh
```

```
cat password.property
```

3. Use the Flume configuration tool on the FusionInsight Manager portal to configure the server parameters and generate the configuration file.
 - a. Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > Flume > Configuration Tool**.
 - b. Set **Agent Name** to **server**. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.
For example, use SpoolDir Source, File Channel, and HDFS Sink.

Figure 7-21 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by seeing [Table 7-41](#) based on the actual environment.

NOTE

- If the server parameters of the Flume role have been configured, you can choose **Cluster > Name of the desired cluster > Services > Flume > Instance** on FusionInsight Manager. Then select the corresponding Flume role instance and click the **Download** button behind the **flume.config.file** parameter on the **Instance Configurations** page to obtain the existing server parameter configuration file. Choose **Cluster > Name of the desired cluster > Services > Flume > Configuration Tool > Import**, import the file, and modify the configuration items related to encrypted transmission.
- It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.
- A unique checkpoint directory needs to be configured for each File Channel.

- d. Click **Export** to save the **properties.properties** configuration file to the local.

Table 7-41 Parameters to be modified of the Flume role server

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
bind	Specifies the IP address to which Avro Source is bound. This parameter cannot be left blank. It must be configured as the IP address that the server configuration file will upload.	192.168.108.11
port	Specifies the IP address to which Avro Source is bound. This parameter cannot be left blank. It must be configured as an unused port.	21154
ssl	Indicates whether to enable the SSL authentication. (You are advised to enable this function to ensure security.) Only Sources of the Avro type have this configuration item. <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the client authentication function is not enabled. 	true
keystore	Indicates the server certificate.	\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_sChat.jks
keystore-password	Specifies the password of the key library, which is the password required to obtain the keystore information. Enter the value of password obtained in Step 3.2 .	-

Parameter	Description	Example Value
truststore	Indicates the SSL certificate trust list of the server.	<code>\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_sChatt.jks</code>
truststore-password	Specifies the trust list password, which is the password required to obtain the truststore information. Enter the value of password obtained in Step 3.2 .	-
dataDirs	Specifies the directory for storing buffer data. The run directory is used by default. Configuring multiple directories on disks can improve transmission efficiency. Use commas (,) to separate multiple directories. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/data directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	<code>/srv/BigData/hadoop/data1/flumeserver/data</code>
checkpointDir	Specifies the directory for storing the checkpoint information, which is under the run directory by default. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/checkpoint directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	<code>/srv/BigData/hadoop/data1/flumeserver/checkpoint</code>

Parameter	Description	Example Value
transactionCapacity	Specifies the transaction size, that is, the number of events in a transaction that can be processed by the current Channel. The size cannot be smaller than the batchSize of Source. Setting the same size as batchSize is recommended.	61200
hdfs.path	Specifies the HDFS data write directory. This parameter cannot be left blank.	hdfs://hacluster/flume/test
hdfs.inUsePrefix	Specifies the prefix of the file that is being written to HDFS.	TMP_
hdfs.batchSize	Specifies the maximum number of events that can be written to HDFS once.	61200
hdfs.kerberosPrincipal	Specifies the Kerberos authentication user, which is mandatory in security versions. This configuration is required only in security clusters.	flume_hdfs
hdfs.kerberosKeytab	Specifies the keytab file path for Kerberos authentication, which is mandatory in security versions. This configuration is required only in security clusters.	/opt/test/conf/user.keytab NOTE Obtain the user.keytab file from the Kerberos certificate file of the user flume_hdfs . In addition, ensure that the user who installs and runs the Flume client has the read and write permissions on the user.keytab file.
hdfs.useLocalTimeStamp	Specifies whether to use the local time. Possible values are true and false .	true

- Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume**. On the displayed page, click the **Flume** role under **Role**.
- Select the Flume role of the node where the configuration file is to be uploaded, choose **Instance Configurations** > **Import** beside the **flume.config.file**, and select the **properties.properties** file.

 NOTE

- An independent server configuration file can be uploaded to each Flume instance.
 - This step is required for updating the configuration file. Modifying the configuration file on the background is an improper operation because the modification will be overwritten after configuration synchronization.
6. Click **Save**, and then click **OK**.
 7. Click **Finish**.

Step 4 Configure the client parameters of the Flume role.

1. Run the following commands to copy the generated client certificate (**flume_cChat.jks**) and client trust list (**flume_cChatt.jks**) to the client directory, for example, **/opt/flume-client/fusionInsight-flume-1.9.0/conf/**. (The Flume client must have been installed.) **10.196.26.1** is the service plane IP address of the node where the client resides.

```
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_cChat.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
```

```
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/flume_cChatt.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
```

 NOTE

When copying the client certificate, you need to enter the password of user **user** of the host (for example, **10.196.26.1**) where the client resides.

2. Log in to the node where the Flume client is decompressed as user **user**. Run the following command to go to the client directory **/opt/flume-client/fusionInsight-flume-1.9.0/bin**.

```
cd opt/flume-client/fusionInsight-flume-1.9.0/bin
```

3. Run the following command to generate and obtain Flume client keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. The password is the same as the password of the certificate whose alias is *flumechatclient* and the password of the *flume_cChat.jks* certificate library.

```
./genPwFile.sh
```

```
cat password.property
```

 NOTE


If the following error message is displayed, run the export **JAVA_HOME=JDKpath** command.

```
JAVA_HOME is null in current user,please install the JDK and set the JAVA_HOME
```

4. Run the **echo \$SCC_PROFILE_DIR** command to check whether the **SCC_PROFILE_DIR** environment variable is empty.
 - If yes, run the **source .sccfile** command.
 - If no, go to [Step 4.5](#).
5. Use the Flume configuration tool on FusionInsight Manager to configure the Flume role client parameters and generate a configuration file.
 - a. Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume** > **Configuration Tool**.

- b. Set **Agent Name** to **client**. Select the source, channel, and sink to be used, drag them to the GUI on the right, and connect them.
Use SpoolDir Source, File Channel, and HDFS Sink.

Figure 7-22 Example for the Flume configuration tool



- c. Double-click the source, channel, and sink. Set corresponding configuration parameters by seeing [Table 7-42](#) based on the actual environment.

NOTE

- If the client parameters of the Flume role have been configured, you can obtain the existing client parameter configuration file from *client installation directory/fusioninsight-flume-1.9.0/conf/properties.properties* to ensure that the configuration is in concordance with the previous. Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Flume > Configuration Tool > Import**, import the file, and modify the configuration items related to encrypted transmission.
 - It is recommended that the numbers of Sources, Channels, and Sinks do not exceed 40 during configuration file import. Otherwise, the response time may be very long.
- d. Click **Export** to save the **properties.properties** configuration file to the local.

Table 7-42 Parameters to be modified of the Flume role client

Parameter	Description	Example Value
Name	The value must be unique and cannot be left blank.	test
spoolDir	Specifies the directory where the file to be collected resides. This parameter cannot be left blank. The directory needs to exist and have the write, read, and execute permissions on the flume running user.	/srv/BigData/hadoop/data1/zb
trackerDir	Specifies the path for storing the metadata of files collected by Flume.	/srv/BigData/hadoop/data1/tracker

Parameter	Description	Example Value
batch-size	Specifies the number of events that Flume sends in a batch.	61200
dataDirs	Specifies the directory for storing buffer data. The run directory is used by default. Configuring multiple directories on disks can improve transmission efficiency. Use commas (,) to separate multiple directories. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/data directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	/srv/BigData/hadoop/data1/flume/data
checkpointDir	Specifies the directory for storing the checkpoint information, which is under the run directory by default. If the directory is inside the cluster, the /srv/BigData/hadoop/dataX/flume/checkpoint directory can be used. dataX ranges from data1 to dataN. If the directory is outside the cluster, it needs to be independently planned.	/srv/BigData/hadoop/data1/flume/checkpoint

Parameter	Description	Example Value
transactionCapacity	Specifies the transaction size, that is, the number of events in a transaction that can be processed by the current Channel. The size cannot be smaller than the batchSize of Source. Setting the same size as batchSize is recommended.	61200
hostname	Specifies the name or IP address of the host whose data is to be sent. This parameter cannot be left blank. Name or IP address must be configured to be the name or IP address that the Avro source associated with it.	192.168.108.11
port	Specifies the IP address to which Avro Sink is bound. This parameter cannot be left blank. It must be consistent with the port that is monitored by the connected Avro Source.	21154

Parameter	Description	Example Value
ssl	<p>Specifies whether to enable the SSL authentication. (You are advised to enable this function to ensure security.)</p> <p>Only Sources of the Avro type have this configuration item.</p> <ul style="list-style-type: none"> ▪ true indicates that the function is enabled. ▪ false indicates that the client authentication function is not enabled. 	true
keystore	Specifies the flume_cChat.jks certificate generated on the server.	/opt/flume-client/fusionInsight-flume-1.9.0/conf/flume_cChat.jks
keystore-password	<p>Specifies the password of the key library, which is the password required to obtain the keystore information.</p> <p>Enter the value of password obtained in Step 4.3.</p>	-
truststore	Indicates the SSL certificate trust list of the server.	/opt/flume-client/fusionInsight-flume-1.9.0/conf/flume_cChatt.jks
truststore-password	<p>Specifies the trust list password, which is the password required to obtain the truststore information.</p> <p>Enter the value of password obtained in Step 4.3.</p>	-

6. Upload the **properties.properties** file to **flume/conf/** under the installation directory of the Flume client.

Step 5 Generate the certificate and trust list of the server and client of the MonitorServer role respectively.

1. Log in to the host with the MonitorServer role assigned as user **omm**.

Go to the **\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin** directory.

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/bin
```

2. Run the following command to generate and export the server and client certificates of the MonitorServer role:

```
sh geneJKS.sh -m Password -n Password
```

The generated certificate is saved in the **\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf** path. Where:

- **ms_sChat.jks** is the certificate library of the MonitorServer role server. **ms_sChat.crt** is the exported file of the **ms_sChat.jks** certificate. **-m** indicates the password of the certificate and certificate library.
- **ms_cChat.jks** is the certificate library of the MonitorServer role client. **ms_cChat.crt** is the exported file of the **ms_cChat.jks** certificate. **-n** indicates the password of the certificate and certificate library.
- **ms_sChatt.jks** and **ms_cChatt.jks** are the SSL certificate trust lists of the MonitorServer server and client, respectively.

Step 6 Set the server parameters of the MonitorServer role.

1. Run the following command to generate and obtain MonitorServer server keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. The password is the same as the password of the certificate whose alias is *mschatserver* and the password of the *ms_sChat.jks* certificate library.

```
./genPwFile.sh
```

```
cat password.property
```

2. Run the following command to open the **\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/service/application.properties** file: Modify related parameters based on the description in [Table 7-43](#), save the modification, and exit.

```
vi ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/service/application.properties
```

Table 7-43 Parameters to be modified of the MonitorServer role server

Parameter	Description	Example Value
ssl_need_kspas swd_decrypt_k ey	Indicates whether to enable the user-defined key encryption and decryption function. (You are advised to enable this function to ensure security.) - true indicates that the function is enabled. - false indicates that the client authentication function is not enabled.	true
ssl_server_enab le	Indicates whether to enable the SSL authentication. (You are advised to enable this function to ensure security.) - true indicates that the function is enabled. - false indicates that the client authentication function is not enabled.	true
ssl_server_key_ store	Set this parameter based on the specific storage location.	\${BIGDATA_HOME}/ FusionInsight_Porter_8.1.0.1 /install/FusionInsight- Flume-1.9.0/flume/conf/ ms_sChat.jks
ssl_server_trust _key_store	Set this parameter based on the specific storage location.	\${BIGDATA_HOME}/ FusionInsight_Porter_8.1.0.1 /install/FusionInsight- Flume-1.9.0/flume/conf/ ms_sChatt.jks
ssl_server_key_ store_password	Indicates the client certificate password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the certificate). Enter the value of password obtained in Step 6.1 .	-

Parameter	Description	Example Value
ssl_server_trust_key_store_password	Indicates the client trust list password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the trust list). Enter the value of password obtained in Step 6.1 .	-
ssl_need_client_auth	Indicates whether to enable the client authentication. (You are advised to enable this function to ensure security.) <ul style="list-style-type: none"> - true indicates that the function is enabled. - false indicates that the client authentication function is not enabled. 	true

3. Restart the MonitorServer instance. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume** > **Instance** > **MonitorServer**, select the configured MonitorServer instance, and choose **More** > **Restart Instance**. Enter the cluster administrator password and click **OK**. After the restart is complete, click **Finish**.

Step 7 Set the client parameters of the MonitorServer role.

1. Run the following commands to copy the generated client certificate (**ms_cChat.jks**) and client trust list (**ms_cChatt.jks**) to the **/opt/flume-client/fusionInsight-flume-1.9.0/conf/** client directory. **10.196.26.1** is the service plane IP address of the node where the client resides.

```
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/ms_cChat.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
scp ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Flume-1.9.0/flume/conf/ms_cChatt.jks user@10.196.26.1:/opt/flume-client/fusionInsight-flume-1.9.0/conf/
```
2. Log in to the node where the Flume client is located as user **user**. Run the following command to go to the client directory **/opt/flume-client/fusionInsight-flume-1.9.0/bin**.

```
cd /opt/flume-client/fusionInsight-flume-1.9.0/bin
```
3. Run the following command to generate and obtain MonitorServer client keystore password, trust list password, and keystore-password encrypted private key information. Enter the password twice and confirm the password. The password is the same as the password of the certificate whose alias is *mschatclient* and the password of the *ms_cChat.jks* certificate library.

./genPwFile.sh

cat password.property

4. Run the following command to open the **/opt/flume-client/fusionInsight-flume-1.9.0/conf/service/application.properties** file. (**/opt/flume-client/fusionInsight-flume-1.9.0** is the directory where the client is installed.) Modify related parameters based on the description in [Table 7-44](#), save the modification, and exit.

vi /opt/flume-client/fusionInsight-flume-1.9.0/conf/service/application.properties

Table 7-44 Parameters to be modified of the MonitorServer role client


Parameter	Description	Example Value
ssl_need_kspas swd_decrypt_key	Indicates whether to enable the user-defined key encryption and decryption function. (You are advised to enable this function to ensure security.) – true indicates that the function is enabled. – false indicates that the client authentication function is not enabled.	true
ssl_client_enable	Indicates whether to enable the SSL authentication. (You are advised to enable this function to ensure security.) – true indicates that the function is enabled. – false indicates that the client authentication function is not enabled.	true
ssl_client_key_store	Set this parameter based on the specific storage location.	\${BIGDATA_HOME}/ FusionInsight_Porter_8.1.0.1 /install/FusionInsight- Flume-1.9.0/flume/conf/ ms_cChat.jks
ssl_client_trust_key_store	Set this parameter based on the specific storage location.	\${BIGDATA_HOME}/ FusionInsight_Porter_8.1.0.1 /install/FusionInsight- Flume-1.9.0/flume/conf/ ms_cChatt.jks

Parameter	Description	Example Value
ssl_client_key_store_password	Specifies the keystore password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the certificate). Enter the value of password obtained in Step 7.3 .	-
ssl_client_trust_key_store_password	Specifies the trustkeystore password. Set this parameter based on the actual situation of certificate creation (the plaintext key used to generate the trust list). Enter the value of password obtained in Step 7.3 .	-
ssl_need_client_auth	Indicates whether to enable the client authentication. (You are advised to enable this function to ensure security.) <ul style="list-style-type: none"> - true indicates that the function is enabled. - false indicates that the client authentication function is not enabled. 	true

Step 8 Verify log transmission.

1. Log in to FusionInsight Manager as a user who has the management permission on HDFS. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**, click the HDFS WebUI link to go to the HDFS WebUI, and choose **Utilities** > **Browse the file system**.
2. Check whether the data is generated in the **/flume/test** directory on the HDFS.

Figure 7-23 Checking HDFS directories and files



----End

7.12 Viewing Flume Client Monitoring Information

Scenario

The Flume client outside the FusionInsight cluster is a part of the end-to-end data collection. Both the Flume client outside the cluster and the Flume server in the cluster need to be monitored. Users can use FusionInsight Manager to monitor the Flume client and view the monitoring indicators of the Source, Sink, and Channel of the client as well as the client process status.

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

Procedure

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume** > **Flume Management** to view the current Flume client list and process status.
- Step 3** Click the **Instance ID**, and view client monitoring metrics in the **Current** area.
- Step 4** Click **History**. The page for querying historical monitoring data is displayed. Select a time range and click **View** to view the monitoring data within the time range.

----End

7.13 Connecting Flume to Kafka in Security Mode

Scenario

This section describes how to connect to Kafka using the Flume client in security mode.

This section applies to MRS 3.x or later.

Procedure

- Step 1** Create a **jaas.conf** file and save it to **$\${Flume\ client\ installation\ directory}\backslash conf$** . The content of the **jaas.conf** file is as follows:

```
KafkaClient {
  com.sun.security.auth.module.Krb5LoginModule required
  useKeyTab=true
  keyTab="/opt/test/conf/user.keytab"
```

```
principal="flume_hdfs@<System domain name>"  
useTicketCache=false  
storeKey=true  
debug=true;  
};
```

Set **keyTab** and **principal** based on site requirements. The configured **principal** must have certain kafka permissions.

- Step 2** Configure services. Set the port number of **kafka.bootstrap.servers** to **21007**, and set **kafka.security.protocol** to **SASL_PLAINTEXT**.
- Step 3** If the domain name of the cluster where Kafka is located is changed, change the value of *-Dkerberos.domain.name* in the **flume-env.sh** file in *`\${Flume client installation directory}/conf/`* based on the site requirements.
- Step 4** Upload the configured **properties.properties** file to *`\${Flume client installation directory}/conf/`*.

----End

7.14 Connecting Flume with Hive in Security Mode

Scenario

This section describes how to use Flume to connect to Hive (version 3.1.0) in the cluster.

This section applies to MRS 3.x or later.

Prerequisites

Flume and Hive have been correctly installed in the cluster. The services are running properly, and no alarm is reported.

Procedure

- Step 1** Import the following JAR packages to the lib directory (client/server) of the Flume instance to be tested as user **omm**:
 - antlr-2.7.7.jar
 - antlr-runtime-3.4.jar
 - calcite-core-1.16.0.jar
 - hadoop-mapreduce-client-core-3.1.1.jar
 - hive-beeline-3.1.0.jar
 - hive-cli-3.1.0.jar
 - hive-common-3.1.0.jar
 - hive-exec-3.1.0.jar
 - hive-hcatalog-core-3.1.0.jar
 - hive-hcatalog-pig-adapter-3.1.0.jar
 - hive-hcatalog-server-extensions-3.1.0.jar
 - hive-hcatalog-streaming-3.1.0.jar

- hive-metastore-3.1.0.jar
- hive-service-3.1.0.jar
- libfb303-0.9.3.jar
- hadoop-plugins-1.0.jar

You can obtain the JAR package from the Hive installation directory and restart the Flume process to ensure that the JAR package is loaded to the running environment.

Step 2 Set Hive configuration items.

On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive**. On the page that is displayed, click the **Configurations** tab then the **All Configurations** sub-tab. On this sub-tab page, choose **HiveServer(Role)** > **Customization** > **hive.server.customized.configs**.

Table 7-45 hive.server.customized.configs parameters

Name	Value
hive.support.concurrency	true
hive.exec.dynamic.partition.mode	nonstrict
hive.txn.manager	org.apache.hadoop.hive.ql.lockmgr.DbTxnManager
hive.compactor.initiator.on	true
hive.compactor.worker.threads	1

Step 3 Prepare the system user **flume_hive** who has the supergroup and Hive permissions, install the client, and create the required Hive table.

Example:

1. Install the cluster client in **/opt/client**.
2. Run the following command to authenticate the user:

```
cd /opt/client
source bigdata_env
kinit flume_hive
```
3. Run the **beeline** command and run the following table creation statement:

```
create table flume_multi_type_part(id string, msg string)
partitioned by (country string, year_month string, day string)
clustered by (id) into 5 buckets
stored as orc TBLPROPERTIES('transactional'='true');
```
4. Run the **select * from Table name;** command to query data in the table.
 In this case, the number of data records in the table is **0**.

Step 4 Prepare related configuration files. Assume that the client installation package is stored in **/opt/FusionInsight_Cluster_1_Services_ClientConfig**.

1. Obtain the following files from the **\$Client decompression directory/Hive/config** directory:

- hivesite-site.xml
 - hive-site.xml
2. Obtain the following files from the *\$Client decompression directory/HDFS/config* directory:
core-site.xml
 3. Create a directory on the host where the Flume instance is started and save the prepared files to the created directory.
Example: **/opt/hivesink-conf/hive-site.xml**.
 4. Copy all property configurations in the **hivesite-site.xml** file to the **hive-site.xml** file and ensure that the configurations are placed before the original configurations.
Data is loaded in sequence in Hive.

 **NOTE**

Ensure that the Flume running user **omm** has the read and write permissions on the directory where the configuration file is stored.

Step 5 Observe the result.

On the Hive client, run the **select * from Table name;** command. Check whether the corresponding data has been written to the Hive table.

----End

Examples

Flume configuration example (SpoolDir--Mem--Hive):

```
server.sources = spool_source
server.channels = mem_channel
server.sinks = Hive_Sink

#config the source
server.sources.spool_source.type = spooldir
server.sources.spool_source.spoolDir = /tmp/testflume
server.sources.spool_source.montime =
server.sources.spool_source.fileSuffix = .COMPLETED
server.sources.spool_source.deletePolicy = never
server.sources.spool_source.trackerDir = flumespool
server.sources.spool_source.ignorePattern = ^$
server.sources.spool_source.batchSize = 20
server.sources.spool_source.inputCharset =UTF-8
server.sources.spool_source.selector.type = replicating
server.sources.spool_source.fileHeader = false
server.sources.spool_source.fileHeaderKey = file
server.sources.spool_source.basenameHeaderKey= basename
server.sources.spool_source.deserializer = LINE
server.sources.spool_source.deserializer.maxBatchLine= 1
server.sources.spool_source.deserializer.maxLineLength= 2048
server.sources.spool_source.channels = mem_channel

#config the channel
server.channels.mem_channel.type = memory
server.channels.mem_channel.capacity =10000
server.channels.mem_channel.transactionCapacity= 2000
server.channels.mem_channel.channelfullcount= 10
server.channels.mem_channel.keep-alive = 3
server.channels.mem_channel.byteCapacity =
server.channels.mem_channel.byteCapacityBufferPercentage= 20

#config the sink
```

```
server.sinks.Hive_Sink.type = hive
server.sinks.Hive_Sink.channel = mem_channel
server.sinks.Hive_Sink.hive.metastore = thrift://{any MetaStore service IP address}:21088
server.sinks.Hive_Sink.hive.hiveSite = /opt/hivesink-conf/hive-site.xml
server.sinks.Hive_Sink.hive.coreSite = /opt/hivesink-conf/core-site.xml
server.sinks.Hive_Sink.hive.metastoreSite = /opt/hivesink-conf/hivemeatastore-site.xml
server.sinks.Hive_Sink.hive.database = default
server.sinks.Hive_Sink.hive.table = flume_multi_type_part
server.sinks.Hive_Sink.hive.partition = Tag,%Y-%m,%d
server.sinks.Hive_Sink.hive.txnsPerBatchAsk= 100
server.sinks.Hive_Sink.hive.autoCreatePartitions= true
server.sinks.Hive_Sink.useLocalTimeStamp = true
server.sinks.Hive_Sink.batchSize = 1000
server.sinks.Hive_Sink.hive.kerberosPrincipal= super1
server.sinks.Hive_Sink.hive.kerberosKeytab= /opt/mykeytab/user.keytab
server.sinks.Hive_Sink.round = true
server.sinks.Hive_Sink.roundValue = 10
server.sinks.Hive_Sink.roundUnit = minute
server.sinks.Hive_Sink.serializer = DELIMITED
server.sinks.Hive_Sink.serializer.delimiter= ";"
server.sinks.Hive_Sink.serializer.serdeSeparator= ','
server.sinks.Hive_Sink.serializer.fieldnames= id,msg
```

7.15 Configuring the Flume Service Model

7.15.1 Overview

This section applies to MRS 3.x or later.

Guide a reasonable Flume service configuration by providing performance differences between Flume common modules, to avoid a nonstandard overall service performance caused when a frontend Source and a backend Sink do not match in performance.

Only single channels are compared for description.

7.15.2 Service Model Configuration Guide

This section applies to MRS 3.x or later.

During Flume service configuration and module selection, the ultimate throughput of a sink must be greater than the maximum throughput of a source. Otherwise, in extreme load scenarios, the write speed of the source to a channel is greater than the read speed of sink from channel. Therefore, the channel is fully occupied due to frequent usage, and the performance is affected.

Avro Source and Avro Sink are usually used in pairs to transfer data between multiple Flume Agents. Therefore, Avro Source and Avro Sink do not become a performance bottleneck in general scenarios.

Inter-Module Performance

Based on comparison between the limit performances of modules, Kafka Sink and HDFS Sink can meet the throughput requirements when the front-end is SpoolDir Source. However, HBase Sink could become performance bottlenecks due to the low write performances thereof. As a result, data is stacked in Channel. If you have to use HBase Sink or other sinks that are prone to become performance bottlenecks, you can use **Channel Selector** or **Sink Group** to meet performance requirements.

Channel Selector

A channel selector allows a source to connect to multiple channels. Data of the source can be distributed or copied by selecting different types of selectors. Currently, a channel selector provided by Flume can be a replicating channel selector or a multiplexing channel selector.

Replicating: indicates that the data of the source is synchronized to all channels.

Multiplexing: indicates that based on the value of a specific field of the header of an event, a channel is selected to send the data. In this way, the data is distributed based on a service type.

- Replicating configuration example:

```
client.sources = kafkasource
client.channels = channel1 channel2
client.sources.kafkasource.type = org.apache.flume.source.kafka.KafkaSource
client.sources.kafkasource.kafka.topics = topic1,topic2
client.sources.kafkasource.kafka.consumer.group.id = flume
client.sources.kafkasource.kafka.bootstrap.servers = 10.69.112.108:21007
client.sources.kafkasource.kafka.security.protocol = SASL_PLAINTEXT
client.sources.kafkasource.batchDurationMillis = 1000
client.sources.kafkasource.batchSize = 800
client.sources.kafkasource.channels = channel1 c el2

client.sources.kafkasource.selector.type = replicating
client.sources.kafkasource.selector.optional = channel2
```

Table 7-46 Parameters in the Replicating configuration example

Parameter	Default Value	Description
Selector.type	replicating	Selector type. Set this parameter to replicating .
Selector.optional	-	Optional channel. Configure this parameter as a list.

- Multiplexing configuration example:

```
client.sources = kafkasource
client.channels = channel1 channel2
client.sources.kafkasource.type = org.apache.flume.source.kafka.KafkaSource
client.sources.kafkasource.kafka.topics = topic1,topic2
client.sources.kafkasource.kafka.consumer.group.id = flume
client.sources.kafkasource.kafka.bootstrap.servers = 10.69.112.108:21007
client.sources.kafkasource.kafka.security.protocol = SASL_PLAINTEXT
client.sources.kafkasource.batchDurationMillis = 1000
client.sources.kafkasource.batchSize = 800
client.sources.kafkasource.channels = channel1 channel2

client.sources.kafkasource.selector.type = multiplexing
client.sources.kafkasource.selector.header = myheader
client.sources.kafkasource.selector.mapping.topic1 = channel1
client.sources.kafkasource.selector.mapping.topic2 = channel2
client.sources.kafkasource.selector.default = channel1
```

Table 7-47 Parameters in the Multiplexing configuration example

Parameter	Default Value	Description
Selector.type	replicating	Selector type. Set this parameter to multiplexing .
Selector.header	Flume.selector.header	-
Selector.default	-	-
Selector.mapping.*	-	-

In a multiplexing selector example, select a field whose name is topic from the header of the event. When the value of the topic field in the header is topic1, send the event to a channel 1; or when the value of the topic field in the header is topic2, send the event to a channel 2.

Selectors need to use a specific header of an event in a source to select a channel, and need to select a proper header based on a service scenario to distribute data.

SinkGroup

When the performance of a backend single sink is insufficient, and high reliability or heterogeneous output is required, you can use a sink group to connect a specified channel to multiple sinks, thereby meeting use requirements. Currently, Flume provides two types of sink processors to manage sinks in a sink group. The types are load balancing and failover.

Failover: Indicates that there is only one active sink in the sink group each time, and the other sinks are on standby and inactive. When the active sink becomes faulty, one of the inactive sinks is selected based on priorities to take over services, so as to ensure that data is not lost. This is used in high-reliability scenarios.

Load balancing: Indicates that all sinks in the sink group are active. Each sink obtains data from the channel and processes the data. In addition, during running, loads of all sinks in the sink group are balanced. This is used in performance improvement scenarios.

- Load balancing configuration examples:

```
client.sources = source1
client.sinks = sink1 sink2
client.channels = channel1

client.sinkgroups = g1
client.sinkgroups.g1.sinks = sink1 sink2
client.sinkgroups.g1.processor.type = load_balance
client.sinkgroups.g1.processor.backoff = true
client.sinkgroups.g1.processor.selector = random

client.sinks.sink1.type = logger
client.sinks.sink1.channel = channel1

client.sinks.sink2.type = logger
client.sinks.sink2.channel = channel1
```


Table 7-48 Parameters of Load Balancing configuration examples

Parameter	Default Value	Description
sinks	-	Specifies the sink list of the sink group. Multiple sinks are separated by spaces.
processor.type	default	Specifies the type of a processor. Set this parameter to load_balance .
processor.backoff	false	Indicates whether to back off failed sinks exponentially.
processor.selector	round_robin	Specifies the selection mechanism. It must be round_robin, random, or a customized class that inherits AbstractSinkSelector.
processor.selector.maxTimeOut	30000	Specifies the time for masking a faulty sink. The default value is 30,000 ms.

- Failover configuration examples:

```

client.sources = source1
client.sinks = sink1 sink2
client.channels = channel1

client.sinkgroups = g1
client.sinkgroups.g1.sinks = sink1 sink2
client.sinkgroups.g1.processor.type = failover
client.sinkgroups.g1.processor.priority.sink1 = 10
client.sinkgroups.g1.processor.priority.sink2 = 5
client.sinkgroups.g1.processor.maxpenalty = 10000

client.sinks.sink1.type = logger
client.sinks.sink1.channel = channel1

client.sinks.sink2.type = logger
client.sinks.sink2.channel = channel1
    
```

Table 7-49 Parameters in the failover configuration example

Parameter	Default Value	Description
sinks	-	Specifies the sink list of the sink group. Multiple sinks are separated by spaces.

Parameter	Default Value	Description
processor.type	default	Specifies the type of a processor. Set this parameter to failover .
processor.priority.<sink Name>	-	Priority. <sinkName> must be defined in description of sinks. A sink having a higher priority is activated earlier. A larger value indicates a higher priority. Note: If there are multiple sinks, their priorities must be different. Otherwise, only one of them takes effect.
processor.maxpenalty	30000	Specifies the maximum backoff time of failed sinks (unit: ms).

Interceptors

The Flume interceptor supports modification or discarding of basic unit events during data transmission. You can specify the class name list of built-in interceptors in Flume or develop customized interceptors to modify or discard events. The following table lists the built-in interceptors in Flume. A complex example is used in this section. Other users can configure and use interceptions as required.

NOTE

1. The interceptor is used between the sources and channels of Flume. Most sources provide parameters for configuring interceptors. You can set the parameters as required.
2. Flume allows multiple interceptors to be configured for a source. The interceptor names are separated by spaces.
3. The specified interceptor sequence is the order in which they are called.
4. The contents inserted by the interceptor in the header can be read and used in sink.

Table 7-50 Types of built-in interceptors in Flume

Interceptor Type	Description
Timestamp Interceptor	The interceptor inserts a timestamp into the header of an event.
Host Interceptor	The interceptor inserts the IP address or host name of the node where the agent is located into the Header of an event.

Interceptor Type	Description
Remove Header Interceptor	The interceptor discards the corresponding event based on the strings that matches the regular expression contained in the event header.
UUID Interceptor	The interceptor generates a UUID string for the header of each event.
Search and Replace Interceptor	The interceptor provides a simple string-based search and replacement function based on Java regular expressions. The rule is the same as that of Java <code>Matcher.replaceAll()</code> .
Regex Filtering Interceptor	The interceptor uses the body of an event as a text file and matches the configured regular expression to filter events. The provided regular expression can be used to exclude or include events.
Regex Extractor Interceptor	The interceptor extracts content from the original events using a regular expression and adds the content to the header of events.

Regex Filtering Interceptor is used as an example to describe how to use the interceptor. (For other types of interceptions, see the configuration provided on the official website.)

Table 7-51 Parameter configuration for **Regex Filtering Interceptor**

Parameter	Default Value	Description
type	-	Specifies the component type name. The value must be regex_filter .
regex	-	Specifies the regular expression used to match events.
excludeEvents	false	By default, the matched events are collected. If this parameter is set to true , the matched events are deleted and the unmatched events are retained.

Configuration example (**netcat tcp** is used as the source, and **logger** is used as the sink). After configuring the preceding parameters, run the **telnet Host name or IP address 44444** command on the host where the Linux operating system is run, and enter a string that complies with the regular expression and another does

not comply with the regular expression. The log shows that only the matched string is transmitted.

```
#define the source, channel, sink
server.sources = r1

server.channels = c1
server.sinks = k1

#config the source
server.sources.r1.type = netcat
server.sources.r1.bind = ${Host IP address}
server.sources.r1.port = 44444
server.sources.r1.interceptors= i1
server.sources.r1.interceptors.i1.type= regex_filter
server.sources.r1.interceptors.i1.regex= (flume)|(myflume)
server.sources.r1.interceptors.i1.excludeEvents= false
server.sources.r1.channels = c1

#config the channel
server.channels.c1.type = memory
server.channels.c1.capacity = 1000
server.channels.c1.transactionCapacity = 100
#config the sink
server.sinks.k1.type = logger
server.sinks.k1.channel = c1
```

7.16 Introduction to Flume Logs

Log Description

Log path: The default path of Flume log files is `/var/log/Bigdata/Role name`.

- FlumeServer: `/var/log/Bigdata/flume/flume`
- FlumeClient: `/var/log/Bigdata/flume-client-n/flume`
- MonitorServer: `/var/log/Bigdata/flume/monitor`

Log archive rule: The automatic Flume log compression function is enabled. By default, when the size of logs exceeds 50 MB , logs are automatically compressed into a log file named in the following format: `<Original log file name>-<yyyy-mm-dd_hh-mm-ss>.[ID].log.zip`. A maximum of 20 latest compressed files are reserved. The number of compressed files can be configured on the Manager portal.

Table 7-52 Flume log list

Type	Name	Description
Run logs	<code>/flume/flumeServer.log</code>	Log file that records FlumeServer running environment information.
	<code>/flume/install.log</code>	FlumeServer installation log file
	<code>/flume/flumeServer-gc.log.<No.></code>	GC archive log file of the FlumeServer process
	<code>/flume/prestartDvietail.log</code>	Work log file before the FlumeServer startup

Type	Name	Description
	/flume/startDetail.log	Startup log file of the Flume process
	/flume/stopDetail.log	Shutdown log file of the Flume process
	/monitor/monitorServer.log	Log file that records MonitorServer running environment information
	/monitor/startDetail.log	Startup log file of the MonitorServer process
	/monitor/stopDetail.log	Shutdown log file of the MonitorServer process
	function.log	External function invoking log file
	/flume/flume- <i>Username-Date-pid-gc</i> .log	GC log file of the Flume process
	/flume/Flume-audit.log	Audit log file of the Flume client
	/flume/startAgent.out	Process parameter log file generated before Flume startup

Log Level

[Table 7-53](#) describes the log levels supported by Flume.

Levels of run logs are FATAL, ERROR, WARN, INFO, and DEBUG from the highest to the lowest priority. Run logs of equal or higher levels are recorded. The higher the specified log level, the fewer the logs recorded.

Table 7-53 Log level

Type	Level	Description
Run log	FATAL	Logs of this level record critical error information about system running.
	ERROR	Logs of this level record error information about system running.

Type	Level	Description
	WARN	Logs of this level record exception information about the current event processing.
	INFO	Logs of this level record normal running status information about the system and events.
	DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of Flume by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.

----End

 **NOTE**

The configurations take effect immediately without the need to restart the service.

Log Format

The following table lists the Flume log formats.

Table 7-54 Log format

Type	Format	Example
Run logs	<i><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></i>	2014-12-12 11:54:57,316 INFO [main] log4j dynamic load is start. org.apache.flume.tools.LogDyna micLoad.start(LogDynamicLoa d.java:59)

Type	Format	Example
	<code><yyyy-MM-dd HH:mm:ss,SSS><Username>< User IP><Time><Operation><Reso urce><Result><Detail></code>	2014-12-12 23:04:16,572 INFO [SinkRunner-PollingRunner- DefaultSinkProcessor] SRCIP=null OPERATION=close

7.17 Flume Client Cgroup Usage Guide

Scenario

This section describes how to join and log out of a cgroup, query the cgroup status, and change the cgroup CPU threshold.

This section applies to MRS 3.x or later.

Procedure

- **Join Cgroup**

Assume that the Flume client installation path is `/opt/FlumeClient`, and the cgroup CPU threshold is 50%. Run the following command to join a cgroup:

```
cd /opt/FlumeClient/fusioninsight-flume-1.9.0/bin
./flume-manage.sh cgroup join 50
```

 **NOTE**

- This command can be used to join a cgroup and change the cgroup CPU threshold.
- The value of the CPU threshold of a cgroup ranges from 1 to 100 x *N*. *N* indicates the number of CPU cores.

- **Check Cgroup status**

Assume that the Flume client installation path is `/opt/FlumeClient`. Run the following commands to query the cgroup status:

```
cd /opt/FlumeClient/fusioninsight-flume-1.9.0/bin
./flume-manage.sh cgroup status
```

- **Exit Cgroup**

Assume that the Flume client installation path is `/opt/FlumeClient`. Run the following commands to exit cgroup:

```
cd /opt/FlumeClient/fusioninsight-flume-1.9.0/bin
./flume-manage.sh cgroup exit
```

 NOTE

- After the client is installed, the default cgroup is automatically created. If the `-s` parameter is not configured during client installation, the default value `-1` is used. The default value indicates that the agent process is not restricted by the CPU usage.
- Joining or exiting a cgroup does not affect the agent process. Even if the agent process is not started, the joining or exiting operation can be performed successfully, and the operation will take effect after the next startup of the agent process.
- After the client is uninstalled, the cgroups created during the client installation are automatically deleted.

7.18 Secondary Development Guide for Flume Third-Party Plug-ins

Scenario

This section describes how to perform secondary development for third-party plug-ins.

This section applies to MRS 3.x or later.

Prerequisites

- You have obtained the third-party JAR package.
- You have installed Flume server or client.

Procedure

Step 1 Compress the self-developed code into a JAR package.

Step 2 Create a directory for the plug-in.

1. Access the `$FLUME_HOME/plugins.d` path and run the following command to create a directory:

```
mkdir thirdPlugin
cd thirdPlugin
mkdir lib libext native
```

The command output is displayed as follows:

```
[root@██████████ plugins.d]#mkdir thirdPlugin
[root@██████████ plugins.d]#ll
total 8
drwxr-x--- 3 root root 4096 ██████████ native
drwxr-xr-x 2 root root 4096 ██████████ thirdPlugin
[root@██████████ plugins.d]#cd thirdPlugin/
[root@██████████ thirdPlugin]#mkdir lib libext native
[root@██████████ thirdPlugin]#ll
total 12
drwxr-xr-x 2 root root 4096 ██████████ lib
drwxr-xr-x 2 root root 4096 ██████████ libext
drwxr-xr-x 2 root root 4096 ██████████ native
[root@██████████ thirdPlugin]#
```


2. Place the third-party JAR package in the `$FLUME_HOME/plugins.d/thirdPlugin/lib` directory. If the JAR package depends on other JAR packages, place the depended JAR packages to the `$FLUME_HOME/plugins.d/thirdPlugin/libext` directory, and place the local library files in `$FLUME_HOME/plugins.d/thirdPlugin/native`.

Step 3 Configure the `properties.properties` file in `$FLUME_HOME/conf/`.

For details about how to set parameters in the `properties.properties` file, see the parameter list in the `properties.properties` file in the corresponding typical scenario [Non-Encrypted Transmission](#) and [Encrypted Transmission](#).

 **NOTE**

- `$FLUME_HOME` indicates the Flume installation path. Set this parameter based on the site requirements (server or client) when configuring third-party plug-ins.
- `thirdPlugin` is the name of the third-party plugin.

----End

7.19 Common Issues About Flume

Flume logs are stored in `/var/log/Bigdata/flume/flume/flumeServer.log`. Most data transmission exceptions and data transmission failures are recorded in logs. You can run the following command:

tailf /var/log/Bigdata/flume/flume/flumeServer.log

- Problem: After the configuration file is uploaded, an exception occurs. After the configuration file is uploaded again, the scenario requirements are still not met, but no exception is recorded in the log.

Solution: Restart the Flume process, run the `kill -9 Process code` to kill the process code, and view the logs.

- Issue: "**java.lang.IllegalArgumentException: Keytab is not a readable file: /opt/test/conf/user.keytab**" is displayed when HDFS is connected.

Solution: Grant the read and write permissions to the Flume running user.

- Problem: The following error is reported when the Flume client is connected to Kafka:

Caused by: java.io.IOException: /opt/FlumeClient/fusioninsight-flume-1.9.0/cof//jaas.conf (No such file or directory)

Solution: Add the `jaas.conf` configuration file and save it to the `conf` directory of the Flume client.

vi jaas.conf

```
KafkaClient {
com.sun.security.auth.module.Krb5LoginModule required
useKeyTab=true
keyTab="/opt/test/conf/user.keytab"
principal="flume_hdfs@<System domain name>"
useTicketCache=false
storeKey=true
debug=true;
};
```

Values of `keyTab` and `principal` vary depending on the actual situation.

- Problem: The following error is reported when the Flume client is connected to HBase:

Caused by: java.io.IOException: /opt/FlumeClient/fusioninsight-flume-1.9.0/cof//jaas.conf (No such file or directory)

Solution: Add the **jaas.conf** configuration file and save it to the **conf** directory of the Flume client.

vi jaas.conf

```
Client {  
  com.sun.security.auth.module.Krb5LoginModule required  
  useKeyTab=true  
  keyTab="/opt/test/conf/user.keytab"  
  principal="flume_hbase@<System domain name>"  
  useTicketCache=false  
  storeKey=true  
  debug=true;  
};
```

Values of **keyTab** and **principal** vary depending on the actual situation.

- Question: After the configuration file is submitted, the Flume Agent occupies resources. How do I restore the Flume Agent to the state when the configuration file is not uploaded?

Solution: Submit an empty **properties.properties** file.

8 Using HBase

8.1 Using HBase from Scratch

HBase is a column-based distributed storage system that features high reliability, performance, and scalability. This section describes how to use HBase from scratch, including how to update the client on the Master node in the cluster, create a table using the client, insert data in the table, modify the table, read data from the table, delete table data, and delete the table.

Background

Suppose a user develops an application to manage users who use service A in an enterprise. The procedure of operating service A on the HBase client is as follows:

- Create the **user_info** table.
- Add users' educational backgrounds and titles to the table.
- Query user names and addresses by user ID.
- Query information by user name.
- Deregister users and delete user data from the user information table.
- Delete the user information table after service A ends.

Table 8-1 User information

ID	Name	Gender	Age	Address
12005000201	A	Male	19	City A
12005000202	B	Female	23	City B
12005000203	C	Male	26	City C
12005000204	D	Male	18	City D
12005000205	E	Female	21	City E
12005000206	F	Male	32	City F

ID	Name	Gender	Age	Address
12005000207	G	Female	29	City G
12005000208	H	Female	30	City H
12005000209	I	Male	26	City I
12005000210	J	Male	25	City J

Prerequisites

The client has been installed in a directory, for example, `/opt/client`. The client directory in the following operations is only an example. Change it to the actual installation directory. Before using the client, download and update the client configuration file, and ensure that the active management node of Manager is available.

Procedure

For MRS clusters earlier than version 3.x, perform the following steps:

Step 1 Download the client configuration file.

1. Log in to MRS Manager. For details, see [Accessing Manager](#). Then, choose **Services**.
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.

Step 2 Log in to the active management node of MRS Manager.

1. On the **Node** tab page, view the **Name** parameter. The node that contains **master1** in its name is the Master1 node. The node that contains **master2** in its name is the Master2 node.
The active and standby management nodes of MRS Manager are installed on Master nodes by default. Because Master1 and Master2 are switched over in active and standby mode, Master1 is not always the active management node of MRS Manager. Run a command in Master1 to check whether Master1 is active management node of MRS Manager. For details about the command, see [Step 2.4](#).
2. Log in to the Master1 node using the password as user **root**.
3. Run the following commands to switch to user **omm**:
sudo su - root
su - omm
4. Run the following command to check the active management node of MRS Manager:
sh \${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh
In the command output, the node whose **HActive** is **active** is the active management node, and the node whose **HActive** is **standby** is the standby

management node. In the following example, **mgtomsdat-sh-3-01-1** is the active management node, and **mgtomsdat-sh-3-01-2** is the standby management node.

```
Ha mode
double
NodeName      HostName      HAVersion      StartTime      HAActive
HAAllResOK    HARunPhase
192-168-0-30  mgtomsdat-sh-3-01-1  V100R001C01    2019-11-18 23:43:02
active      normal      Activated
192-168-0-24  mgtomsdat-sh-3-01-2  V100R001C01    2019-11-21 07:14:02
standby    normal      Deactivated
```

5. Log in to the active management node, for example, **192-168-0-30** of MRS Manager as user **root**, and run the following command to switch to user **omm**:

```
sudo su - omm
```

- Step 3** Run the following command to switch to the client installation directory, for example, **/opt/client**:

```
cd /opt/client
```

- Step 4** Run the following command to update the client configuration for the active management node.

```
sh refreshConfig.sh /opt/client Full path of the client configuration file package
```

For example, run the following command:

```
sh refreshConfig.sh /opt/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.
```

- Step 5** Use the client on a Master node.

1. On the active management node where the client is updated, for example, node **192-168-0-30**, run the following command to go to the client directory:

```
cd /opt/client
```

2. Run the following command to configure environment variables:

```
source bigdata_env
```

3. If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create HBase tables. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit MRS cluster user
```

For example, **kinit hbaseuser**.

4. Run the following HBase client command:

```
hbase shell
```

- Step 6** Run the following commands on the HBase client to implement service A.

1. Create the **user_info** user information table according to [Table 8-1](#) and add data to it.

```
create 'user_info',{NAME => 'i'}
```

For example, to add information about the user whose ID is 12005000201, run the following commands:

```
put 'user_info','12005000201','i:name','A'  
put 'user_info','12005000201','i:gender','Male'  
put 'user_info','12005000201','i:age','19'  
put 'user_info','12005000201','i:address','City A'
```

2. Add users' educational backgrounds and titles to the **user_info** table.

For example, to add educational background and title information about user 12005000201, run the following commands:

```
put 'user_info','12005000201','i:degree','master'  
put 'user_info','12005000201','i:pose','manager'
```

3. Query user names and addresses by user ID.

For example, to query the name and address of user 12005000201, run the following command:

```
scan 'user_info',  
{STARTROW=>'12005000201',STOPROW=>'12005000201',COLUMNS=>['i:name','i:address']}
```

4. Query information by user name.

For example, to query information about user A, run the following command:

```
scan 'user_info',{FILTER=>"SingleColumnValueFilter('i','name',=,'binary:A')"}  
}
```

5. Delete user data from the user information table.

All user data needs to be deleted. For example, to delete data of user 12005000201, run the following command:

```
delete 'user_info','12005000201','i'
```

6. Delete the user information table.

```
disable 'user_info'  
drop 'user_info'
```

----End

For MRS 3.x or later clusters, perform the following steps:

- Step 1** Use the client on the active management node.

1. Log in to the node where the client is installed as the client installation user and run the following command to switch to the client directory:

```
cd /opt/client
```

2. Run the following command to configure environment variables:

```
source bigdata_env
```

3. If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create HBase tables. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit MRS cluster user
```

For example, **kinit hbaseuser**.

4. Run the following HBase client command:

hbase shell

Step 2 Run the following commands on the HBase client to implement service A.

1. Create the **user_info** user information table according to [Table 8-1](#) and add data to it.

```
create 'user_info',{NAME => 'i'}
```

For example, to add information about the user whose ID is **12005000201**, run the following commands:

```
put 'user_info','12005000201','i:name','A'
```

```
put 'user_info','12005000201','i:gender','Male'
```

```
put 'user_info','12005000201','i:age','19'
```

```
put 'user_info','12005000201','i:address','City A'
```

2. Add users' educational backgrounds and titles to the **user_info** table.

For example, to add educational background and title information about user 12005000201, run the following commands:

```
put 'user_info','12005000201','i:degree','master'
```

```
put 'user_info','12005000201','i:pose','manager'
```

3. Query user names and addresses by user ID.

For example, to query the name and address of user 12005000201, run the following command:

```
scan'user_info',  
{STARTROW=>'12005000201',STOPROW=>'12005000201',COLUMNS=>['i:na  
me','i:address']}
```

4. Query information by user name.

For example, to query information about user A, run the following command:

```
scan'user_info',{FILTER=>"SingleColumnValueFilter('i','name',=,'binary:A')"
```

5. Delete user data from the user information table.

All user data needs to be deleted. For example, to delete data of user 12005000201, run the following command:

```
delete'user_info','12005000201','i'
```

6. Delete the user information table.

```
disable'user_info'
```

```
drop 'user_info'
```

----End

8.2 Using an HBase Client

Scenario

This section describes how to use the HBase client in an O&M scenario or a service scenario.

Prerequisites

- The client has been installed. For example, the installation directory is **/opt/hadoopclient**. The client directory in the following operations is only an example. Change it to the actual installation directory.
- Service component users have been created by the MRS cluster administrator. A machine-machine user needs to download the **keytab** file and a human-machine user needs to change the password upon the first login.
- If a non-**root** user uses the HBase client, ensure that the owner of the HBase client directory is this user. Otherwise, run the following command to change the owner.

```
chown user:group -R Client installation directory/HBase
```

Using the HBase Client (Versions Earlier Than MRS 3.x)

Step 1 Log in to the node where the client is installed as the client installation user.

Step 2 Run the following command to go to the client directory:

```
cd /opt/hadoopclient
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create HBase tables. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit Component service user
```

For example, **kinit hbaseuser**.

Step 5 Run the following HBase client command:

```
hbase shell
```

```
----End
```

Using the HBase Client (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 go to the client directory:

```
cd /opt/hadoopclient
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create HBase tables. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit Component service user
```


For example, **kinit hbaseuser**.

Step 5 Run the following HBase client command:

hbase shell

----End

Common HBase client commands

The following table lists common HBase client commands.

Table 8-2 HBase client commands

Command	Description
create	Used to create a table, for example, create 'test', 'f1', 'f2', 'f3' .
disable	Used to disable a specified table, for example, disable 'test' .
enable	Used to enable a specified table, for example, enable 'test' .
alter	Used to alter the table structure. You can run the alter command to add, modify, or delete column family information and table-related parameter values, for example, alter 'test', {NAME => 'f3', METHOD => 'delete'} .
describe	Used to obtain the table description, for example, describe 'test' .
drop	Used to delete a specified table, for example, drop 'test' . Before deleting a table, you must stop it.
put	Used to write the value of a specified cell, for example, put 'test','r1','f1:c1','myvalue1' . The cell location is unique and determined by the table, row, and column.
get	Used to get the value of a row or the value of a specified cell in a row, for example, get 'test','r1' .
scan	Used to query table data, for example, scan 'test' . The table name and scanner must be specified in the command.

8.3 Creating HBase Roles

Scenario

Create and configure an HBase role on Manager as an MRS cluster administrator. The HBase role can set HBase administrator permissions and read (R), write (W), create (C), execute (X), or manage (A) permissions for HBase tables and column families.

Users can create a table, query/delete/insert/update data, and authorize others to access HBase tables after they set the corresponding permissions for the specified databases or tables on HDFS.

 **NOTE**

- This section applies to MRS 3.x or later clusters.
- HBase roles can be created in security mode, but cannot be created in normal mode.
- If the current component uses Ranger for permission control, you need to configure related policies based on Ranger for permission management. For details, see [Adding a Ranger Access Permission Policy for HBase](#).

Prerequisites

- The MRS cluster administrator has understood service requirements.
- You have logged in to Manager.

Procedure

Step 1 On Manager, choose **System > Permission > Role**.

Step 2 On the displayed page, click **Create Role** and enter a **Role Name** and **Description**.

Step 3 Set **Permission**. For details, see [Table 8-3](#).

HBase permissions:

- HBase Scope: Authorizes HBase tables. The minimum permission is read (R) and write (W) for columns.
- SUPER_USER_GROUP: HBase administrator permissions.

 **NOTE**

Users have the read (R), write (W), create (C), execute (X), and administrate (A) permissions for the tables created by themselves.

Table 8-3 Setting a role

Task	Role Authorization
Setting the HBase administrator permission	In Configure Resource Permission , choose <i>Name of the desired cluster</i> > HBase and select HBase Administrator Permission .
Setting the permission for users to create tables	<ol style="list-style-type: none"> 1. In Configure Resource Permission, choose <i>Name of the desired cluster</i> > HBase > HBase Scope. 2. Click global. 3. In the Permission column of the specified namespace, select Create and Execute. For example, select Create and Execute for the default namespace default.

Task	Role Authorization
Setting the permission for users to write data to tables	<ol style="list-style-type: none"> In Configure Resource Permission, choose <i>Name of the desired cluster</i> > HBase > HBase Scope > global. In the Permission column of the specified namespace, select Write. For example, select Write for the default namespace default. By default, HBase sub-objects inherit the permission from the parent object.
Setting the permission for users to read data from tables	<ol style="list-style-type: none"> In Configure Resource Permission, choose <i>Name of the desired cluster</i> > HBase > HBase Scope > global. In the Permission column of the specified namespace, select Read. For example, select Read for the default namespace default. By default, HBase sub-objects inherit the permission from the parent object.
Setting the permission for users to manage namespaces or tables	<ol style="list-style-type: none"> In Configure Resource Permission, choose <i>Name of the desired cluster</i> > HBase > HBase Scope > global. In the Permission column of the specified namespace, select admin. For example, select admin for the default namespace default.

Task	Role Authorization
<p>Setting the permission for reading data from or writing data to columns</p>	<ol style="list-style-type: none"> 1. In Configure Resource Permission, select <i>Name of the desired cluster</i> > HBase > HBase Scope > global and click the specified namespace to display the tables in the namespace. 2. Click a table. 3. Click a column family. 4. Confirm whether you want to create a role? <ul style="list-style-type: none"> - If yes, enter the column name in the Resource Name text box. Use commas (,) to separate multiple columns. Select Read or Write. If there are no columns with the same name in the HBase table, a newly created column with the same name as the existing column has the same permission as the existing one. The column permission is set successfully. - If no, modify the column permission of the existing HBase role. The columns for which the permission has been separately set are displayed in the table. Go to Step 3.5. 5. To add column permissions for a role, enter the column name in the Resource Name text box and set the column permissions. To modify column permissions for a role, enter the column name in the Resource Name text box and set the column permissions. Alternatively, you can directly modify the column permissions in the table. If the column permissions are modified in the table and column permissions with the same name are added, the settings cannot be saved. You are advised to modify the column permission of a role directly in the table. The search function is supported.

Step 4 Click **OK**, and return to the **Role** page.

----End

8.4 Configuring HBase Replication

Scenario

As a key feature to ensure high availability of the HBase cluster system, HBase cluster replication provides HBase with remote data replication in real time. It provides basic O&M tools, including tools for maintaining and re-establishing active/standby relationships, verifying data, and querying data synchronization progress. To achieve real-time data replication, you can replicate data from the HBase cluster to another one.

Prerequisites

- The active and standby clusters have been successfully installed and started (the cluster status is **Running** on the **Active Clusters** page), and you have the administrator rights of the clusters.
- The network between the active and standby clusters is normal and ports can be used properly.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same Manager, configure mutual trust between clusters. If the active cluster is deployed in normal mode, mutual trust between clusters is not needed.
- If historical data exists in the active cluster and needs to be synchronized to the standby cluster, cross-cluster replication must be configured for the active and standby clusters. For details, see [Enabling Cross-Cluster Copy](#).
- Time is consistent between the active and standby clusters and the Network Time Protocol (NTP) service on the active and standby clusters uses the same time source.
- Mapping relationships between the names of all hosts in the active and standby clusters and service IP addresses have been configured in the `/etc/hosts` file by appending `192.***.***.*** host1` to the `hosts` file.
- The network bandwidth between the active and standby clusters is determined based on service volume, which cannot be less than the possible maximum service volume.

Constraints

- Despite that HBase cluster replication provides the real-time data replication function, the data synchronization progress is determined by several factors, such as the service loads in the active cluster and the health status of processes in the standby cluster. In normal cases, the standby cluster should not take over services. In extreme cases, system maintenance personnel and other decision makers determine whether the standby cluster takes over services according to the current data synchronization indicators.
- Currently, the replication function supports only one active cluster and one standby cluster in HBase.
- Typically, do not perform operations on data synchronization tables in the standby cluster, such as modifying table properties or deleting tables. If any misoperation on the standby cluster occurs, data synchronization between the active and standby clusters will fail and data of the corresponding table in the standby cluster will be lost.
- If the replication function of HBase tables in the active cluster is enabled for data synchronization, after modifying the structure of a table in the active cluster, you need to manually modify the structure of the corresponding table in the standby cluster to ensure table structure consistency.

Procedure

Enable the replication function for the active cluster to synchronize data written by Put.

Step 1 Log in to the MRS console, click a cluster name and choose **Components**.

Step 2 Go to the **All Configurations** page of the HBase service. For details, see [Modifying Cluster Service Configuration Parameters](#).

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 3 (Optional) Set configuration items listed in [Table 8-4](#). You can set the parameters based on the description or use the default values.

Table 8-4 Optional configuration items

Navigation Path	Parameter	Default Value	Description
HMaster > Performance	hbase.master.logcleaner.ttl	600000	Time to live (TTL) of HLog files. If the value is set to 604800000 (unit: millisecond), the retention period of HLog is 7 days.
	hbase.master.cleaner.interval	60000	Interval for the HMaster to delete historical HLog files. The HLog that exceeds the configured period will be automatically deleted. You are advised to set it to the maximum value to save more HLogs.
RegionServer > Replication	replication.source.size.capacity	16777216	Maximum size of edits, in bytes. If the edit size exceeds the value, HLog edits will be sent to the standby cluster.
	replication.source.nb.capacity	25000	Maximum number of edits, which is another condition for triggering HLog edits to be sent to the standby cluster. After data in the active cluster is synchronized to the standby cluster, the active cluster reads and sends data in HLog according to this parameter value. This parameter is used together with replication.source.size.capacity .
	replication.source.maxretriesmultiplier	10	Maximum number of retries when an exception occurs during replication.

Navigation Path	Parameter	Default Value	Description
	replication.source.sleepforretries	1000	Retry interval (unit: ms)
	hbase.regionserver.replication.handler.count	6	Number of replication RPC server instances on RegionServer

Enable the replication function for the active cluster to synchronize data written by bulkload.

Step 4 Determine whether to enable bulkload replication.

 **NOTE**

If bulkload import is used and data needs to be synchronized, you need to enable Bulkload replication.

If yes, go to [Step 5](#).

If no, go to [Step 9](#).

Step 5 Go to the **All Configurations** page of the HBase service parameters by referring to [Modifying Cluster Service Configuration Parameters](#).

Step 6 On the HBase configuration interface of the active and standby clusters, search for **hbase.replication.cluster.id** and modify it. It specifies the HBase ID of the active and standby clusters. For example, the HBase ID of the active cluster is set to **replication1** and the HBase ID of the standby cluster is set to **replication2** for connecting the active cluster to the standby cluster. To save data overhead, the parameter value length is not recommended to exceed 30.

Step 7 On the HBase configuration interface of the standby cluster, search for **hbase.replication.conf.dir** and modify it. It specifies the HBase configurations of the active cluster client used by the standby cluster and is used for data replication when the bulkload data replication function is enabled. The parameter value is a path name, for example, **/home**.

 **NOTE**

- In versions earlier than MRS 3.x, you do not need to set this parameter. Skip [Step 7](#).
- When bulkload replication is enabled, you need to manually place the HBase client configuration files (**core-site.xml**, **hdfs-site.xml**, and **hbase-site.xml**) in the active cluster on all RegionServer nodes in the standby cluster. The actual path for placing the configuration file is **\${hbase.replication.conf.dir}/\${hbase.replication.cluster.id}**. For example, if **hbase.replication.conf.dir** of the standby cluster is set to **/home** and **hbase.replication.cluster.id** of the active cluster is set to **replication1**, the actual path for placing the configuration files in the standby cluster is **/home/replication1**. You also need to change the corresponding directory and file permissions by running the **chown -R omm:wheel /home/replication1** command.
- You can obtain the client configuration file from a directory on the client in the active cluster, for example, **/opt/client/HBase/hbase/conf**.

Step 8 On the HBase configuration page of the active cluster, search for and change the value of **hbase.replication.bulkload.enabled** to **true** to enable bulkload replication.

Restarting the HBase service and install the client

Step 9 Save the configurations and restart HBase.

Step 10 In the active and standby clusters. Update the client configuration file.

Synchronize table data of the active cluster. (Skip this step if the active cluster has no data.)

Step 11 Access the HBase shell of the active cluster as user **hbase**.

1. On the active management node where the client has been updated, run the following command to go to the client directory:

```
cd /opt/client
```

2. Run the following command to configure environment variables:

```
source bigdata_env
```

3. If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit hbase
```

NOTE

The system prompts you to enter the password after you run **kinit hbase**. The default password of user **hbase** is **Hbase@123**.

4. Run the following HBase client command:

```
hbase shell
```

Step 12 Check whether historical data exists in the standby cluster. If historical data exists and data in the active and standby clusters must be consistent, delete data from the standby cluster first.

1. On the HBase shell of the standby cluster, run the **list** command to view the existing tables in the standby cluster.
2. Delete data tables from the standby cluster based on the output list.

```
disable 'tableName'
```

```
drop 'tableName'
```

Step 13 After HBase replication is configured and data synchronization is enabled, check whether tables and data exist in the active cluster and whether the historical data needs to be synchronized to the standby cluster.

Run the **list** command to check the existing tables in the active cluster and run the **scan 'tableName'** command to check whether the tables contain historical data.

- If tables exist and data needs to be synchronized, go to [Step 14](#).
- If no, no further action is required.

Step 14 The HBase replication configuration does not support automatic synchronization of historical data in tables. You need to back up the historical data of the active cluster and then manually synchronize the historical data to the standby cluster.

Manual synchronization refers to the synchronization of a single table that is implemented by Export, distcp, and Import.

The process for manually synchronizing data of a single table is as follows:

1. Export table data from the active cluster.

hbase org.apache.hadoop.hbase.mapreduce.Export - Dhbase.mapreduce.include.deleted.rows=true *Table name Directory where the source data is stored*

Example: **hbase org.apache.hadoop.hbase.mapreduce.Export - Dhbase.mapreduce.include.deleted.rows=true t1 /user/hbase/t1**

2. Copy the data that has been exported to the standby cluster.

hadoop distcp *Directory for storing source data in the active cluster* **hdfs://ActiveNameNodeIP:9820/** *Directory for storing source data in the standby cluster*

ActiveNameNodeIP indicates the IP address of the active NameNode in the standby cluster.

Example: **hadoop distcp /user/hbase/t1 hdfs://192.168.40.2:9820/user/hbase/t1**

3. Import data to the standby cluster as the HBase table user of the standby cluster.

hbase org.apache.hadoop.hbase.mapreduce.Import - Dimport.bulk.output=*Directory where the output data is stored in the standby cluster* *Table name Directory where the source data is stored in the standby cluster*

hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles *Directory where the output data is stored in the standby cluster* *Table name*

For example, **hbase org.apache.hadoop.hbase.mapreduce.Import - Dimport.bulk.output=/user/hbase/output_t1 t1 /user/hbase/t1** and

hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles /user/hbase/output_t1 t1

Add the replication relationship between the active and standby clusters.

- Step 15** Run the following command on the HBase Shell to create the replication synchronization relationship between the active cluster and the standby cluster:

```
add_peer 'Standby cluster ID', CLUSTER_KEY => 'ZooKeeper address of the standby cluster',{HDFS_CONFS => true}
```

- *Standby cluster ID* indicates an ID for the active cluster to recognize the standby cluster. It is recommended that the ID contain letters and digits.
- The ZooKeeper address of the standby cluster includes the service IP address of ZooKeeper, the port for listening to client connections, and the HBase root directory of the standby cluster on ZooKeeper.
- **{HDFS_CONFS => true}** indicates that the default HDFS configuration of the active cluster will be synchronized to the standby cluster. This parameter is used for HBase of the standby cluster to access HDFS of the active cluster. If bulkload replication is disabled, you do not need to use this parameter.

Suppose the standby cluster ID is replication2 and the ZooKeeper address of the standby cluster is **192.168.40.2,192.168.40.3,192.168.40.4:2181:/hbase**.

- Run the **add_peer 'replication2',CLUSTER_KEY => '192.168.40.2,192.168.40.3,192.168.40.4:2181:/hbase',CONFIG =>**

{ "hbase.regionserver.kerberos.principal" => "<val>",
"hbase.master.kerberos.principal" => "<val2>" } command for a security cluster and the `add_peer 'replication2',CLUSTER_KEY => '192.168.40.2,192.168.40.3,192.168.40.4:2181:/hbase'` command for a common cluster.

The `hbase.master.kerberos.principal` and `hbase.regionserver.kerberos.principal` parameters are the Kerberos users of HBase in the security cluster. You can search the `hbase-site.xml` file on the client for the parameter values. For example, if the client is installed in the `/opt/client` directory of the master node, you can run the `grep "kerberos.principal" /opt/client/HBase/hbase/conf/hbase-site.xml -A1` command to obtain the parameter value.

Figure 8-1 Obtaining the principal of HBase

```
[root@hadoop1000055 opt]# grep "kerberos.principal" /opt/client/HBase/hbase/conf/hbase-site.xml -A1
<name>hbase.regionserver.kerberos.principal</name>
<value>hbase/hadoop.hadoop.com@HADOOP.COM</value>
--
<name>hbase.master.kerberos.principal</name>
<value>hbase/hadoop.hadoop.com@HADOOP.COM</value>
--
```

 **NOTE**

1. Obtain the ZooKeeper service IP address.
Log in to the MRS console, click the cluster name, and choose **Components > ZooKeeper > Instances** to obtain the ZooKeeper service IP address.
2. On the ZooKeeper service parameter configuration page, search for `clientPort`, which is the port for the client to connect to the server.
3. Run the `list_peers` command to check whether the replication relationship between the active and standby clusters is added. If the following information is displayed, the relationship is successfully added.

```
hbase(main):003:0> list_peers
PEER_ID CLUSTER_KEY ENDPOINT_CLASSNAME STATE REPLICATE_ALL NAMESPACES
TABLE_CFS BANDWIDTH SERIAL
replication2 192.168.0.13,192.168.0.177,192.168.0.25:2181:/hbase ENABLED true 0 false
```

Specify the data writing status for the active and standby clusters.

Step 16 On the HBase shell of the active cluster, run the following command to retain the data writing status:

set_clusterState_active

The command is run successfully if the following information is displayed:

```
hbase(main):001:0> set_clusterState_active
=> true
```

Step 17 On the HBase shell of the standby cluster, run the following command to retain the data read-only status:

set_clusterState_standby

The command is run successfully if the following information is displayed:

```
hbase(main):001:0> set_clusterState_standby
=> true
```

Enable the HBase replication function to synchronize data.

Step 18 Check whether a namespace exists in the HBase service instance of the standby cluster and the namespace has the same name as the namespace of the HBase table for which the replication function is to be enabled.

On the HBase shell of the standby cluster, run the **list_namespace** command to query the namespace.

- If the same namespace exists, go to [Step 19](#).
- If the same namespace does not exist, on the HBase shell of the standby cluster, run the following command to create a namespace with the same name and go to [Step 19](#):

```
create_namespace'ns1
```

Step 19 On the HBase shell of the active cluster, run the following command to enable real-time replication for tables in the active cluster. This ensures that modified data in the active cluster can be synchronized to the standby cluster in real time.

You can only synchronize data of one HTable at one time.

```
enable_table_replication 'Table name'
```

 **NOTE**

- If the standby cluster does not contain a table with the same name as the table for which real-time synchronization is to be enabled, the table is automatically created.
- If a table with the same name as the table for which real-time synchronization is to be enabled exists in the standby cluster, the structures of the two tables must be the same.
- If the encryption algorithm SMS4 or AES is configured for '*Table name*', the function for synchronizing data from the active cluster to the standby cluster cannot be enabled for the HBase table.
- If the standby cluster is offline or has tables with the same name but different structures, the replication function cannot be enabled.

If the standby cluster is offline, start it.

If the standby cluster has a table with the same name but different structure, modify the table structure to make it as the same as the table structure of the active cluster. On the HBase shell of the standby cluster, run the **alter** command to change the password by referring to the example.

Step 20 On the HBase shell of the active cluster, run the following command to enable the real-time replication function for the active cluster to synchronize the HBase permission table:

```
enable_table_replication 'hbase:acl'
```

 **NOTE**

After the permission of the active HBase source data table is modified, to ensure that the standby cluster can properly read data, modify the role permission for the standby cluster.

Check the data synchronization status for the active and standby clusters.

Step 21 Run the following command on the HBase client to check the synchronized data of the active and standby clusters. After the replication function is enabled, you can run this command to check whether the newly synchronized data is consistent.

```
hbase org.apache.hadoop.hbase.mapreduce.replication.VerifyReplication --  
starttime=Start time --endtime=End time Column family name ID of the standby  
cluster Table name
```

 NOTE

- The start time must be earlier than the end time.
- The value of **starttime** and **endtime** must be in the timestamp format. You need to run **date -d "2015-09-30 00:00:00" +%s** to change a common time format to a timestamp format. The command output is a 10-digit number (accurate to second), but HBase identifies a 13-digit number (accurate to millisecond). Therefore, you need to add three zeros (000) to the end of the command output.

Switch over active and standby clusters.

 NOTE

1. If the standby cluster needs to be switched over to the active cluster, reconfigure the active/standby relationship by referring to [Step 10](#) and [Step 15](#) to [Step 20](#).
2. Do not perform [Step 11](#) to [Step 14](#).

----End

Related Commands

Table 8-5 HBase replication

Operation	Command	Description
Set up the active/standby relationship.	add_peer <i>'Standby cluster ID', 'Standby cluster address'</i> Examples: add_peer '1', 'zk1,zk2,zk3:2181:/hbase' add_peer '1', 'zk1,zk2,zk3:2181:/hbase1'	Set up the relationship between the active cluster and the standby cluster. To enable bulkload replication, run the add_peer <i>'Standby cluster ID', CLUSTER_KEY => 'Standby cluster address'</i> command, configure hbase.replication.conf.dir , and manually copy the HBase client configuration file in the active cluster to all RegionServer nodes in the standby cluster. For details, see Step 4 to 11 .
Remove the active/standby relationship.	remove_peer <i>'Standby cluster ID'</i> Example: remove_peer '1'	Remove standby cluster information from the active cluster.
Query the active/standby relationship.	list_peers	Query standby cluster information (mainly Zookeeper information) in the active cluster.

Operation	Command	Description
Enable the real-time user table synchronization function.	enable_table_replication <i>'Table name'</i> Example: enable_table_replication 't1'	Synchronize user tables from the active cluster to the standby cluster.
Disable the real-time user table synchronization function.	disable_table_replication <i>'Table name'</i> Example: disable_table_replication 't1'	Do not synchronize user tables from the active cluster to the standby cluster.
Verify data of the active and standby clusters.	bin/hbase org.apache.hadoop.hbase.mapreduce.replication.VerifyReplication --starttime --endtime Column family name Standby cluster ID Table name	Verify whether data of the specified table is the same between the active cluster and the standby cluster. The description of the parameters in this command is as follows: <ul style="list-style-type: none"> • Start time: If start time is not specified, the default value 0 will be used. • End time: If end time is not specified, the time when the current operation is submitted will be used by default. • Table name: If a table name is not entered, all user tables for which the real-time synchronization function is enabled will be verified by default.
Switch the data writing status.	set_clusterState_active set_clusterState_standby	Specifies whether data can be written to the cluster HBase tables.

Operation	Command	Description
Add or update the active cluster HDFS configurations saved in the peer cluster.	<code>set_replication_hdfs_confs 'PeerId', {'key1' => 'value1', 'key2' => 'value2'}</code>	<p>Enable replication for data including bulkload data. When HDFS parameters are modified in the active cluster, the modification cannot be automatically synchronized to the standby cluster. You need to manually run the command to synchronize the changes. The affected parameters are as follows:</p> <ul style="list-style-type: none"> • fs.defaultFS • dfs.client.failover.proxy.provider.hacluster • dfs.client.failover.connection.retries.on.timeouts • dfs.client.failover.connection.retries <p>For example, if the value of fs.defaultFS is changed to hdfs://hacluster_sale, run the <code>set_replication_hdfs_confs '1', {'fs.defaultFS' => 'hdfs://hacluster_sale'}</code> command to synchronize the HDFS configuration to the standby cluster whose ID is 1.</p>

8.5 Configuring HBase Parameters

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

If the default parameter settings of the MRS service cannot meet your requirements, you can modify the parameter settings as required.

Step 1 Go to the cluster details page and choose **Components**.

NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 2 Choose **HBase > Service Configuration** and switch **Basic** to **All**. On the displayed HBase configuration page, modify parameter settings.

Table 8-6 HBase parameters

Parameter	Description	Value
hbase.regionserver.hfile.durable.sync	Whether to enable the HFile durability to make data persistence on disks. If this parameter is set to true , HBase performance is affected because each HFile is synchronized to disks by hadoop fsync when being written to HBase. This parameter exists only in MRS 1.9.2 or earlier.	Possible values are as follows: <ul style="list-style-type: none"> • true • false The default value is true .
hbase.regionserver.wal.durable.sync	Specifies whether to enable WAL file durability to make the WAL data persistence on disks. If this parameter is set to true , HBase performance is affected because each edited WAL file is synchronized to disks by hadoop fsync when being written to HBase. This parameter exists only in MRS 1.9.2 or earlier.	Possible values are as follows: <ul style="list-style-type: none"> • true • false The default value is true .

----End

8.6 Enabling Cross-Cluster Copy

Scenario

DistCp is used to copy the data stored on HDFS from a cluster to another cluster. DistCp depends on the cross-cluster copy function, which is disabled by default. This function needs to be enabled in both clusters.

Modify parameters on MRS to enable cross-cluster copy.

Impact on the System

Yarn needs to be restarted to enable the cross-cluster copy function and cannot be accessed during the restart.

Prerequisites

The **hadoop.rpc.protection** parameter of the two HDFS clusters must be set to the same data transmission mode, which can be **privacy** (encryption enabled) or **authentication** (encryption disabled).

 NOTE

Go to the **All Configurations** page by referring to [Modifying Cluster Service Configuration Parameters](#) and search for **hadoop.rpc.protection**.

For versions earlier than MRS 3.x, choose **Components > HDFS > Service Configuration** on the cluster details page. Switch **Basic** to **All**, and search for **hadoop.rpc.protection**.

Procedure

- Step 1** Go to the **All Configurations** page of the Yarn service. For details, see [Modifying Cluster Service Configuration Parameters](#).

 NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- Step 2** In the navigation pane, choose **Yarn > Distcp**.

- Step 3** Set **haclusterX.remotenn1** of **dfs.namenode.rpc-address** to the service IP address and RPC port number 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. Enter a value in the *IP address:port* format.

 NOTE

You can log in to FusionInsight Manager in MRS 3.x clusters, and choose **Cluster > Name of the desired cluster > Services > HDFS > Instance** to obtain the service IP address of the NameNode instance.

For versions earlier than MRS 3.x, on the cluster details page, choose **Components > HDFS > Instances** to obtain the service IP address of the NameNode instance.

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 9820 and cannot be modified on MRS Manager.

For example, **10.1.1.1:9820** and **10.1.1.2:9820**.

- Step 4** Save the configuration. On the **Dashboard** tab page, and choose **More > Restart Service** to restart the Yarn service.

Operation succeeded is displayed. Click **Finish**. The Yarn service is started successfully.

- Step 5** Log in to the other cluster and repeat the preceding operations.

----End

8.7 Using the ReplicationSyncUp Tool

Prerequisites

1. Active and standby clusters have been installed and started.

2. Time is consistent between the active and standby clusters and the NTP service on the active and standby clusters uses the same time source.
3. When the HBase service of the active cluster is stopped, the ZooKeeper and HDFS services must be started and run.
4. ReplicationSyncUp must be run by the system user who starts the HBase process.
5. In security mode, ensure that the HBase system user of the standby cluster has the read permission on HDFS of the active cluster. This is because that it will update the ZooKeeper nodes and HDFS files of the HBase system.
6. When HBase of the active cluster is faulty, the ZooKeeper, file system, and network of the active cluster are still available.

Scenarios

The replication mechanism can use WAL to synchronize the state of a cluster with the state of another cluster. After HBase replication is enabled, if the active cluster is faulty, ReplicationSyncUp synchronizes incremental data from the active cluster to the standby cluster using the information from the ZooKeeper node. After data synchronization is complete, the standby cluster can be used as an active cluster.

Parameter Configuration

Parameter	Description	Default Value
hbase.replication.bulkload.enabled	Whether to enable the bulkload data replication function. The parameter value type is Boolean. To enable the bulkload data replication function, set this parameter to true for the active cluster.	false
hbase.replication.cluster.id	ID of the source HBase cluster. After the bulkload data replication is enabled, this parameter is mandatory and must be defined in the source cluster. The parameter value type is String.	-

Tool Usage

Run the following command on the client of the active cluster:

```
hbase org.apache.hadoop.hbase.replication.regionserver.ReplicationSyncUp -Dreplication.sleep.before.failover=1
```

NOTE

replication.sleep.before.failover indicates sleep time required for replication of the remaining data when RegionServer fails to start. You are advised to set this parameter to 1 second to quickly trigger replication.

Precautions

1. When the active cluster is stopped, this tool obtains the WAL processing progress and WAL processing queue from the ZooKeeper Node (RS znode) and copies the queues that are not copied to the standby cluster.
2. RegionServer of each active cluster has its own znode under the replication node of ZooKeeper in the standby cluster. It contains one znode of each peer cluster.
3. If RegionServer is faulty, each RegionServer in the active cluster receives a notification through the watcher and attempts to lock the znode of the faulty RegionServer, including its queues. The successfully created RegionServer transfers all queues to the znode of its own queue. After queues are transferred, they are deleted from the old location.
4. When the active cluster is stopped, ReplicationSyncUp synchronizes data between active and standby clusters using the information from the ZooKeeper node. In addition, WALs of the RegionServer znode will be moved to the standby cluster.

Restrictions and Limitations

If the standby cluster is stopped or the peer relationship is closed, the tool runs normally but the peer relationship cannot be replicated.

8.8 Using HIndex

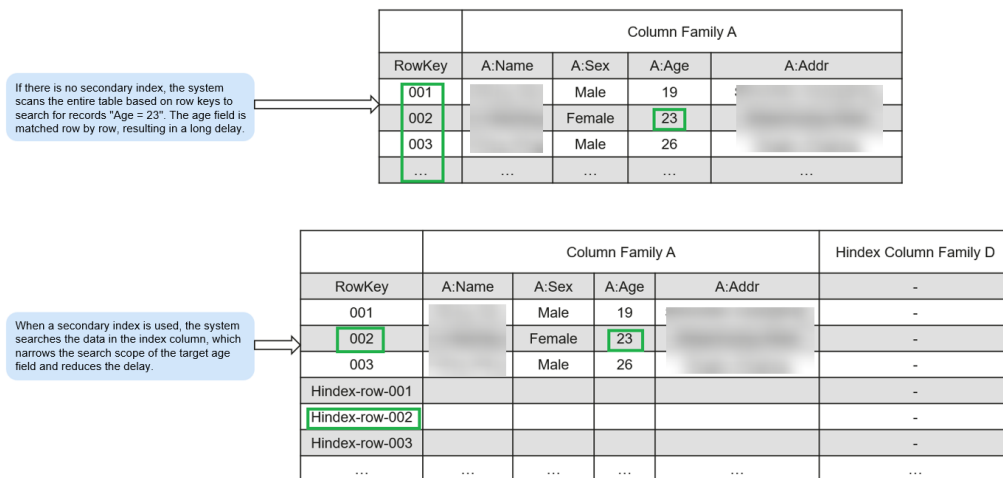
8.8.1 Introduction to HIndex

Scenarios

HBase is a distributed storage database of the Key-Value type. Data in tables is sorted by dictionary based on row keys. If you query data by specifying a row key or scan data in a specific row key range, HBase can help you quickly locate the data to be read. In most cases, you need to query data whose column value is *XXX*. HBase provides the filter function to enable you to query data with a specific column value. All data is scanned in the sequence of row keys and is matched with the specific column value until the required data is found. To obtain the required data, the filter will scan some unnecessary data. As a result, the filter function cannot meet the requirements for high-performance, frequent queries.

HBase HIndex is designed to address these issues. HBase HIndex provides HBase with the capability of indexing based on specific column values, making queries faster.

Figure 8-2 HBase HIndex

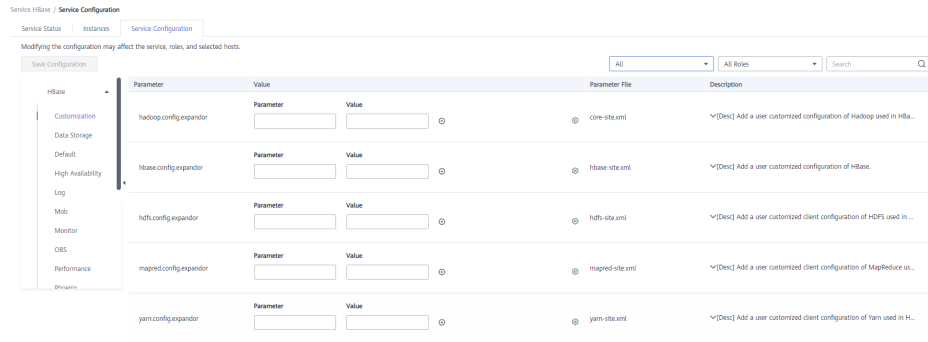


NOTE

- Rolling upgrade is not supported for index data.
- Composite index: You must add or delete all columns that participate in composite indexes. Otherwise, the data may be inconsistent.
- You should not explicitly configure any split policy to a data table where an index has been created.
- The mutation operations are not supported, such as increment and append.
- Index of the column with **maxVersions** greater than 1 is not supported.
- The value size of a column for which an index is added cannot exceed 32 KB.
- When the user data is deleted because TTL of the column family is invalid, the corresponding index data will not be deleted immediately. The index data will be deleted during major compaction.
- After an index is created, the TTL of the user column family must not be changed.
 - If the TTL of the column family is changed to a larger value after an index is created, delete the index and create one again. Otherwise, some generated index data may be deleted before the deletion of user data.
 - If the TTL of the column family is changed to a smaller value after an index is created, the index may be deleted after the deletion of user data.
- After disaster recovery is enabled for HBase tables, a secondary index is created in the active cluster and index table changes are not automatically synchronized to the standby cluster. To implement disaster recovery in this case, perform the following operations:
 1. After the secondary index is created in the active table, create a secondary index with the same schema and name using the same method in the standby cluster.
 2. In the active cluster, manually set **REPLICATION_SCOPE** of the index column family (default value: **d**) to 1.

Parameter Configuration

1. Log in to the MRS console, click a cluster name and choose **Components**.
2. Go to the **All Configurations** page of the HBase service. For details, see [Modifying Cluster Service Configuration Parameters](#).



3. View parameters on the HBase configurations page.

Navigation Path	Parameter	Default Value	Description
HMaster > System	hbase.coprocessor.master.classes	org.apache.hadoop.hbase.hindex.server.master.HIndexMasterCoprocesor,com.xxx.hadoop.hbase.backup.services.RecoveryCoprocesor,org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor,org.apache.hadoop.hbase.security.access.ReadOnlyClusterEnabler,org.apache.hadoop.hbase.rsgroup.RSGroupAdminEndpoint	This coprocessor is used to handle Master-level operations after the HIndex function is enabled, for example, creating an index meta table, adding an index, and deleting an index, a table, and index metadata.

Navigation Path	Parameter	Default Value	Description
RegionServer > RegionServer	hbase.coprocessor.regionserver.classes	org.apache.hadoop.hbase.hindex.server.regionserver.HIndexRegionServerCoprocessor,org.apache.hadoop.hbase.JMXListener,org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor	This coprocessor is used to handle the operations that the Master delivers to RegionServer after the HIndex function is enabled.

Navigation Path	Parameter	Default Value	Description
	hbase.coprocessor.region.classes	org.apache.hadoop.hbase.hindex.server.regionserver.HIndexRegionCoprocessor,org.apache.hadoop.hbase.security.token.TokenProvider,com.xxx.hadoop.hbase.backup.services.RecoveryCoprocessor,org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor,org.apache.hadoop.hbase.security.access.SecureBulkLoadEndpoint,org.apache.hadoop.hbase.security.access.ReadOnlyClusterEnabler,org.apache.hadoop.hbase.coprocessor.MetaTableMetrics	This coprocessor is used to operate data in the Region after the HIndex function is enabled.

Navigation Path	Parameter	Default Value	Description
	hbase.coprocessor.wal.classes	org.apache.hadoop.hbase.hindex.server.regionserver.HIndexRegionServerCoprocessor,org.apache.hadoop.hbase.JMXListener,org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor	<p>This coprocessor is used for Replication, which filters index data to prevent the index data from being sent to the peer cluster. The peer cluster generates index data by itself.</p> <p>This parameter is supported only in versions earlier than MRS 3.x.</p>

 NOTE

1. The preceding default values need to be configured after the HBase HIndex function is enabled. In MRS clusters that support the HBase HIndex function, the values have been configured by default.
2. Ensure that the **master** parameter is configured on HMaster and the **region** and **regionserver** parameters are configured on RegionServer.

Related Interfaces

The APIs that use HIndex are in the **org.apache.hadoop.hbase.hindex.client.HIndexAdmin** class. The following table describes the related APIs.

Operation	API	Description	Precautions
Add an index.	addIndices()	Add an index to a table without data. Calling this API will add the specified index to a table but skips index data generation. Therefore, after this operation, the index cannot be used for the scanning and filtering operations. This API applies to scenarios where users want to add indexes in batches to tables that have a large amount of pre-existing user data. The specific operation is to use external tools such as the TableIndexer tool to build index data.	<ul style="list-style-type: none"> • An index cannot be modified once it is added. To modify the index, you need to delete the old index and then create a new one. • Do not create two indexes on the same column with different index names. Otherwise, storage and processing resources will be wasted. • Indexes cannot be added to a system table. • The append and increment operations are not supported when data is put into the index column. • If any fault occurs on the client except DoNotRetryIOException, you need to try again. • An index column family is selected from the following conditions in sequence based on availability: <ul style="list-style-type: none"> – Typically, the default index column family is d. However, if the value of hindex.default.family.name is set, the value will be used. – Symbol #, @, \$, or %
	addIndicesWithData()	Add an index to a table with data. This API is used to add the specified index to the table and create index data for the existing user data. Alternatively, the API can be called to generate an index and then generate index data when the user data is being stored. Therefore, after this operation, the index can be used for the scanning and filtering operations immediately.	

Operation	API	Description	Precautions
			<ul style="list-style-type: none"> - #0, @ 0, \$ 0, %0, #1, @ 1 ...to #255, @ 255, \$ 255, %255 - Throw exceptions. • You can use the HIndex TableIndexer tool to add indexes without building index data.
Delete an index.	dropIndices()	<p>This API is used to delete an index only. It deletes the specified index from a table but skips the corresponding index data. After this operation, the index cannot be used for the scanning and filtering operations. The cluster automatically deletes old index data during major compaction.</p> <p>This API applies to scenarios where a table contains a large amount of index data and dropIndicesWithData() is unavailable. In addition, you can use the TableIndexer tool to delete indexes and index data.</p>	<ul style="list-style-type: none"> • An index can be disabled when it is in the ACTIVE, INACTIVE, or DROPPING state. • If you use dropIndices() to delete an index, ensure that the index data has been deleted before the index is added to the table with the same index name (that is, major compaction has been completed). • If you delete an index, the following information will also be deleted: <ul style="list-style-type: none"> - A column family with an index - Any one of column families in a combination index • Indexes and index data can be deleted together using the HIndex TableIndexer tool.
	dropIndicesWithData()	<p>Delete index data. This API deletes the specified index and all index data corresponding to the index in a user table. After this operation, the index is completely deleted from the table and is no longer used for the scanning and filtering operations.</p>	

Operation	API	Description	Precautions
Enable/ Disable an index.	disableIndices()	This API disables all indexes specified by a user so that they are no longer used for the scanning and filtering operations.	<ul style="list-style-type: none"> • An index can be enabled when the index is in the ACTIVE, INACTIVE, or BUILDING state. • An index can be disabled when the index is in the ACTIVE or INACTIVE state. • Before disabling an index, ensure that the index data is consistent with the user data. If no new data is added to the table when the index is disabled, the index data is consistent with the user data. • When enabling an index, you can use the TableIndexer tool to build index data to ensure data consistency.
	enableIndices()	This API enables all indexes specified by a user so that they can be used for the scanning and filtering operations.	
View the created index.	listIndices()	This API is used to list all indexes of a specified table.	N/A

Querying Data Based on Indexes

You can use a filter to query data in a user table with an index. The query result of a user table with a single or combination index is the same as that of a table without an index, but the table with an index provides higher data query performance than the table without an index.

The index usage rules are as follows:

- Scenario 1: A single index is created for one or more columns.
 - When this column is used for AND or OR query filtering, an index can improve query performance.
Example: Filter_Condition(IndexCol1)AND / OR Filter_Condition(IndexCol2)
 - When you use **Index Column AND Non-Index Column** for filtering in the query, the index can improve query performance.

- Example: `Filter_Condition(IndexCol1)AND
Filter_Condition(IndexCol2)AND Filter_Condition(NonIndexCol1)`
- When you use **Index Column OR Non-Index Column** for filtering in the query but do not use an index, query performance will not be improved.
Example: `Filter_Condition(IndexCol1)AND / OR
Filter_Condition(IndexCol2) OR Filter_Condition(NonIndexCol1)`
 - Scenario 2: A combination index is created for multiple columns.
 - When the columns to be queried are all or part of the combination index and have the same order as the combination index, using the index improves query performance.
For example, create a combination index for C1, C2, and C3.
 - The index takes effect in the following situations:
`Filter_Condition(IndexCol1)AND Filter_Condition(IndexCol2)AND
Filter_Condition(IndexCol3)`
`Filter_Condition(IndexCol1)AND Filter_Condition(IndexCol2)`
`FILTER_CONDITION(IndexCol1)`
 - The index does not take effect in the following situations:
`Filter_Condition(IndexCol2)AND Filter_Condition(IndexCol3)`
`Filter_Condition(IndexCol1)AND Filter_Condition(IndexCol3)`
`FILTER_CONDITION(IndexCol2)`
`FILTER_CONDITION(IndexCol3)`
 - When you use **Index Column AND Non-Index Column** for filtering in the query, the index can improve query performance.
Examples:
`Filter_Condition(IndexCol1)AND Filter_Condition(NonIndexCol1)`
`Filter_Condition(IndexCol1)AND Filter_Condition(IndexCol2)AND
Filter_Condition(NonIndexCol1)`
 - When you use **Index Column OR Non-Index Column** for filtering in the query but do not use an index, query performance will not be improved.
Examples:
`Filter_Condition(IndexCol1)OR Filter_Condition(NonIndexCol1)`
`(Filter_Condition(IndexCol1)AND
Filter_Condition(IndexCol2))OR(Filter_Condition(NonIndexCol1))`
 - When multiple columns are used for query, you can specify a value range for only the last column in the combination index and set other columns to specified values
For example, create a combination index for C1, C2, and C3. In a range query, only the value range of C3 can be set. The filter criteria are "C1 = XXX, C2 = XXX, and C3 = Value range."

Query Policy Selection

Use **SingleColumnValueFilter** or **SingleColumnRangeFilter**. It will provide the definite value **column_family:qualifierpair** (called **col1**) in filter criteria.

If **col1** is the first index column in the table, any index in the table can be a candidate index used during the query. The following provides an example:

If there is an index on **col1**, the index can be used as a candidate index because **col1** is the first and the only column of the index. If there is another index on **col1** and **col2**, you can consider this index as a candidate index because **col1** is the first column in the index list. However, if there is an index on **col2** and **col1**, this index cannot be used as a candidate index because the first column in the index list is not **col1**.

The most suitable method to use the index now is that when there are multiple candidate indexes, select the most suitable index for scanning data.

You can use the following solutions to learn how to select the best index policy.

- It is better to fully match.
Scenario: There are two indexes available, one for **col1&col2** and the other for **col1**.
In this scenario, the second index is better than the first one, because it scans less index data.
- If there are multiple candidate multi-column indexes, select an index with fewer index columns.
Scenario: There are two indexes available, one for **col1&col2** and the other for **col1&col2&col3**.
In this case, you had better use the index on **col1&col2**, because it scans less index data.

NOTE

- During a query based on an index, the index state must be **ACTIVE**. You can call the **listIndices()** API to view the index state.
- To query the correct data based on the index, ensure the consistency between index data and user data.
- Run the following command to perform a complex query on the HBase shell client (assuming that an index has been created for the specified column):
scan 'tablename', {FILTER => "SingleColumnValueFilter(family, qualifier, compareOp, comparator, filterIfMissing, latestVersionOnly)"}
Example: **scan 'test', {FILTER => "SingleColumnValueFilter('info', 'age', =, 'binary:26', true, true)"}**
In the preceding scenario, if you want to save the row where no column is found in the result, you should not create any index in any such column, because if the column to be queried does not exist, the row will be filtered out when SCVF is used to scan the index columns. When the SCVF whose **filterIfMissing** is **false** (default value) scans non-index columns, rows where no column is queried will also be returned in the result. Therefore, to avoid inconsistent query results, you are advised to set **filterIfMissing** to **true** after creating SCVF for the index column.
- Run the following command on the HBase shell client to view the index data created for user data:
scan 'tablename', {ATTRIBUTES => {'FETCH_INDEX_DATA' => 'true'}}

8.8.2 Loading Index Data in Batches

Scenarios

HBase provides the ImportTsv&LoadIncremental tool to load user data in batches. HBase also provides the HIndexImportTsv tool to load both the user data and index data in batches. HIndexImportTsv inherits all functions of the HBase batch data loading tool ImportTsv. If a table is not created before the HIndexImportTsv tool is executed, an index will be created when the table is created, and index data is generated when user data is generated.

Procedure

1. Run the following commands to import data to HDFS:

```
hdfs dfs -mkdir <inputdir>
```

```
hdfs dfs -put <local_data_file> <inputdir>
```

For example, define the data file **data.txt** as follows:

```
12005000201,Zhang San,Male,19,City a, Province a  
12005000202,Li Wanting,Female,23,City b, Province b  
12005000203,Wang Ming,Male,26,City c, Province c  
12005000204,Li Gang,Male,18,City d, Province d  
12005000205,Zhao Enru,Female,21,City e, Province e  
12005000206,Chen Long,Male,32,City f, Province f  
12005000207,Zhou Wei,Female,29,City g, Province g  
12005000208,Yang Yiwen,Female,30,City h, Province h  
12005000209,Xu Bing,Male,26,City i, Province i  
12005000210,Xiao Kai,Male,25,City j, Province j
```

Run the following commands:

```
hdfs dfs -mkdir /datadirImport
```

```
hdfs dfs -put data.txt /datadirImport
```

2. Go to HBase shell and run the following command to create the **bulkTable** table:

```
create 'bulkTable', {NAME => 'info',COMPRESSION => 'SNAPPY',  
DATA_BLOCK_ENCODING => 'FAST_DIFF'},{NAME=>'address'}
```

After the execution is complete, exit the HBase shell.

3. Run the following commands to generate an HFile file (StoreFiles):

```
hbase org.apache.hadoop.hbase.index.mapreduce.HIndexImportTsv -  
Dimporttsv.separator=<separator>
```

```
-Dimporttsv.bulk.output=</path/for/output> -
```

```
Dindexspecs.to.add=<indexspecs> -Dimporttsv.columns=<columns>  
tableName <inputdir>
```

- **-Dimport.separator**: indicates a separator, for example, -
Dimport.separator=','.
- **-Dimport.bulk.output=</path/for/output>**: indicates the output path of the execution result. You need to specify a path that does not exist.
- **<columns>**: Indicates the mapping of the imported data in a table, for example, -
**Dimporttsv.columns=HBASE_ROW_KEY,info:name,info:gender,info:age,
address:city,address:province.**

- **<tablename>**: Indicates the name of a table to be operated.
- **<inputdir>**: Indicates the directory where data is loaded in batches.
- **-Dindexspecs.to.add=<indexspecs>**: Indicates the mapping between an index name and a column, for example, -
Dindexspecs.to.add='index_bulk=>info:[age->String]'. The index composition can be represented as follows:

```
indexNameN=>familyN :[columnQualifierN-> columnQualifierDataType],
[columnQualifierM-> columnQualifierDataType];familyM:
[columnQualifierO-> columnQualifierDataType]# indexNameN=>
familyM: [columnQualifierO-> columnQualifierDataType]
```

Column qualifiers are separated by commas (,).

Example: "index1 => f1:[c1-> String],[c2-> String]"

Column families are separated by semicolons (;).

Example: "index1 => f1:[c1-> String],[c2-> String]; f2:[c3-> Long]"

Multiple indexes are separated by pound keys (#).

Example: "index1 => f1:[c1-> String],[c2-> String]; f2:[c3-> Long]#index2
=> f2:[c3-> Long]"

The following data types are supported by columns.

Available data types are as follows: STRING, INTEGER, FLOAT, LONG, DOUBLE, SHORT, BYTE, CHAR

NOTE

Data types can also be transferred in lowercase.

For example, run the following command:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.HIndexImportTsv -  
Dimporttsv.separator=',' -Dimporttsv.bulk.output=/dataOutput -  
Dindexspecs.to.add='index_bulk=>info:[age->String]' -  
Dimporttsv.columns=HBASE_ROW_KEY,info:name,info:gender,info:age,add  
ress:city,address:province bulkTable /datadirImport/data.txt
```

Command output:

```
[root@shap000000406 opt]# hbase org.apache.hadoop.hbase.hindex.mapreduce.HIndexImportTsv -
Dimporttsv.separator=',' -Dimporttsv.bulk.output=/dataOutput -Dindexspecs.to.add='index_bulk=>info:
[age->String]' -
Dimporttsv.columns=HBASE_ROW_KEY,info:name,info:gender,info:age,address:city,address:province
bulkTable /datadirImport/data.txt
2018-05-08 21:29:16,059 INFO [main] mapreduce.HFileOutputFormat2: Incremental table bulkTable
output configured.
2018-05-08 21:29:16,069 INFO [main] client.ConnectionManager$HConnectionImplementation:
Closing master protocol: MasterService
2018-05-08 21:29:16,069 INFO [main] client.ConnectionManager$HConnectionImplementation:
Closing zookeeper sessionId=0x80007c2cb4fd5b4d
2018-05-08 21:29:16,072 INFO [main] zookeeper.ZooKeeper: Session: 0x80007c2cb4fd5b4d closed
2018-05-08 21:29:16,072 INFO [main-EventThread] zookeeper.ClientCnxn: EventThread shut down
for session: 0x80007c2cb4fd5b4d
2018-05-08 21:29:16,379 INFO [main] client.ConfiguredRMFailoverProxyProvider: Failing over to 147
2018-05-08 21:29:17,328 INFO [main] input.FileInputFormat: Total input files to process : 1
2018-05-08 21:29:17,413 INFO [main] mapreduce.JobSubmitter: number of splits:1
2018-05-08 21:29:17,430 INFO [main] Configuration.deprecation: io.bytes.per.checksum is
deprecated. Instead, use dfs.bytes-per-checksum
2018-05-08 21:29:17,687 INFO [main] mapreduce.JobSubmitter: Submitting tokens for job:
job_1525338489458_0002
2018-05-08 21:29:18,100 INFO [main] impl.YarnClientImpl: Submitted application
application_1525338489458_0002
2018-05-08 21:29:18,136 INFO [main] mapreduce.Job: The url to track the job: http://
```

```
shap000000407:8088/proxy/application_1525338489458_0002/
2018-05-08 21:29:18,136 INFO [main] mapreduce.Job: Running job: job_1525338489458_0002
2018-05-08 21:29:28,248 INFO [main] mapreduce.Job: Job job_1525338489458_0002 running in uber
mode : false
2018-05-08 21:29:28,249 INFO [main] mapreduce.Job: map 0% reduce 0%
2018-05-08 21:29:38,344 INFO [main] mapreduce.Job: map 100% reduce 0%
2018-05-08 21:29:51,421 INFO [main] mapreduce.Job: map 100% reduce 100%
2018-05-08 21:29:51,428 INFO [main] mapreduce.Job: Job job_1525338489458_0002 completed
successfully
2018-05-08 21:29:51,523 INFO [main] mapreduce.Job: Counters: 50
```

4. Run the following command to import the generated HFile to HBase:

```
hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles </
path/for/output> <tablename>
```

For example, run the following command:

```
hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles /
dataOutput bulkTable
```

Command output:

```
[root@shap000000406 opt]# hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles /
dataOutput bulkTable
2018-05-08 21:30:01,398 WARN [main] mapreduce.LoadIncrementalHFiles: Skipping non-directory
hdfs://hacluster/dataOutput/_SUCCESS
2018-05-08 21:30:02,006 INFO [LoadIncrementalHFiles-0] hfile.CacheConfig: Created cacheConfig:
CacheConfig:disabled
2018-05-08 21:30:02,006 INFO [LoadIncrementalHFiles-2] hfile.CacheConfig: Created cacheConfig:
CacheConfig:disabled
2018-05-08 21:30:02,006 INFO [LoadIncrementalHFiles-1] hfile.CacheConfig: Created cacheConfig:
CacheConfig:disabled
2018-05-08 21:30:02,085 INFO [LoadIncrementalHFiles-2] compress.CodecPool: Got brand-new
decompressor [.snappy]
2018-05-08 21:30:02,120 INFO [LoadIncrementalHFiles-0] mapreduce.LoadIncrementalHFiles: Trying
to load hfile=hdfs://hacluster/dataOutput/address/042426c252f74e859858c7877b95e510
first=12005000201 last=12005000210
2018-05-08 21:30:02,120 INFO [LoadIncrementalHFiles-2] mapreduce.LoadIncrementalHFiles: Trying
to load hfile=hdfs://hacluster/dataOutput/info/f3995920ae0247a88182f637aa031c49
first=12005000201 last=12005000210
2018-05-08 21:30:02,128 INFO [LoadIncrementalHFiles-1] mapreduce.LoadIncrementalHFiles: Trying
to load hfile=hdfs://hacluster/dataOutput/d/c53b252248af42779f29442ab84f86b8 first=\x00index_bulk
\x00\x00\x00\x00\x00\x00\x00\x0018\x00\x0012005000204 last=\x00index_bulk
\x00\x00\x00\x00\x00\x00\x00\x0032\x00\x0012005000206
2018-05-08 21:30:02,231 INFO [main] client.ConnectionManager$HConnectionImplementation:
Closing master protocol: MasterService
2018-05-08 21:30:02,231 INFO [main] client.ConnectionManager$HConnectionImplementation:
Closing zookeeper sessionId=0x81007c2cf0f55cc5
2018-05-08 21:30:02,235 INFO [main] zookeeper.ZooKeeper: Session: 0x81007c2cf0f55cc5 closed
2018-05-08 21:30:02,235 INFO [main-EventThread] zookeeper.ClientCnxn: EventThread shut down
for session: 0x81007c2cf0f55cc5
```

8.8.3 Using an Index Generation Tool

Scenarios

To quickly create indexes for user data, HBase provides the TableIndexer tool for you to create, add, and delete indexes using MapReduce functions. The application scenarios are as follows:

- You want to add an index for a specified column in a table where a large amount of data exists. However, if you use the **addIndicesWithData()** API to add an index, index data corresponding to the related user data will be generated, which is time-consuming. If you use **addIndices()** to create an index, index data corresponding to user data will not be generated. Therefore,

to create index data for user data, you can use the TableIndexer tool to create an index.

- If the index data is inconsistent with the user data, the tool can be used to rebuild index data.

If you temporarily disable the index, put new data to the disabled index column, and then directly enable the index from the disabled state, index data and user data may be inconsistent. Therefore, you must rebuild all index data before using it again.

- You can use the TableIndexer tool to completely delete a large amount of existing index data from a user table.
- For user tables that do not have indexes, this tool allows you to add and build indexes at the same time.

How to Use

- **Adding a new index to a user table**

The command is as follows:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.TableIndexer -  
Dtablename.to.index=tablename -Dindexspecs.to.add='idx_0=>cf_0:[q_0-  
>string],[q_1];cf_1:[q_2],[q_3]#idx_1=>cf_1:[q_4]'
```

The following parameters are required.

- **tablename.to.index**: Indicates the name of a table for which an index is created.
- **indexspecs.to.add**: Indicates the mapping between the index name and the column in the corresponding user table.
- **scan.caching** (optional): Contains an integer value, indicating the number of cached rows to be transmitted to the scanner during data table scanning.

The parameters in the preceding command are described as follows:

- **idx_1**: Indicates an index name.
- **cf_0**: Indicates the name of a column family.
- **q_0**: Indicates the name of a column.
- **string**: Indicates a data type. The parameter value can be STRING, INTEGER, FLOAT, LONG, DOUBLE, SHORT, BYTE, or CHAR.

NOTE

- The pound key (#) is used to separate indexes. The semicolon (;) is used to separate column families. The comma (,) is used to separate column qualifiers.
- The column name and its data type must be included in '[]'.
- Column names and their data types are separated by '->'.
- If the data type of a specific column is not specified, the default data type (string) is used.
- If **scan.caching** is not configured, the default value **1000** is used.
- The user table must exist.
- The index specified in the table must not exist.
- If a column family named **d** exists in the user table, you must use the TableIndexer tool to build index data.

After the preceding command is executed, the specified index is added to the table and is in INACTIVE state. This behavior is similar to the **addIndices()** API.

- **Creating index data for existing indexes in a user table**

The command is as follows:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.TableIndexer -  
Dtablename.to.index=tablename -Dindexnames.to.build='idx_0#idx_1'
```

The following parameters are required.

- **tablename.to.index**: Indicates the name of a table for which an index is created.
- **indexspecs.to.build**: Indicates an index name.
- **scan.caching** (optional): Contains an integer value, indicating the number of cached rows to be transmitted to the scanner during data table scanning.

The parameters in the preceding command are described as follows:

- **idx_1**: Indicates an index name.

 **NOTE**

- The pound key (#) is used to separate index names.
- If **scan.caching** is not configured, the default value **1000** is used.
- The user table must exist.

After the preceding command is executed, the specified index is set to the ACTIVE state. Users can use them when scanning data.

- **Deleting the existing indexes and their data from a user table**

The command is as follows:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.TableIndexer -  
Dtablename.to.index=tablename -Dindexnames.to.drop='idx_0#idx_1'
```

The following parameters are required.

- **tablename.to.index**: Indicates the name of a table for which an index is created.
- **indexnames.to.drop**: Indicates the name of the index that should be deleted with its data (must exist in the table).
- **scan.caching** (optional): Contains an integer value, indicating the number of cached rows to be transmitted to the scanner during data table scanning.

The parameters in the preceding command are described as follows:

- **idx_1**: Indicates an index name.

 **NOTE**

- The pound key (#) is used to separate index names.
- If **scan.caching** is not configured, the default value **1000** is used.
- The user table must exist.

After the preceding command is executed, the specified index is deleted from the table.

- **Adding new indexes to user tables and building data based on existing data**

The command is as follows:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.TableIndexer -  
Dtablename.to.index=tablename -Dindexspecs.to.add='idx_0 => cf_0:[q_0-  
> string],[q_1];cf_1:[q_2],[q_3]#idx_1 => cf_1:[q_4]' -  
Dindexnames.to.build='idx_0'
```

 NOTE

- The parameters are the same as the previous ones.
- The user table must exist.
- The indexes specified in **indexspecs.to.add** must not exist in the table.
- The index names specified in **indexnames.to.build** must exist in the table or be part of the value of **indexspecs.to.add**.

After the preceding command is executed, all indexes specified in **indexspecs.to.add** will be added to this table, and index data will be built for all specified indexes using **indexnames.to.build**.

8.8.4 Migrating Index Data

Scenario

The indexes used in MRS 1.7 or later are incompatible with secondary indexes used by HBase in earlier MRS versions. Therefore, you need to perform the following operations to migrate index data from an earlier version (MRS 1.5 or earlier) to MRS 1.7 or later.

Prerequisites

1. During data migration, the cluster of the old version must be MRS 1.5 or earlier, and the cluster of the new version must be MRS 1.7 or later.
2. Before data migration, you must have old index data.
3. A cross-cluster mutual trust relationship must be configured and the inter-cluster replication function must be enabled for a security cluster. For a common cluster, only the inter-cluster replication function needs to be enabled. For details, see .

Procedure

Migrate user data from an old cluster to a new cluster. To migrate data, you need to manually synchronize data of the old and new clusters in a single table by export, distcp, and import.

For example, the current old cluster has a user table (**t1**, index name: **idx_t1**) and its corresponding index table (**t1_idx**). Perform the following operations to migrate data.

1. Export table data from the old cluster.

```
hbase org.apache.hadoop.hbase.mapreduce.Export -Dhbase.mapreduce.include.deleted.rows=true  
<tableName> <path/for/data>
```

 - **<tableName>**: Indicates a table name, for example, **t1**.
 - **<path/for/data>**: Indicates the path for storing source data, for example, **/user/hbase/t1**.

Example: **hbase org.apache.hadoop.hbase.mapreduce.Export -Dhbase.mapreduce.include.deleted.rows=true t1 /user/hbase/t1**

2. Copy the exported data to the new cluster as follows:

```
hadoop distcp <path/for/data> hdfs://ActiveNameNodeIP:9820/<path/for/newData>
```

- *<path/for/data>*: Indicates the path for storing source data in the old cluster, for example, **/user/hbase/t1**.
- *<path/for/newData>*: Indicates the path for storing source data in the new cluster, for example, **/user/hbase/t1**.

ActiveNameNodeIP indicates the IP address of the active NameNode in the new cluster.

Example: **hadoop distcp /user/hbase/t1 hdfs://192.168.40.2:9820/user/hbase/t1**

 **NOTE**

- Manually copy the exported data to HDFS of the new cluster, for example, **/user/hbase/t1**.
3. Use the HBase table user of the new cluster to generate HFiles in the new cluster.

```
hbase org.apache.hadoop.hbase.mapreduce.Import -Dimport.bulk.output=<path/for/hfiles>  
<tableName><path/for/newData>
```

- *<path/for/hfiles>*: Indicates the path of the HFiles generated in the new cluster, for example, **/user/hbase/output_t1**.
- *<tableName>*: Indicates a table name, for example, **t1**.
- *<path/for/newData>*: Indicates the path for storing source data in the new cluster, for example, **/user/hbase/t1**.

Example:

hbase org.apache.hadoop.hbase.mapreduce.Import -Dimport.bulk.output=/user/hbase/output_t1 t1 /user/hbase/t1

4. Import the generated HFiles to the table in the new cluster.

The command is as follows:

```
hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles <path/for/hfiles> <tableName>
```

- *<path/for/hfiles>*: Indicates the path of the HFiles generated in the new cluster, for example, **/user/hbase/output_t1**.
- *<tableName>*: Indicates a table name, for example, **t1**.

Example:

hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles /user/hbase/output_t1 t1

 **NOTE**

1. The preceding shows the process of migrating user data. You only need to perform the first three steps to migrate the index data of the old cluster and change the corresponding table name to an index table name (for example, **t1_idx**).
 2. Skip **4** when migrating index data.
5. Import index data to a table in the new cluster.
 - a. Add an index the same as that of the user table of the previous version to the user table of the new cluster (the user table cannot contain a column family named **d**).

The command is as follows:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.TableIndexer -  
Dtablename.to.index=<tableName> -Dindexspecs.to.add=<indexspecs>
```

- **-Dtablename.to.index=<tableName>**: Indicates a table name, for example, **-Dtablename.to.index=t1**.
- **-Dindexspecs.to.add=<indexspecs>**: Indicates the mapping between an index name and a column, for example, **-Dindexspecs.to.add='idx_t1=>info:[name->String]'**.

Example:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.TableIndexer -  
Dtablename.to.index=t1 -Dindexspecs.to.add='idx_t1=>info:[name->  
String]'
```

 NOTE

If a column family named **d** exists in the user table, you must use the TableIndexer tool to build index data.

- b. Run the LoadIncrementalHFiles tool to load the index data of the old cluster to a table in the new cluster.

The command is as follows:

```
hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles </path/for/hfiles>  
<tableName>
```

- **</path/for/hfiles>**: Indicates the path of index data on HDFS. The path is the index generation path specified in **-Dimport.bulk.output**, for example, **/user/hbase/output_t1_idx**.
- **<tableName>**: Indicates a table name of the new cluster, for example, **t1**.

Example:

```
hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles /  
user/hbase/output_t1_idx t1
```

8.9 Configuring HBase DR

Scenario

HBase disaster recovery (DR), a key feature that is used to ensure high availability (HA) of the HBase cluster system, provides the real-time remote DR function for HBase. HBase DR provides basic O&M tools, including tools for maintaining and re-establishing DR relationships, verifying data, and querying data synchronization progress. To implement real-time DR, back up data of an HBase cluster to another HBase cluster. DR in the HBase table common data writing and BulkLoad batch data writing scenarios is supported.

 NOTE

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

Prerequisites

- The active and standby clusters are successfully installed and started, and you have the administrator permissions on the clusters.
- Ensure that the network connection between the active and standby clusters is normal and ports are available.
- If the active cluster is deployed in security mode and is not managed by one FusionInsight Manager, cross-cluster trust relationship has been configured for the active and standby clusters.. If the active cluster is deployed in normal mode, no cross-cluster mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters.
- Time is consistent between the active and standby clusters and the NTP service on the active and standby clusters uses the same time source.
- The mapping between host names and service IP addresses of all nodes in the active and standby clusters have been configured in the **hosts** file of these nodes.

NOTE

If the client of the active cluster is installed on a node outside the cluster, the mapping between host names and service IP addresses of all nodes in the active and standby clusters must have been configured in the **hosts** file of these nodes.

- The network bandwidth between the active and standby clusters is determined based on service volume, which cannot be less than the possible maximum service volume.
- The MRS versions of the active and standby clusters must be the same.
- The scale of the standby cluster must be greater than or equal to that of the active cluster.

Constraints

- Although DR provides the real-time data replication function, the data synchronization progress is affected by many factors, such as the service volume in the active cluster and the health status of the standby cluster. In normal cases, the standby cluster should not take over services. In extreme cases, system maintenance personnel and other decision makers determine whether the standby cluster takes over services according to the current data synchronization indicators.
- HBase clusters must be deployed in active/standby mode.
- Table-level operations on the DR table of the standby cluster are forbidden, such as modifying the table attributes and deleting the table. Misoperations on the standby cluster will cause data synchronization failure of the active cluster. As a result, table data in the standby cluster is lost.
- If the DR data synchronization function is enabled for HBase tables of the active cluster, the DR table structure of the standby cluster needs to be modified to ensure table structure consistency between the active and standby clusters during table structure modification.

Procedure

Configuring the common data writing DR parameters for the active cluster

- Step 1** Log in to Manager of the active cluster.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations** and click **All Configurations**. The HBase configuration page is displayed.
- Step 3** (Optional) [Table 8-7](#) describes the optional configuration items during HBase DR. You can set the parameters based on the description or use the default values.

Table 8-7 Optional configuration items

Navigation Path	Parameter	Default Value	Description
HMaster > Performance	hbase.master.logcleaner.ttl	600000	Specifies the retention period of HLog. If the value is set to 604800000 (unit: millisecond), the retention period of HLog is 7 days.
	hbase.master.cleaner.interval	60000	Interval for the HMaster to delete historical HLog files. The HLog that exceeds the configured period will be automatically deleted. You are advised to set it to the maximum value to save more HLogs.
RegionServer > Replication	replication.source.size.capacity	16777216	Maximum size of edits, in bytes. If the edit size exceeds the value, HLog edits will be sent to the standby cluster.
	replication.source.nb.capacity	25000	Maximum number of edits, which is another condition for triggering HLog edits to be sent to the standby cluster. After data in the active cluster is synchronized to the standby cluster, the active cluster reads and sends data in HLog according to this parameter value. This parameter is used together with replication.source.size.capacity .
	replication.source.maxretriesmultiplier	10	Maximum number of retries when an exception occurs during replication.
	replication.source.sleepforretries	1000	Retry interval (Unit: ms)

Navigation Path	Parameter	Default Value	Description
	hbase.regionserver.replication.handler.count	6	Number of replication RPC server instances on RegionServer

Configuring the BulkLoad batch data writing DR parameters for the active cluster

Step 4 Determine whether to enable the BulkLoad batch data writing DR function.

If yes, go to [Step 5](#).

If no, go to [Step 8](#).

Step 5 Choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations** and click **All Configurations**. The HBase configuration page is displayed.

Step 6 Search for **hbase.replication.bulkload.enabled** and change its value to **true** to enable the BulkLoad batch data writing DR function.

Step 7 Search for **hbase.replication.cluster.id** and change the HBase ID of the active cluster. The ID is used by the standby cluster to connect to the active cluster. The value can contain uppercase letters, lowercase letters, digits, and underscores (_), and cannot exceed 30 characters.

Restarting the HBase service and install the client

Step 8 Click **Save**. In the displayed dialog box, click **OK**. Restart the HBase service.

Step 9 In the active and standby clusters, choose **Cluster** > *Name of the desired cluster* > **Service** > **HBase** > **More** > **Download Client** to download the client and install it.

Adding the DR relationship between the active and standby clusters

Step 10 Log in as user **hbase** to the HBase shell page of the active cluster.

Step 11 Run the following command on HBase Shell to create the DR synchronization relationship between the active cluster HBase and the standby cluster HBase.

```
add_peer 'Standby cluster ID', CLUSTER_KEY => "ZooKeeper service IP address in the standby cluster", CONFIG => {"hbase.regionserver.kerberos.principal" => "Standby cluster RegionServer principal", "hbase.master.kerberos.principal" => "Standby cluster HMaster principal"}
```

- The standby cluster ID indicates the ID for the active cluster to recognize the standby cluster. Enter an ID. The value can be specified randomly. Digits are recommended.
- The ZooKeeper address of the standby cluster includes the service IP address of ZooKeeper, the port for listening to client connections, and the HBase root directory of the standby cluster on ZooKeeper.

- Search for **hbase.master.kerberos.principal** and **hbase.regionserver.kerberos.principal** in the HBase **hbase-site.xml** configuration file of the standby cluster.

For example, to add the DR relationship between the active and standby clusters, run the **add_peer** '*Standby cluster ID*', **CLUSTER_KEY** => **"192.168.40.2,192.168.40.3,192.168.40.4:24002:/hbase"**, **CONFIG** => **{"hbase.regionserver.kerberos.principal" => "hbase/hadoop.hadoop.com@HADOOP.COM", "hbase.master.kerberos.principal" => "hbase/hadoop.hadoop.com@HADOOP.COM"}**

Step 12 (Optional) If the BulkLoad batch data write DR function is enabled, the HBase client configuration of the active cluster must be copied to the standby cluster.

- Create the **/hbase/replicationConf/hbase.replication.cluster.id of the active cluster** directory in the HDFS of the standby cluster.
- HBase client configuration file, which is copied to the **/hbase/replicationConf/hbase.replication.cluster.id of the active cluster** directory of the HDFS of the standby cluster.

Example: **hdfs dfs -put HBase/hbase/conf/core-site.xml HBase/hbase/conf/hdfs-site.xml HBase/hbase/conf/yarn-site.xml hdfs://NameNode IP.25000/hbase/replicationConf/source_cluster**

Enabling HBase DR to synchronize data

Step 13 Check whether a naming space exists in the HBase service instance of the standby cluster and the naming space has the same name as the naming space of the HBase table for which the DR function is to be enabled.

- If the same namespace exists, go to [Step 14](#).
- If no, create a naming space with the same name in the HBase shell of the standby cluster and go to [Step 14](#).

Step 14 In the HBase shell of the active cluster, run the following command as user **hbase** to enable the real-time DR function for the table data of the active cluster to ensure that the data modified in the active cluster can be synchronized to the standby cluster in real time.

You can only synchronize the data of one HTable at a time.

enable_table_replication '*table name*'

 NOTE

- If the standby cluster does not contain a table with the same name as the table for which real-time synchronization is to be enabled, the table is automatically created.
- If a table with the same name as the table for which real-time synchronization is to be enabled exists in the standby cluster, the structures of the two tables must be the same.
- If the encryption algorithm SMS4 or AES is configured for '*Table name*', the function for synchronizing data from the active cluster to the standby cluster cannot be enabled for the HBase table.
- If the standby cluster is offline or has tables with the same name but different structures, the DR function cannot be enabled.
- If the DR data synchronization function is enabled for some Phoenix tables in the active cluster, the standby cluster cannot have common HBase tables with the same names as the Phoenix tables in the active cluster. Otherwise, the DR function fails to be enabled or the tables with the names in the standby cluster cannot be used properly.
- If the DR data synchronization function is enabled for Phoenix tables in the active cluster, you need to enable the DR data synchronization function for the metadata tables of the Phoenix tables. The metadata tables include SYSTEM.CATALOG, SYSTEM.FUNCTION, SYSTEM.SEQUENCE, and SYSTEM.STATS.
- If the DR data synchronization function is enabled for HBase tables of the active cluster, after adding new indexes to HBase tables, you need to manually add secondary indexes to DR tables in the standby cluster to ensure secondary index consistency between the active and standby clusters.

Step 15 (Optional) If HBase does not use Ranger, run the following command as user **hbase** in the HBase shell of the active cluster to enable the real-time permission to control data DR function for the HBase tables in the active cluster.

```
enable_table_replication 'hbase:acl'
```

Creating Users

Step 16 Log in to FusionInsight Manager of the standby cluster, choose **System > Permission > Role > Create Role** to create a role, and add the same permission for the standby data table to the role based on the permission of the HBase source data table of the active cluster.

Step 17 Choose **System > Permission > User > Create** to create a user. Set the **User Type** to **Human-Machine** or **Machine-Machine** based on service requirements and add the user to the created role. Access the HBase DR data of the standby cluster as the newly created user.

 NOTE

- After the permission of the active HBase source data table is modified, to ensure that the standby cluster can properly read data, modify the role permission for the standby cluster.
- If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for HBase](#).

Synchronizing the table data of the active cluster

Step 18 After HBase DR is configured and data synchronization is enabled, check whether tables and data exist in the active cluster and whether the historical data needs to be synchronized to the standby cluster.

- If yes, a table exists and data needs to be synchronized. Log in as the HBase table user to the node where the HBase client of the active cluster is installed

and run the `kinit` username to authenticate the identity. The user must have the read and write permissions on tables and the execute permission on the `hbase:meta` table. Then go to [Step 19](#).

- If no, no further action is required.

Step 19 The HBase DR configuration does not support automatic synchronization of historical data in tables. You need to back up the historical data of the active cluster and then manually restore the historical data in the standby cluster.

Manual recovery refers to the recovery of a single table, which can be performed through `Export`, `DistCp`, or `Import`.

To manually recover a single table, perform the following steps:

1. Export table data from the active cluster.

hbase org.apache.hadoop.hbase.mapreduce.Export -
Dhbase.mapreduce.include.deleted.rows=true *Table name Directory where the source data is stored*

Example: **hbase org.apache.hadoop.hbase.mapreduce.Export -**
Dhbase.mapreduce.include.deleted.rows=true t1 /user/hbase/t1

2. Copy the data that has been exported to the standby cluster.

hadoop distcp *directory where the source data is stored on the active cluster*
hdfs://ActiveNameNodeIP:8020/directory *where the source data is stored on the standby cluster*

ActiveNameNodeIP indicates the IP address of the active NameNode in the standby cluster.

Example: **hadoop distcp /user/hbase/t1 hdfs://192.168.40.2:8020/user/hbase/t1**

3. Import data to the standby cluster as the HBase table user of the standby cluster.

On the HBase shell screen of the standby cluster, run the following command as user `hbase` to retain the data writing status:

set_clusterState_active

The command is run successfully if the following information is displayed:

```
hbase(main):001:0> set_clusterState_active  
=> true
```

hbase org.apache.hadoop.hbase.mapreduce.Import -
Dimport.bulk.output=Directory *where the output data is stored in the standby cluster* **Table name** *Directory where the source data is stored in the standby cluster*

hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles
Directory *where the output data is stored in the standby cluster* **Table name**

Example:

```
hbase(main):001:0> set_clusterState_active  
=> true
```

hbase org.apache.hadoop.hbase.mapreduce.Import -
Dimport.bulk.output=/user/hbase/output_t1 t1 /user/hbase/t1

hbase org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles /user/hbase/output_t1 t1

Step 20 Run the following command on the HBase client to check the synchronized data of the active and standby clusters. After the DR data synchronization function is enabled, you can run this command to check whether the newly synchronized data is consistent.

hbase org.apache.hadoop.hbase.mapreduce.replication.VerifyReplication --starttime=*Start time* --endtime=*End time* *Column family name* *ID of the standby cluster* *Table name*

 **NOTE**

- The start time must be earlier than the end time.
- The values of **starttime** and **endtime** must be in the timestamp format. You need to run **date -d "2015-09-30 00:00:00" +%s** to change a common time format to a timestamp format.

Specify the data writing status for the active and standby clusters.

Step 21 On the HBase shell screen of the active cluster, run the following command as user **hbase** to retain the data writing status:

set_clusterState_active

The command is run successfully if the following information is displayed:

```
hbase(main):001:0> set_clusterState_active  
=> true
```

Step 22 On the HBase shell screen of the standby cluster, run the following command as user **hbase** to retain the data read-only status:

set_clusterState_standby

The command is run successfully if the following information is displayed:

```
hbase(main):001:0> set_clusterState_standby  
=> true
```

----End

Related Commands

Table 8-8 HBase DR

Operation	Command	Description
Set up a DR relationship.	<pre>add_peer '<i>Standby cluster ID</i>, CLUSTER_KEY => "<i>Standby cluster ZooKeeper service IP address</i>", CONFIG => {"hbase.regionserver.kerberos.principal" => "<i>Standby cluster RegionServer principal</i>", "hbase.master.kerberos.principal" => "<i>Standby cluster HMaster principal</i>"}</pre> <p>add_peer '1','zk1,zk2,zk3:2181:/hbase1' 2181: port number of ZooKeeper in the cluster</p>	<p>Set up the relationship between the active cluster and the standby cluster.</p> <p>If BulkLoad batch data write DR is enabled:</p> <ul style="list-style-type: none"> • Create the /hbase/replicationConf/hbase.replication.cluster.id of the active cluster directory in the HDFS of the standby cluster. • HBase client configuration file, which is copied to the /hbase/replicationConf/hbase.replication.cluster.id of the active cluster directory of the HDFS of the standby cluster.
Remove the DR relationship.	<pre>remove_peer '<i>Standby cluster ID</i>'</pre> <p>Example: remove_peer '1'</p>	Remove standby cluster information from the active cluster.
Querying the DR Relationship	list_peers	Query standby cluster information (mainly Zookeeper information) in the active cluster.
Enable the real-time user table synchronization function.	<pre>enable_table_replication '<i>Table name</i>'</pre> <p>Example: enable_table_replication 't1'</p>	Synchronize user tables from the active cluster to the standby cluster.
Disable the real-time user table synchronization function.	<pre>disable_table_replication '<i>Table name</i>'</pre> <p>Example: disable_table_replication 't1'</p>	Do not synchronize user tables from the active cluster to the standby cluster.

Operation	Command	Description
<p>Verify data of the active and standby clusters.</p>	<p>bin/hbase org.apache.hadoop.hbase.mapreduce.replication.VerifyReplication <i>--starttime=Start time --endtime=End time Column family name Standby cluster ID Table name</i></p>	<p>Verify whether data of the specified table is the same between the active cluster and the standby cluster.</p> <p>The description of the parameters in this command is as follows:</p> <ul style="list-style-type: none"> • Start time: If start time is not specified, the default value 0 will be used. • End time: If end time is not specified, the time when the current operation is submitted will be used by default. • Table name: If a table name is not entered, all user tables for which the real-time synchronization function is enabled will be verified by default.
<p>Switch the data writing status.</p>	<p>set_clusterState_active set_clusterState_standby</p>	<p>Specifies whether data can be written to the cluster HBase tables.</p>

Operation	Command	Description
Add or update the active cluster HDFS configurations saved in the peer cluster.	hdfs dfs -put -f HBase/hbase/conf/core-site.xml HBase/hbase/conf/hdfs-site.xml HBase/hbase/conf/yarn-site.xml hdfs://Standby cluster NameNode IP:PORT/hbase/replicationConf/Active cluster/hbase.replication.cluster.id	<p>Enable DR for data including bulkload data. When HDFS parameters are modified in the active cluster, the modification cannot be automatically synchronized from the active cluster to the standby cluster. You need to manually run the command to synchronize configuration. The affected parameters are as follows:</p> <ul style="list-style-type: none"> • fs.defaultFS • dfs.client.failover.proxy.provider.hacluster • dfs.client.failover.connection.retries.on.timeouts • dfs.client.failover.connection.retries <p>For example, change fs.defaultFS to hdfs://hacluster_sale, HBase client configuration file, which is copied to the /hbase/replicationConf/hbase.replication.cluster.id of the active cluster directory of the HDFS of the standby cluster.</p>

8.10 Configuring HBase Data Compression and Encoding

Scenario

HBase encodes data blocks in HFiles to reduce duplicate keys in KeyValues, reducing used space. Currently, the following data block encoding modes are supported: NONE, PREFIX, DIFF, FAST_DIFF, and ROW_INDEX_V1. NONE indicates that data blocks are not encoded. HBase also supports compression algorithms for HFile compression. The following algorithms are supported by default: NONE, GZ, SNAPPY, and ZSTD. NONE indicates that HFiles are not compressed.

The two methods are used on the HBase column family. They can be used together or separately.

Prerequisites

- The HBase client has been installed in a directory, for example, **/opt/client**.

- If authentication has been enabled for HBase, you must have the corresponding operation permissions. For example, you must have the creation (C) or administration (A) permission on the corresponding namespace or higher-level items to create a table, and the creation (C) or administration (A) permission on the created table or higher-level items to modify a table. For details about how to grant permissions, see [Creating HBase Roles](#).

Procedure

Setting data block encoding and compression algorithms during creation

- **Method 1: Using hbase shell**
 - a. Log in to the node where the client is installed as the client installation user.
 - b. Run the following command to go to the client directory:
cd /opt/client
 - c. Run the following command to configure environment variables:
source bigdata_env
 - d. If the Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step:
kinit Component service user
For example, **kinit hbaseuser**.
 - e. Run the following HBase client command:
hbase shell
 - f. Create a table.
create 't1', {NAME => 'f1', COMPRESSION => 'SNAPPY', DATA_BLOCK_ENCODING => 'FAST_DIFF'}

NOTE

- *t1*: indicates the table name.
 - *f1*: indicates the column family name.
 - *SNAPPY*: indicates the column family uses the SNAPPY compression algorithm.
 - *FAST_DIFF*: indicates FAST_DIFF is used for encoding.
 - The parameter in the braces specifies the column family. You can specify multiple column families using multiple braces and separate them by commas (,). For details about table creation statements, run the **help 'create'** statement in the HBase shell.
- **Method 2: Using Java APIs**

The following code snippet shows only how to set the encoding and compression modes of a column family when creating a table. For complete code for creating a table and how to use the code to create a table, see "HBase Development Guide" > "Modifying a Table" in .

```
TableDescriptorBuilder htd = TableDescriptorBuilder.newBuilder(TableName.valueOf("t1")); // Create a descriptor for table t1.
ColumnFamilyDescriptorBuilder hcd =
ColumnFamilyDescriptorBuilder.newBuilder(Bytes.toBytes("f1")); // Create a builder for column family f1.
```

```
hcd.setDataBlockEncoding(DataBlockEncoding.FAST_DIFF);// Set the encoding mode of column family f1 to FAST_DIFF.  
hcd.setCompressionType(Compression.Algorithm.SNAPPY);// Set the compression algorithm of column family f1 to SNAPPY.  
htd.setColumnFamily(hcd.build())// Add the column family f1 to the descriptor of table t1.
```

Setting or modifying the data block encoding mode and compression algorithm for an existing table

- **Method 1: Using hbase shell**

- a. Log in to the node where the client is installed as the client installation user.
- b. Run the following command to go to the client directory:
cd /opt/client
- c. Run the following command to configure environment variables:
source bigdata_env
- d. If the Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step:

kinit *Component service user*

For example, **kinit hbaseuser**.

- e. Run the following HBase client command:

hbase shell

- f. Run the following command to modify the table:

```
alter 't1', {NAME => 'f1', COMPRESSION => 'SNAPPY',  
DATA_BLOCK_ENCODING => 'FAST_DIFF'}
```

- **Method 2: Using Java APIs**

The following code snippet shows only how to modify the encoding and compression modes of a column family in an existing table. For complete code for modifying a table and how to use the code to modify a table, see "HBase Development Guide".

```
TableDescriptor htd = admin.getDescriptor(TableName.valueOf("t1"));// Obtain the descriptor of table t1.  
ColumnFamilyDescriptor originCF = htd.getColumnFamily(Bytes.toBytes("f1"));// Obtain the descriptor of column family f1.  
builder.ColumnFamilyDescriptorBuilder hcd = ColumnFamilyDescriptorBuilder.newBuilder(originCF);// Create a builder based on the existing column family attributes.  
hcd.setDataBlockEncoding(DataBlockEncoding.FAST_DIFF);// Change the encoding mode of the column family to FAST_DIFF.  
hcd.setCompressionType(Compression.Algorithm.SNAPPY);// Change the compression algorithm of the column family to SNAPPY.  
admin.modifyColumnFamily(TableName.valueOf("t1"), hcd.build());// Submit to the server to modify the attributes of column family f1.
```

After the modification, the encoding and compression modes of the existing HFile will take effect after the next compaction.

8.11 Performing an HBase DR Service Switchover

Scenario

MRS cluster administrators can configure HBase cluster DR to improve system availability. If the active cluster in the DR environment is faulty and the connection

to the HBase upper-layer application is affected, you need to configure the standby cluster information for the HBase upper-layer application so that the application can run in the standby cluster.

 **NOTE**

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

Impact on the System

After a service switchover, data written to the standby cluster is not synchronized to the active cluster by default. Add the active cluster is recovered, the data newly generated in the standby cluster needs to be synchronized to the active cluster by backup and recovery. If automatic data synchronization is required, you need to switch over the active and standby HBase DR clusters.

Procedure

Step 1 Log in to FusionInsight Manager of the standby cluster.

Step 2 Download and install the HBase client.

Step 3 On the HBase client of the standby cluster, run the following command as user **hbase** to enable the data writing status in the standby cluster.

```
kinit hbase
```

```
hbase shell
```

```
set_clusterState_active
```

The command is run successfully if the following information is displayed:

```
hbase(main):001:0> set_clusterState_active  
=> true
```

Step 4 Check whether the original configuration files **hbase-site.xml**, **core-site.xml**, and **hdfs-site.xml** of the HBase upper-layer application are modified to adapt to the application running.

- If yes, update the related content to the new configuration file and replace the old configuration file.
- If no, use the new configuration file to replace the original configuration file of the HBase upper-layer application.

Step 5 Configure the network connection between the host where the HBase upper-layer application is located and the standby cluster.

 **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.

1. Ensure that the host where the client is installed can communicate with the hosts listed in the **hosts** file in the directory where the client installation package is decompressed.
2. If the host where the client is located is not a node in the cluster, you need to set the mapping between the host name and the IP address (service plan) in

the `/etc/hosts` file on the host. The host names and IP addresses must be mapped one by one.

Step 6 Set the time of the host where the HBase upper-layer application is located to be the same as that of the standby cluster. The time difference must be less than 5 minutes.

Step 7 Check the authentication mode of the active cluster.

- If the security mode is used, go to [Step 8](#).
- If the normal mode is used, no further action is required.

Step 8 Obtain the **keytab** and **krb5.conf** configuration files of the HBase upper-layer application user.

1. On FusionInsight Manager of the standby cluster, choose **System** > **Permission** > **User**.
2. Locate the row that contains the target user, click **More** > **Download Authentication Credential** in the **Operation** column, and download the **keytab** file to the local PC.
3. Decompress the package to obtain **user.keytab** and **krb5.conf**.

Step 9 Use the **user.keytab** and **krb5.conf** files to replace the original files in the HBase upper-layer application.

Step 10 Stop upper-layer applications.

Step 11 Determine whether to switch over the active and standby HBase clusters. If the switchover is not performed, data will not be synchronized.

- If yes, switch over the active and standby HBase DR clusters. For details, see [Performing an HBase DR Active/Standby Cluster Switchover](#). Then, go to [Step 12](#).
- If no, go to [Step 12](#).

Step 12 Start the upper-layer services.

----End

8.12 Performing an HBase DR Active/Standby Cluster Switchover

Scenario

The HBase cluster in the current environment is a DR cluster. Due to some reasons, the active and standby clusters need to be switched over. That is, the standby cluster becomes the active cluster, and the active cluster becomes the standby cluster.

NOTE

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

Impact on the System

After the active and standby clusters are switched over, data cannot be written to the original active cluster, and the original standby cluster becomes the active cluster to take over upper-layer services.

Procedure

Ensuring that upper-layer services are stopped

- Step 1** Ensure that the upper-layer services have been stopped. If not, perform operations by referring to [Performing an HBase DR Service Switchover](#).

Disabling the write function of the active cluster

- Step 2** Download and install the HBase client.

- Step 3** On the HBase client of the standby cluster, run the following command as user **hbase** to disable the data write function of the standby cluster:

```
kinit hbase
```

```
hbase shell
```

```
set_clusterState_standby
```

The command is run successfully if the following information is displayed:

```
hbase(main):001:0> set_clusterState_standby  
=> true
```

Checking whether the active/standby synchronization is complete

- Step 4** Run the following command to ensure that the current data has been synchronized (SizeOfLogQueue=0 and SizeOfLogToReplicate=0 are required). If the values are not 0, wait and run the following command repeatedly until the values are 0.

```
status 'replication'
```

Disabling synchronization between the active and standby clusters

- Step 5** Query all synchronization clusters and obtain the value of **PEER_ID**.

```
list_peers
```

- Step 6** Delete all synchronization clusters.

```
remove_peer 'Standby cluster ID'
```

Example:

```
remove_peer '1'
```

- Step 7** Query all synchronized tables.

```
list_replicated_tables
```

- Step 8** Disable all synchronized tables queried in the preceding step.

```
disable_table_replication 'Table name'
```

Example:

```
disable_table_replication 't1'
```

Performing an active/standby switchover

Step 9 Reconfigure HBase DR. For details, see [Configuring HBase DR](#).

----End

8.13 Community BulkLoad Tool

The Apache HBase official website provides the function of importing data in batches. For details, see the description of the **Import** and **ImportTsv** tools at <http://hbase.apache.org/2.2/book.html#tools>.

8.14 Configuring the MOB

Scenario

In the actual application scenario, data in various sizes needs to be stored, for example, image data and documents. Data whose size is smaller than 10 MB can be stored in HBase. HBase can yield the best read-and-write performance for data whose size is smaller than 100 KB. If the size of data stored in HBase is greater than 100 KB or even reaches 10 MB and the same number of data files are inserted, the total data amount is large, causing frequent compaction and split, high CPU consumption, high disk I/O frequency, and low performance.

MOB data (100 KB to 10 MB data) is stored in a file system (such as the HDFS) in the HFile format. Files are centrally managed using the `expiredMobFileCleaner` and `Sweeper` tools. The addresses and size of files are stored in the HBase store as values. This greatly decreases the compaction and split frequency in HBase and improves performance.

The MOB function of HBase is enabled by default. For details about related configuration items, see [Table 8-9](#). To use the MOB function, you need to specify the MOB mode for storing data in the specified column family when creating a table or modifying table attributes.

NOTE

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

Configuration Description

To enable the HBase MOB function, you need to specify the MOB mode for storing data in the specified column family when creating a table or modifying table attributes.

Use code to declare that the MOB mode for storing data is used:

```
HColumnDescriptor hcd = new HColumnDescriptor("f");  
hcd.setMobEnabled(true);
```

Use code to declare that the MOB mode for storing data is used, the unit of MOB_THRESHOLD is byte:

```
hbase(main):009:0> create 't3',{NAME => 'd', MOB_THRESHOLD => '102400', IS_MOB => 'true'}
0 row(s) in 0.3450 seconds
=> Hbase::Table - t3
hbase(main):010:0> describe 't3'
Table t3 is ENABLED

t3

COLUMN FAMILIES DESCRIPTION

{NAME => 'd', MOB_THRESHOLD => '102400', VERSIONS => '1', KEEP_DELETED_CELLS => 'FALSE',
DATA_BLOCK_ENCODING => 'NONE',
TTL => 'FOREVER', MIN_VERSIONS => '0', REPLICATION_SCOPE => '0', BLOOMFILTER => 'ROW',
IN_MEMORY => 'false', IS_MOB => 'true', COMPRESSION => 'NONE', BLOCKCACHE => 'true', BLOCKSIZE =>
'65536'}
```

Navigation path for setting parameters:

On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations** > **All Configurations**. Enter a parameter name in the search box.

Table 8-9 Parameter description

Parameter	Description	Default Value
hbase.mob.file.cache.size	Size of the opened file handle cache. If this parameter is set to a large value, more file handles can be cached, reducing the frequency of opening and closing files. However, if this parameter is set to a large value, too many file handles will be opened. The default value is 1000 . This parameter is configured on the ResionServer.	1000
hbase.mob.cache.evict.period	Expiration time of cached MOB files in the MOB cache, in seconds.	3600
hbase.mob.cache.evict.remain.ratio	Ratio of the number of retained files after MOB cache reclamation to the number of cached files. hbase.mob.cache.evict.remain.ratio is an algorithm factor. When the number of cached MOB files reaches the product of hbase.mob.file.cache.size hbase.mob.cache.evict.remain.ratio , cache reclamation is triggered.	0.5

Parameter	Description	Default Value
hbase.master.mob.ttl.cleaner.period	Interval for deleting expired files, in seconds. The default value is one day (86,400 seconds). NOTE If the validity period of an MOB file expires, that is, the file has been created for more than 24 hours, the MOB file will be deleted by the tool for deleting expired MOB files.	86400

8.15 Configuring Secure HBase Replication

Scenario

This topic provides the procedure to configure the secure HBase replication during cross-realm Kerberos setup in security mode.

Prerequisites

- Mapping for all the FQDNs to their realms should be defined in the Kerberos configuration file.
- The passwords and keytab files of **ONE.COM** and **TWO.COM** must be the same.

Procedure

Step 1 Create krbtgt principals for the two realms.

For example, if you have two realms called **ONE.COM** and **TWO.COM**, you need to add the following principals: **krbtgt/ONE.COM@TWO.COM** and **krbtgt/TWO.COM@ONE.COM**.

Add these two principals at both realms.

```
kadmin: addprinc -e "<enc_type_list>" krbtgt/ONE.COM@TWO.COM
kadmin: addprinc -e "<enc_type_list>" krbtgt/TWO.COM@ONE.COM
```

NOTE

There must be at least one common keytab mode between these two realms.

Step 2 Add rules for creating short names in Zookeeper.

Dzookeeper.security.auth_to_local is a parameter of the ZooKeeper server process. Following is an example rule that illustrates how to add support for the realm called **ONE.COM**. The principal has two members (such as **service/instance@ONE.COM**).

```
Dzookeeper.security.auth_to_local=RULE:[2:\$1@\$0](.*@\QONE.COM\E)s/@\QONE.COM\E//DEFAULT
```

The above code example adds support for the **ONE.COM** realm in a different realm. Therefore, in the case of replication, you must add a rule for the master cluster realm in the slave cluster realm. **DEFAULT** is for defining the default rule.

Step 3 Add rules for creating short names in the Hadoop processes.

The following is the **hadoop.security.auth_to_local** property in the **core-site.xml** file in the slave cluster HBase processes. For example, to add support for the **ONE.COM** realm:

```
<property>
<name>hadoop.security.auth_to_local</name>
<value>RULE:[2:$1@$0](.*@QONE.COM\E$)s/@QONE.COM\E$//DEFAULT</value>
</property>
```

 **NOTE**

If replication for bulkload data is enabled, then the same property for supporting the slave realm needs to be added in the **core-site.xml** file in the master cluster HBase processes.

Example:

```
<property>
<name>hadoop.security.auth_to_local</name>
<value>RULE:[2:$1@$0](.*@QONE.COM\E$)s/@QONE.COM\E$//DEFAULT</value>
</property>
```

----End

8.16 Configuring Region In Transition Recovery Chore Service

Scenario

In a faulty environment, there are possibilities that a region may be stuck in transition for longer duration due to various reasons like slow region server response, unstable network, ZooKeeper node version mismatch. During region transition, client operation may not work properly as some regions will not be available.

Configuration

A chore service should be scheduled at HMaster to identify and recover regions that stay in the transition state for a long time.

The following table describes the parameters for enabling this function.

Table 8-10 Parameters

Parameter	Description	Default Value
hbase.region.assignment.auto.recovery.enabled	Configuration parameter used to enable/disable the region assignment recovery thread feature.	true

8.17 Using a Secondary Index

Scenario

HIndex enables HBase indexing based on specific column values, making the retrieval of data highly efficient and fast.

Constraints

- Column families are separated by semicolons (;).
- Columns and data types must be contained in square brackets ([]).
- The column data type is specified by using -> after the column name.
- If the column data type is not specified, the default data type (string) is used.
- The number sign (#) is used to separate two index details.
- The following is an optional parameter:
-Dscan.caching: number of cached rows when the data table is scanned.
The default value is set to 1000.
- Indexes are created for a single region to repair damaged indexes.
This function is not used to generate new indexes.

Procedure

Step 1 Install the HBase client. For details, see [Using an HBase Client](#).

Step 2 Go to the client installation directory, for example, `/opt/client`.

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 If the cluster is in security mode, run the following command to authenticate the user. In normal mode, user authentication is not required.

```
kinit Component service user
```

Step 5 Run the following command to access HIndex:

```
hbase org.apache.hadoop.hbase.hindex.mapreduce.TableIndexer
```

Table 8-11 Common HIndex commands

Description	Command
Add Index	TableIndexer-Dtablename.to.index=table1-Dindexspecs.to.add='IDX1=>cf1:[q1->datatype],[q2],[q3];cf2:[q1->datatype],[q2->datatype]#IDX2=>cf1:[q5]'
Create Index	TableIndexer -Dtablename.to.index=table1 -Dindexnames.to.build='IDX1#IDX2'

Description	Command
Delete Index	TableIndexer -Dtablename.to.index=table1 - Dindexnames.to.drop='IDX1#IDX2'
Disable Index	TableIndexer -Dtablename.to.index=table1 - Dindexnames.to.disable='IDX1#IDX2'
Add and Create Index	TableIndexer -Dtablename.to.index=table1 - Dindexspecs.to.add='IDX1=>cf1:[q1->datatype],[q2],[q3];cf2: [q1->datatype],[q2->datatype]#IDX2=>cf1:[q5]' - Dindexnames.to.build='IDX1'
Create Index for a Single Region	TableIndexer -Dtablename.to.index=table1 - Dregion.to.index=regionEncodedName - Dindexnames.to.build='IDX1#IDX2'

 NOTE

- **IDX1**: indicates the index name.
- **cf1**: indicates the column family name.
- **q1**: indicates the column name.
- **datatype**: indicates the data type, including String, Integer, Double, Float, Long, Short, Byte and Char.

----End

8.18 HBase Log Overview

Log Description

Log path: The default storage path of HBase logs is `/var/log/Bigdata/hbase/Role name`.

- HMaster: `/var/log/Bigdata/hbase/hm` (run logs) and `/var/log/Bigdata/audit/hbase/hm` (audit logs)
- RegionServer: `/var/log/Bigdata/hbase/rs` (run logs) and `/var/log/Bigdata/audit/hbase/rs` (audit logs)
- ThriftServer: `/var/log/Bigdata/hbase/ts2` (run logs, **ts2** is the instance name) and `/var/log/Bigdata/audit/hbase/ts2` (audit logs, **ts2** is the instance name)

Log archive rule: The automatic log compression and archiving function of HBase is enabled. By default, when the size of a log file exceeds 30 MB, the log file is automatically compressed. The naming rule of a compressed log file is as follows: `<Original log name>-<yyyy-mm-dd_hh-mm-ss>.[ID].log.zip` A maximum of 20 latest compressed files are reserved. The number of compressed files can be configured on the Manager portal.

Table 8-12 HBase log list

Type	Name	Description
Run logs	hbase-<SSH_USER>-<process_name>-<hostname>.log	HBase system log that records the startup time, startup parameters, and most logs generated when the HBase system is running.
	hbase-<SSH_USER>-<process_name>-<hostname>.out	Log that records the HBase running environment information.
	<process_name>-<SSH_USER>-<DATE>-<PID>-gc.log	Log that records HBase junk collections.
	checkServiceDetail.log	Log that records whether the HBase service starts successfully.
	hbase.log	Log generated when the HBase service health check script and some alarm check scripts are executed.
	sendAlarm.log	Log that records alarms reported after execution of HBase alarm check scripts.
	hbase-haCheck.log	Log that records the active and standby status of HMaster
	stop.log	Log that records the startup and stop processes of HBase.
Audit logs	hbase-audit-<process_name>.log	Log that records HBase security audit.

Log Level

Table 8-13 describes the log levels supported by HBase. 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 specified level are printed. The number of printed logs decreases as the specified log level increases.

Table 8-13 Log levels

Level	Description
FATAL	Logs of this level record fatal error information about the current event processing that may result in a system crash.

Level	Description
ERROR	Logs of this level record error information about the current event processing, which indicates that system running is abnormal.
WARN	Logs of this level record abnormal information about the current event processing. These abnormalities will not result in system faults.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of the HBase service. For details, see [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the left menu bar, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.

 **NOTE**

The configurations take effect immediately without the need to restart the service.

----End

Log Formats

The following table lists the HBase log formats.

Table 8-14 Log formats

Type	Component	Format	Example
Run logs	HMaster	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> <Thread that generates the log> <Message in the log> <Location of the log event>	2020-01-19 16:04:53,558 INFO main env:HBASE_THRIFT_OPTS= org.apache.hadoop.hbase.util.ServerCommandLine.log ProcessInfo(ServerCommandLine.java:113)

Type	Component	Format	Example
	RegionServer	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> <Thread that generates the log> <Message in the log> <Location of the log event>	2020-01-19 16:05:18,589 INFO regionserver16020-SendThread(linux-k6da:2181) Client will use GSSAPI as SASL mechanism. org.apache.zookeeper.client.ZooKeeperSaslClient\$1.run(ZooKeeperSaslClient.java:285)
	ThriftServer	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> <Thread that generates the log> <Message in the log> <Location of the log event>	2020-02-16 09:42:55,371 INFO main loaded properties from hadoop-metrics2.properties org.apache.hadoop.metrics2.impl.MetricsConfig.loadFirst(MetricsConfig.java:111)
Audit logs	HMaster	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> <Thread that generates the log> <Message in the log> <Location of the log event>	2020-02-16 09:42:40,934 INFO master:linux-k6da:16000 Master: [master:linux-k6da:16000] start operation called. org.apache.hadoop.hbase.master.HMaster.run(HMaster.java:581)
	RegionServer	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> <Thread that generates the log> <Message in the log> <Location of the log event>	2020-02-16 09:42:51,063 INFO main RegionServer: [regionserver16020] start operation called. org.apache.hadoop.hbase.regionserver.HRegionServer.startRegionServer(HRegionServer.java:2396)
	ThriftServer	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> <Thread that generates the log> <Message in the log> <Location of the log event>	2020-02-16 09:42:55,512 INFO main thrift2 server start operation called. org.apache.hadoop.hbase.thrift2.ThriftServer.main(ThriftServer.java:421)

8.19 HBase Performance Tuning

8.19.1 Improving the BulkLoad Efficiency

Scenario

BulkLoad uses MapReduce jobs to directly generate files that comply with the internal data format of HBase, and then loads the generated StoreFiles to a running cluster. Compared with HBase APIs, BulkLoad saves more CPU and network resources.

ImportTSV is an HBase table data loading tool.

NOTE

This section applies to MRS 3.x and later versions.

Prerequisites

When using BulkLoad, the output path of the file has been specified using the **Dimporttsv.bulk.output** parameter.

Procedure

Add the following parameter to the BulkLoad command when performing a batch loading task:

Table 8-15 Parameter for improving BulkLoad efficiency

Parameter	Description	Value
- Dimporttsv.mapper.class	<p>The construction of key-value pairs is moved from the user-defined mapper to reducer to improve performance. The mapper only needs to send the original text in each row to the reducer. The reducer parses the record in each row and creates a key-value pair.</p> <p>NOTE When this parameter is set to org.apache.hadoop.hbase.mapreduce.TsvImporterByteMapper, this parameter is used only when the batch loading command without the <i>HBASE_CELL_VISIBILITY OR HBASE_CELL_TTL</i> option is executed. The org.apache.hadoop.hbase.mapreduce.TsvImporterByteMapper provides better performance.</p>	<p>org.apache.hadoop.hbase.mapreduce.TsvImporterByteMapper and org.apache.hadoop.hbase.mapreduce.TsvImporterTextMapper</p>

8.19.2 Improving Put Performance

Scenario

In the scenario where a large number of requests are continuously put, setting the following two parameters to **false** can greatly improve the Put performance.

- **hbase.regionserver.wal.durable.sync**
- **hbase.regionserver.hfile.durable.sync**

When the performance is improved, there is a low probability that data is lost if three DataNodes are faulty at the same time. Exercise caution when configuring the parameters in scenarios that have high requirements on data reliability.

 **NOTE**

This section applies to MRS 3.x and later versions.

Procedure

Navigation path for setting parameters:

On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations** > **All Configurations**. Enter the parameter name in the search box, and change the value.

Table 8-16 Parameters for improving put performance

Parameter	Description	Value
hbase.wal.hsync	Specifies whether to enable WAL file durability to make the WAL data persistence on disks. If this parameter is set to true , the performance is affected because each WAL file is synchronized to the disk by the Hadoop fsync.	false
hbase.hfile.hsync	Specifies whether to enable the HFile durability to make data persistence on disks. If this parameter is set to true, the performance is affected because each Hfile file is synchronized to the disk by the Hadoop fsync.	false

8.19.3 Optimizing Put and Scan Performance

Scenario

HBase has many configuration parameters related to read and write performance. The configuration parameters need to be adjusted based on the read/write request loads. This section describes how to optimize read and write performance by modifying the RegionServer configurations.

 **NOTE**

This section applies to MRS 3.x and later versions.

Procedure

- JVM GC parameters

Suggestions on setting the RegionServer **GC_OPTS** parameter:

- Set **-Xms** and **-Xmx** to the same value based on your needs. Increasing the memory can improve the read and write performance. For details, see the description of **hfile.block.cache.size** in [Table 8-18](#) and **hbase.regionserver.global.memstore.size** in [Table 8-17](#).
- Set **-XX:NewSize** and **-XX:MaxNewSize** to the same value. You are advised to set the value to **512M** in low-load scenarios and **2048M** in high-load scenarios.
- Set **X-XX:CMSInitiatingOccupancyFraction** to be less than and equal to 90, and it is calculated as follows: **100 x (hfile.block.cache.size + hbase.regionserver.global.memstore.size + 0.05)**.
- **-XX:MaxDirectMemorySize** indicates the non-heap memory used by the JVM. You are advised to set this parameter to **512M** in low-load scenarios and **2048M** in high-load scenarios.

NOTE

The **-XX:MaxDirectMemorySize** parameter is not used by default. If you need to set this parameter, add it to the **GC_OPTS** parameter.

- Put parameters

RegionServer processes the data of the put request and writes the data to memstore and HLog.

- When the size of memstore reaches the value of **hbase.hregion.memstore.flush.size**, memstore is updated to HDFS to generate HFiles.
- Compaction is triggered when the number of HFiles in the column cluster of the current region reaches the value of **hbase.hstore.compaction.min**.
- If the number of HFiles in the column cluster of the current region reaches the value of **hbase.hstore.blockingStoreFiles**, the operation of refreshing the memstore and generating HFiles is blocked. As a result, the put request is blocked.

Table 8-17 Put parameters

Parameter	Description	Default Value
hbase.wal.hsync	Indicates whether each WAL is persistent to disks. For details, see Improving Put Performance .	true
hbase.hfile.hsync	Indicates whether HFile write operations are persistent to disks. For details, see Improving Put Performance .	true

Parameter	Description	Default Value
hbase.hregion.memstore.flush.size	If the size of MemStore (unit: Byte) exceeds a specified value, MemStore is flushed to the corresponding disk. The value of this parameter is checked by each thread running hbase.server.thread.wakefrequency . It is recommended that you set this parameter to an integer multiple of the HDFS block size. You can increase the value if the memory is sufficient and the put load is heavy.	134217728
hbase.regionserver.global.memstore.size	Updates the size of all MemStores supported by the RegionServer before locking or forcible flush. It is recommended that you set this parameter to hbase.hregion.memstore.flush.size x Number of regions with active writes/ RegionServer GC -Xmx . The default value is 0.4 , indicating that 40% of RegionServer GC -Xmx is used.	0.4
hbase.hstore.flusher.count	Indicates the number of memstore flush threads. You can increase the parameter value in heavy-put-load scenarios.	2
hbase.regionserver.thread.compaction.small	Indicates the number of small compaction threads. You can increase the parameter value in heavy-put-load scenarios.	10

Parameter	Description	Default Value
hbase.hstore.blockingStoreFiles	If the number of HStoreFile files in a Store exceeds the specified value, the update of the HRegion will be locked until a compression is completed or the value of base.hstore.blockingWaitTime is exceeded. Each time MemStore is flushed, a StoreFile file is written into MemStore. Set this parameter to a larger value in heavy-put-load scenarios.	15

- Scan parameters

Table 8-18 Scan parameters

Parameter	Description	Default Value
hbase.client.scanner.timeout.period	Client and RegionServer parameters, indicating the lease timeout period of the client executing the scan operation. You are advised to set this parameter to an integer multiple of 60000 ms. You can set this parameter to a larger value when the read load is heavy. The unit is milliseconds.	60000
hfile.block.cache.size	Indicates the data cache percentage in the RegionServer GC -Xmx. You can increase the parameter value in heavy-read-load scenarios, in order to improve cache hit ratio and performance. It indicates the percentage of the maximum heap (-Xmx setting) allocated to the block cache of HFiles or StoreFiles.	When offheap is disabled, the default value is 0.25 . When offheap is enabled, the default value is 0.1 .

- Handler parameters

Table 8-19 Handler parameters

Parameter	Description	Default Value
hbase.regionserver.handler.count	Indicates the number of RPC server instances on RegionServer. The recommended value ranges from 200 to 400.	200
hbase.regionserver.metahandler.count	Indicates the number of program instances for processing prioritized requests. The recommended value ranges from 200 to 400.	200

8.19.4 Improving Real-time Data Write Performance

Scenario

Scenarios where data needs to be written to HBase in real time, or large-scale and consecutive put scenarios

 **NOTE**

This section applies to MRS 3.x and later versions.

Prerequisites

The HBase put or delete interface can be used to save data to HBase.

Procedure

- **Data writing server tuning**

Parameter portal:

Go to the **All Configurations** page of the HBase service. For details, see [Modifying Cluster Service Configuration Parameters](#).

Table 8-20 Configuration items that affect real-time data writing

Parameter	Description	Default Value
hbase.wal.hsync	Controls the synchronization degree when HLogs are written to the HDFS. If the value is true , HDFS returns only when data is written to the disk. If the value is false , HDFS returns when data is written to the OS cache. Set the parameter to false to improve write performance.	true
hbase.hfile.hsync	Controls the synchronization degree when HFiles are written to the HDFS. If the value is true , HDFS returns only when data is written to the disk. If the value is false , HDFS returns when data is written to the OS cache. Set the parameter to false to improve write performance.	true

Parameter	Description	Default Value
GC_OPTS	<p>You can increase HBase memory to improve HBase performance because read and write operations are performed in HBase memory. HeapSize and NewSize need to be adjusted. When you adjust HeapSize, set Xms and Xmx to the same value to avoid performance problems when JVM dynamically adjusts HeapSize. Set NewSize to 1/8 of HeapSize.</p> <ul style="list-style-type: none"> • HMaster: If HBase clusters enlarge and the number of Regions grows, properly increase the GC_OPTS parameter value of the HMaster. • RegionServer: A RegionServer needs more memory than an HMaster. If sufficient memory is available, increase the HeapSize value. <p>NOTE When the value of HeapSize for the active HMaster is 4 GB, the HBase cluster can support 100,000 regions. Empirically, each time 35,000 regions are added to the cluster, the value of HeapSize must be increased by 2 GB. It is recommended that the value of HeapSize for the active HMaster not exceed 32 GB.</p>	<ul style="list-style-type: none"> • HMaster -server - Xms4G - Xmx4G - XX:NewSize= 512M - XX:MaxNewSi ze=512M - XX:Metaspac eSize=128M - XX:MaxMetas paceSize=512 M - XX:+UseConc MarkSweepGC - XX:+CMSPara llelRemarkEn abled - XX:CMSInitiat ingOccupancy Fraction=65 - XX:+PrintGCD etails - Dsun.rmi.dgc. client.gcInter val=0x7FFFFFF FFFFFFFFFE - Dsun.rmi.dgc. server.gcInter val=0x7FFFFFF FFFFFFFFFE - XX:- OmitStackTra ceInFastThro w - XX:+PrintGCT imeStamps - XX:+PrintGCD ateStamps - XX:+UseGCLo gFileRotation - XX:NumberO fGCLogFiles= 10 - XX:GCLogFile Size=1M

Parameter	Description	Default Value
		<ul style="list-style-type: none"> • Region Server -server - Xms6G - Xmx6G - XX:NewSize=1024M - XX:MaxNewSize=1024M - XX:MetaspaceSize=128M - XX:MaxMetaspaceSize=512M - XX:+UseConcMarkSweepGC - XX:+CMSParallelRemarkEnabled - XX:CMSInitiatingOccupancyFraction=65 - XX:+PrintGCDetails - Dsun.rmi.dgc.client.gcInterval=0x7FFFFFFF - Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFF - XX:-OmitStackTraceInFastThrow - XX:+PrintGCTimeStamps - XX:+PrintGCDateStamps - XX:+UseGCLogFileRotation - XX:NumberOfGCLogFiles=10 - XX:GCLogFileSize=1M

Parameter	Description	Default Value
hbase.regionserver.handler.count	<p>Indicates the number of RPC server instances started on RegionServer. If the parameter is set to an excessively large value, threads will compete fiercely. If the parameter is set to an excessively small value, requests will be waiting for a long time in RegionServer, reducing the processing capability. You can add threads based on resources.</p> <p>It is recommended that the value be set to 100 to 300 based on the CPU usage.</p>	200
hbase.hregion.max.filesize	<p>Indicates the maximum size of an HStoreFile, in bytes. If the size of any HStoreFile exceeds the value of this parameter, the managed Hregion is divided into two parts.</p>	10737418240
hbase.hregion.memstore.flush.size	<p>On the RegionServer, when the size of memstore that exists in memory of write operations exceeds memstore.flush.size, MemStoreFlusher performs the Flush operation to write the memstore to the corresponding store in the format of HFile.</p> <p>If RegionServer memory is sufficient and active Regions are few, increase the parameter value and reduce compaction times to improve system performance.</p> <p>The Flush operation may be delayed after it takes place. Write operations continue and memstore keeps increasing during the delay. The maximum size of memstore is: memstore.flush.size x hbase.hregion.memstore.block.multiplier. When the memstore size exceeds the maximum value, write operations are blocked. Properly increasing the value of hbase.hregion.memstore.block.multiplier can reduce the blocks and make performance become more stable. Unit: byte</p>	134217728

Parameter	Description	Default Value
<p>hbase.regionserver.global.memstore.size</p>	<p>Updates the size of all MemStores supported by the RegionServer before locking or forcible flush. On the RegionServer, the MemStoreFlusher thread performs the flush. The thread regularly checks memory occupied by write operations. When the total memory volume occupied by write operations exceeds the threshold, MemStoreFlusher performs the flush. Larger memstore will be flushed first and then smaller ones until the occupied memory is less than the threshold.</p> <p>Threshold = hbase.regionserver.global.memstore.size x hbase.regionserver.global.memstore.size.lower.limit x HBase_HEAPSIZE</p> <p>NOTE The sum of the parameter value and the value of hfile.block.cache.size cannot exceed 0.8, that is, memory occupied by read and write operations cannot exceed 80% of HeapSize, ensuring stable running of other operations.</p>	<p>0.4</p>

Parameter	Description	Default Value
hbase.hstore.blockingStoreFiles	<p>Check whether the number of files is larger than the value of hbase.hstore.blockingStoreFiles before you flush regions.</p> <p>If it is larger than the value of hbase.hstore.blockingStoreFiles, perform a compaction and configure hbase.hstore.blockingWaitTime to 90s to make the flush delay for 90s. During the delay, write operations continue and the memstore size keeps increasing and exceeds the threshold (memstore.flush.size x hbase.hregion.memstore.block.multiplier), blocking write operations. After compaction is complete, a large number of writes may be generated. As a result, the performance fluctuates sharply.</p> <p>Increase the value of hbase.hstore.blockingStoreFiles to reduce block possibilities.</p>	15
hbase.regionserver.thread.compaction.throttle	<p>The compression whose size is greater than the value of this parameter is executed by the large thread pool. The unit is bytes. Indicates a threshold of a total file size for compaction during a Minor Compaction. The total file size affects execution duration of a compaction. If the total file size is large, other compactions or flushes may be blocked.</p>	1610612736
hbase.hstore.compaction.min	<p>Indicates the minimum number of HStoreFiles on which minor compaction is performed each time. When the size of a file in a Store exceeds the value of this parameter, the file is compacted. You can increase the value of this parameter to reduce the number of times that the file is compacted. If there are too many files in the Store, read performance will be affected.</p>	6

Parameter	Description	Default Value
hbase.hstore.compaction.max	Indicates the maximum number of HStoreFiles on which minor compaction is performed each time. The functions of the parameter and hbase.hstore.compaction.max.size are similar. Both are used to limit the execution duration of one compaction.	10
hbase.hstore.compaction.max.size	If the size of an HFile is larger than the parameter value, the HFile will not be compacted in a Minor Compaction but can be compacted in a Major Compaction. The parameter is used to prevent HFiles of large sizes from being compacted. After a Major Compaction is forbidden, multiple HFiles can exist in a Store and will not be merged into one HFile, without affecting data access performance. The unit is byte.	9223372036854775807
hbase.hregion.majorcompaction	Main compression interval of all HStoreFile files in a region. The unit is milliseconds. Execution of Major Compactions consumes much system resources and will affect system performance during peak hours. If service updates, deletion, and reclamation of expired data space are infrequent, set the parameter to 0 to disable Major Compactions. If you must perform a Major Compaction to reclaim more space, increase the parameter value and configure the hbase.offpeak.end.hour and hbase.offpeak.start.hour parameters to make the Major Compaction be triggered in off-peak hours.	604800000

Parameter	Description	Default Value
<ul style="list-style-type: none"> hbase.regionserver.maxlogs hbase.regionserver.hlog.blocksize 	<ul style="list-style-type: none"> Indicates the threshold for the number of HLog files that are not flushed on a RegionServer. If the number of HLog files is greater than the threshold, the RegionServer forcibly performs flush operations. Indicates the maximum size of an HLog file. If the size of an HLog file is greater than the value of this parameter, a new HLog file is generated. The old HLog file is disabled and archived. <p>The two parameters determine the number of HLogs that are not flushed in a RegionServer. When the data volume is less than the total size of memstore, the flush operation is forcibly triggered due to excessive HLog files. In this case, you can adjust the values of the two parameters to avoid forcible flush. Unit: byte</p>	<ul style="list-style-type: none"> 32 134217728

- Data writing client tuning**

It is recommended that data is written in Put List mode if necessary, which greatly improves write performance. The length of each put list needs to be set based on the single put size and parameters of the actual environment. You are advised to do some basic tests before configuring parameters.

- Data table writing design optimization**

Table 8-21 Parameters affecting real-time data writing

Parameter	Description	Default Value
COMPRESSION	<p>The compression algorithm compresses blocks in HFiles. For compressible data, configure the compression algorithm to efficiently reduce disk I/Os and improve performance.</p> <p>NOTE Some data cannot be efficiently compressed. For example, a compressed figure can hardly be compressed again. The common compression algorithm is SNAPPY, because it has a high encoding/decoding speed and acceptable compression rate.</p>	NONE

Parameter	Description	Default Value
BLOCKSIZE	Different block sizes affect HBase data read and write performance. You can configure sizes for blocks in an HFile. Larger blocks have a higher compression rate. However, they have poor performance in random data read, because HBase reads data in a unit of blocks. Set the parameter to 128 KB or 256 KB to improve data write efficiency without greatly affecting random read performance. The unit is byte.	65536
IN_MEMORY	Whether to cache table data in the memory first, which improves data read performance. If you will frequently access some small tables, set the parameter.	false

8.19.5 Improving Real-time Data Read Performance

Scenario

HBase data needs to be read.

Prerequisites

The get or scan interface of HBase has been invoked and data is read in real time from HBase.

Procedure

- **Data reading server tuning**

Parameter portal:

Go to the **All Configurations** page of the HBase service. For details, see [Modifying Cluster Service Configuration Parameters](#).

Table 8-22 Configuration items that affect real-time data reading

Parameter	Description	Default Value
GC_OPTS	<p>You can increase HBase memory to improve HBase performance because read and write operations are performed in HBase memory.</p> <p>HeapSize and NewSize need to be adjusted. When you adjust HeapSize, set Xms and Xmx to the same value to avoid performance problems when JVM dynamically adjusts HeapSize. Set NewSize to 1/8 of HeapSize.</p> <ul style="list-style-type: none"> • HMaster: If HBase clusters enlarge and the number of Regions grows, properly increase the GC_OPTS parameter value of the HMaster. • RegionServer: A RegionServer needs more memory than an HMaster. If sufficient memory is available, increase the HeapSize value. <p>NOTE When the value of HeapSize for the active HMaster is 4 GB, the HBase cluster can support 100,000 regions. Empirically, each time 35,000 regions are added to the cluster, the value of HeapSize must be increased by 2 GB. It is recommended that the value of HeapSize for the active HMaster not exceed 32 GB.</p>	<p>For versions earlier than MRS 3.x:</p> <ul style="list-style-type: none"> • HMaster: <ul style="list-style-type: none"> -server - Xms2G - Xmx2G - XX:NewSize=256M - XX:MaxNewSize=256M - XX:MetaspaceSize=128M - XX:MaxMetaspaceSize=512M - XX:MaxDirectMemorySize=512M - XX:+UseConcMarkSweepGC - XX:+CMSParallelRemarkEnabled - XX:CMSInitiatingOccupancyFraction=65 - XX:+PrintGCDetails - Dsun.rmi.dgc.client.gcninterval=0x7FFFFFFF - Dsun.rmi.dgc.server.gcninterval=0x7FFFFFFF - XX:-OmitStackTraceInFastThread - XX:+PrintGCTimeStamps

Parameter	Description	Default Value
		<p>- XX:+PrintGC DateStamps - XX:+UseGCL ogFileRotati on - XX:Number OfGCLogFil es=10 - XX:GCLogFil eSize=1M</p> <ul style="list-style-type: none"> ● RegionServe r: -server - Xms4G - Xmx4G - XX:NewSize =512M - XX:MaxNew Size=512M - XX:Metaspa ceSize=128 M - XX:MaxMet aspaceSize= 512M - XX:MaxDire ctMemorySi ze=512M - XX:+UseCon cMarkSwee pGC - XX:+CMSPar allelRemark Enabled - XX:CMSIniti atingOccup ancyFractio n=65 - XX:+PrintGC Details - Dsun.rmi.dg c.client.gcln terval=0x7F FFFFFFFF FFE - Dsun.rmi.dg c.server.gcln

Parameter	Description	Default Value
		<p> terval=0x7F FFFFFFFF FFE -XX:- OmitStackTr aceInFastTh row - XX:+PrintGC TimeStamps - XX:+PrintGC DateStamps - XX:+UseGCL ogFileRotati on - XX:Number OfGCLogFil es=10 - XX:GCLogFil eSize=1M </p> <p>For MRS 3.x or later:</p> <ul style="list-style-type: none"> HMaster -server - Xms4G - Xmx4G - XX:NewSize =512M - XX:MaxNew Size=512M - XX:Metaspa ceSize=128 M - XX:MaxMet aspaceSize= 512M - XX:+UseCon cMarkSwee pGC - XX:+CMSPar allelRemark Enabled - XX:CMSIniti atingOccup ancyFractio n=65 - XX:+PrintGC Details -

Parameter	Description	Default Value
		<p>Dsun.rmi.dgc.client.gclnterval=0x7FFFFFFF - Dsun.rmi.dgc.server.gclnterval=0x7FFFFFFF -XX:-OmitStackTraceInFastThrow -XX:+PrintGCTimeStamps -XX:+PrintGCDateStamps -XX:+UseGcLogFileRotation -XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=1M</p> <ul style="list-style-type: none"> Region Server <ul style="list-style-type: none"> -server -Xms6G -Xmx6G -XX:NewSize=1024M -XX:MaxNewSize=1024M -XX:MetaspaceSize=128M -XX:MaxMetaspaceSize=512M -XX:+UseConcMarkSweepGC -XX:+CMSParallelRemarkEnabled -XX:CMSInit

Parameter	Description	Default Value
		atingOccupancyFraction=65 - XX:+PrintGCDetails - Dsun.rmi.dgc.client.gcInterval=0x7FFFFFFF - Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFF -XX:-OmitStackTracesInFastThrow - XX:+PrintGCTimeStamps - XX:+PrintGCDateStamps - XX:+UseGCLogFileRotation - XX:NumberOfGCLogFiles=10 - XX:GCLogFileSize=1M
hbase.regionserver.handler.count	Indicates the number of requests that RegionServer can process concurrently. If the parameter is set to an excessively large value, threads will compete fiercely. If the parameter is set to an excessively small value, requests will be waiting for a long time in RegionServer, reducing the processing capability. You can add threads based on resources. It is recommended that the value be set to 100 to 300 based on the CPU usage.	200

Parameter	Description	Default Value
hfile.block.cache.size	HBase cache sizes affect query efficiency. Set cache sizes based on query modes and query record distribution. If random query is used to reduce the hit ratio of the buffer, you can reduce the buffer size.	When offheap is disabled, the default value is 0.25 . When offheap is enabled, the default value is 0.1 .

 **NOTE**

If read and write operations are performed at the same time, the performance of the two operations affects each other. If flush and compaction operations are frequently performed due to data writes, a large number of disk I/O operations are occupied, affecting read performance. If a large number of compaction operations are blocked due to write operations, multiple HFiles exist in the region, affecting read performance. Therefore, if the read performance is unsatisfactory, you need to check whether the write configurations are proper.

- **Data reading client tuning**

When scanning data, you need to set **caching** (the number of records read from the server at a time. The default value is **1**.). If the default value is used, the read performance will be extremely low.

If you do not need to read all columns of a piece of data, specify the columns to be read to reduce network I/O.

If you only need to read the row key, add a filter (FirstKeyOnlyFilter or KeyOnlyFilter) that only reads the row key.

- **Data table reading design optimization**

Table 8-23 Parameters affecting real-time data reading

Parameter	Description	Default Value
COMPRESSION	The compression algorithm compresses blocks in HFiles. For compressible data, configure the compression algorithm to efficiently reduce disk I/Os and improve performance. NOTE Some data cannot be efficiently compressed. For example, a compressed figure can hardly be compressed again. The common compression algorithm is SNAPPY, because it has a high encoding/decoding speed and acceptable compression rate.	NONE

Parameter	Description	Default Value
BLOCKSIZE	Different block sizes affect HBase data read and write performance. You can configure sizes for blocks in an HFile. Larger blocks have a higher compression rate. However, they have poor performance in random data read, because HBase reads data in a unit of blocks. Set the parameter to 128 KB or 256 KB to improve data write efficiency without greatly affecting random read performance. The unit is byte.	65536
DATA_BLOCK_ENCODING	Encoding method of the block in an HFile. If a row contains multiple columns, set FAST_DIFF to save data storage space and improve performance.	NONE

8.19.6 Optimizing JVM Parameters

Scenario

When the number of clusters reaches a certain scale, the default settings of the Java virtual machine (JVM) cannot meet the cluster requirements. In this case, the cluster performance deteriorates or the clusters may be unavailable. Therefore, JVM parameters must be properly configured based on actual service conditions to improve the cluster performance.

Procedure

Navigation path for setting parameters:

The JVM parameters related to the HBase role must be configured in the **hbase-env.sh** file in the `/${BIGDATA_HOME}/FusionInsight_HD_*/install/FusionInsight-HBase-2.2.3/hbase/conf/` directory of the node where the HBase service is installed.

Each role has JVM parameter configuration variables, as shown in [Table 8-24](#).

Table 8-24 HBase-related JVM parameter configuration variables

Variable	Affected Role
HBASE_OPTS	All roles of HBase
SERVER_GC_OPTS	All roles on the HBase server, such as Master and RegionServer
CLIENT_GC_OPTS	Client process of HBase

Variable	Affected Role
HBASE_MASTER_OPTS	Master of HBase
HBASE_REGIONSERVER_OPTS	RegionServer of HBase
HBASE_THRIFT_OPTS	Thrift of HBase

Configuration example:

```
export HADOOP_NAMENODE_OPTS="-Dhadoop.security.logger=${HADOOP_SECURITY_LOGGER:-INFO,RFAS} -Dhdfs.audit.logger=${HDFS_AUDIT_LOGGER:-INFO,NullAppender} $HADOOP_NAMENODE_OPTS"
```

8.20 Common Issues About HBase

8.20.1 Why Does a Client Keep Failing to Connect to a Server for a Long Time?

Question

A HBase server is faulty and cannot provide services. In this case, when a table operation is performed on the HBase client, why is the operation suspended and no response is received for a long time?

Answer

Problem Analysis

When the HBase server malfunctions, the table operation request from the HBase client is tried for several times and times out. The default timeout value is **Integer.MAX_VALUE (2147483647 ms)**. The table operation request is retired constantly during such a long period of time and is suspended at last.

Solution

The HBase client provides two configuration items to configure the retry and timeout of the client. [Table 8-25](#) describes them.


Set the following parameters in the *Client installation path*/HBase/hbase/conf/**hbase-site.xml** configuration file:

Table 8-25 Configuration parameters of retry and timeout

Parameter	Description	Default Value
hbase.client.operation.timeout	Client operation timeout period You need to manually add the information to the configuration file.	2147483647 ms
hbase.client.retries.number	Maximum retry times supported by all retryable operations.	35

Figure 8-3 describes the working principles of retry and timeout.

Figure 8-3 Process for HBase client operation retry timeout



The process indicates that a suspension occurs if the preceding parameters are not configured based on site requirements. It is recommended that a proper timeout period be set based on scenarios. If the operation takes a long time, set a long timeout period. If the operation takes a short time, set a short timeout period. The number of retries can be set to **(hbase.client.retries.number)*60*1000(ms)**. The timeout period can be slightly greater than **hbase.client.operation.timeout**.

8.20.2 Operation Failures Occur in Stopping BulkLoad On the Client

Question

Why submitted operations fail by stopping BulkLoad on the client during BulkLoad data importing?

Answer

When BulkLoad is enabled on the client, a partitioner file is generated and used to demarcate the range of Map task data inputting. The file is automatically deleted when BulkLoad exists on the client. In general, if all map tasks are enabled and running, the termination of BulkLoad on the client does not cause the failure of submitted operations. However, due to the retry and speculative execution mechanism of Map tasks, a Map task is performed again if failures of the Reduce task to download the data of the completed Map task exceed the limit. In this case, if BulkLoad already exists on the client, the retry Map task fails and the operation failure occurs because the partitioner file is missing. Therefore, it is recommended not to stop BulkLoad on the client during BulkLoad data importing.

8.20.3 Why May a Table Creation Exception Occur When HBase Deletes or Creates the Same Table Consecutively?

Question

When HBase consecutively deletes and creates the same table, why may a table creation exception occur?

Answer

Execution process: Disable Table > Drop Table > Create Table > Disable Table > Drop Table > And more

1. When a table is disabled, HMaster sends an RPC request to RegionServer, and RegionServer brings the region offline. When the time required for closing a region on RegionServer exceeds the timeout period for HBase HMaster to wait for the region to enter the RIT state, HMaster considers that the region is offline by default. Actually, the region may be in the flush memstore phase.
2. After an RPC request is sent to close a region, HMaster checks whether all regions in the table are offline. If the closure times out, HMaster considers that the regions are offline and returns a message indicating that the regions are successfully closed.
3. After the closure is successful, the data directory corresponding to the HBase table is deleted.
4. After the table is deleted, the data directory is recreated by the region that is still in the flush memstore phase.
5. When the table is created again, the **temp** directory is copied to the HBase data directory. However, the HBase data directory is not empty. As a result, when the HDFS rename API is called, the data directory changes to the last

layer of the **temp** directory and is appended to the HBase data directory, for example, **\$rootDir/data/\$nameSpace/\$tableName/\$tableName**. In this case, the table fails to be created.

Troubleshooting Method

When this problem occurs, check whether the HBase data directory corresponding to the table exists. If it exists, rename the directory.

The HBase data directory consists of **\$rootDir/data/\$nameSpace/\$tableName**, for example, **hdfs://hacluster/hbase/data/default/TestTable**. **\$rootDir** is the HBase root directory, which can be obtained by configuring **hbase.rootdir.perms** in **hbase-site.xml**. The **data** directory is a fixed directory of HBase. **\$nameSpace** indicates the nameSpace name. **\$tableName** indicates the table name.

8.20.4 Why Other Services Become Unstable If HBase Sets up A Large Number of Connections over the Network Port?

Question

Why other services become unstable if HBase sets up a large number of connections over the network port?

Answer

When the OS command **lsof** or **netstat** is run, it is found that many TCP connections are in the CLOSE_WAIT state and the owner of the connections is HBase RegionServer. This can cause exhaustion of network ports or limit exceeding of HDFS connections, resulting in instability of other services. The HBase CLOSE_WAIT phenomenon is the HBase mechanism.

The reason why HBase CLOSE_WAIT occurs is as follows: HBase data is stored in the HDFS as HFile, which can be called StoreFiles. HBase functions as the client of the HDFS. When HBase creates a StoreFile or starts loading a StoreFile, it creates an HDFS connection. When the StoreFile is created or loaded successfully, the HDFS considers that the task is completed and transfers the connection close permission to HBase. However, HBase may choose not to close the connection to ensure real-time response; that is, HBase may maintain the connection so that it can quickly access the corresponding data file upon request. In this case, the connection is in the CLOSE_WAIT, which indicates that the connection needs to be closed by the client.

When a StoreFile will be created: HBase executes the Flush operation.

When Flush is executed: The data written by HBase is first stored in memstore. The Flush operation is performed only when the usage of memstore reaches the threshold or the **flush** command is run to write data into the HDFS.

To resolve the issue, use either of the following methods:

Because of the HBase connection mechanism, the number of StoreFiles must be restricted to reduce the occupation of HBase ports. This can be achieved by triggering HBase's the compaction action, that is, HBase file merging.

Method 1: On HBase shell client, run **major_compact**.

Method 2: Compile HBase client code to invoke the compact method of the HBaseAdmin class to trigger HBase's compaction action.

If the HBase port occupation issue cannot be resolved through compact, it indicates that the HBase usage has reached the bottleneck. In such a case, you are advised to perform the following:

- Check whether the initial number of Regions configured in the table is appropriate.
- Check whether useless data exists.

If useless data exists, delete the data to reduce the number of storage files for the HBase. If the preceding conditions are not met, then you need to consider a capacity expansion.

8.20.5 Why Does the HBase BulkLoad Task (One Table Has 26 TB Data) Consisting of 210,000 Map Tasks and 10,000 Reduce Tasks Fail?

Question

The HBase bulkLoad task (a single table contains 26 TB data) has 210,000 maps and 10,000 reduce tasks (in MRS 3.x or later), and the task fails.

Answer

ZooKeeper I/O bottleneck observation methods:

1. On the monitoring page of Manager, check whether the number of ZooKeeper requests on a single node exceeds the upper limit.
2. View ZooKeeper and HBase logs to check whether a large number of I/O Exception Timeout or SocketTimeout Exception exceptions occur.

Optimization suggestions:

1. Change the number of ZooKeeper instances to 5 or more. You are advised to set **peerType** to **observer** to increase the number of observers.
2. Control the number of concurrent maps of a single task or reduce the memory for running tasks on each node to lighten the node load.
3. Upgrade ZooKeeper data disks, such as SSDs.

8.20.6 How Do I Restore a Region in the RIT State for a Long Time?

Question

How do I restore a region in the RIT state for a long time?

Answer

Log in to the HMaster Web UI, choose **Procedure & Locks** in the navigation tree, and check whether any process ID is in the **Waiting** state. 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
```

8.20.7 Why Does HMaster Exits Due to Timeout When Waiting for the Namespace Table to Go Online?

Question

Why does HMaster exit due to timeout when waiting for the namespace table to go online?

Answer

During the HMaster active/standby switchover or startup, HMaster performs WAL splitting and region recovery for the RegionServer that failed or was stopped previously.

Multiple threads are running in the background to monitor the HMaster startup process.

- **TableNamespaceManager**
This is a help class, which is used to manage the allocation of namespace tables and monitoring table regions during HMaster active/standby switchover or startup. If the namespace table is not online within the specified time (**hbase.master.namespace.init.timeout**, which is 3,600,000 ms by default), the thread terminates HMaster abnormally.
- **InitializationMonitor**
This is an initialization thread monitoring class of the primary HMaster, which is used to monitor the initialization of the primary HMaster. If a thread fails to be initialized within the specified time (**hbase.master.initializationmonitor.timeout**, which is 3,600,000 ms by default), the thread terminates HMaster abnormally. If **hbase.master.initializationmonitor.haltontimeout** is started, the default value is **false**.

During the HMaster active/standby switchover or startup, if the **WAL hlog** file exists, the WAL splitting task is initialized. If the WAL hlog splitting task is complete, it initializes the table region allocation task.

HMaster uses ZooKeeper to coordinate log splitting tasks and valid RegionServers and track task development. If the primary HMaster exits during the log splitting

task, the new primary HMaster attempts to resend the unfinished task, and RegionServer starts the log splitting task from the beginning.

The initialization of the HMaster is delayed due to the following reasons:

- Network faults occur intermittently.
- Disks run into bottlenecks.
- The log splitting task is overloaded, and RegionServer runs slowly.
- RegionServer (region opening) responds slowly.

In the preceding scenarios, you are advised to add the following configuration parameters to enable HMaster to complete the restoration task earlier. Otherwise, the Master will exit, causing a longer delay of the entire restoration process.

- Increase the online waiting timeout period of the namespace table to ensure that the Master has enough time to coordinate the splitting tasks of the RegionServer worker and avoid repeated tasks.

hbase.master.namespace.init.timeout (default value: 3,600,000 ms)

- Increase the number of concurrent splitting tasks through RegionServer worker to ensure that RegionServer worker can process splitting tasks in parallel (RegionServers need more cores). Add the following parameters to *Client installation path /HBase/hbase/conf/hbase-site.xml*:

hbase.regionserver.wal.max Splitters (default value: 2)

- If all restoration processes require time, increase the timeout period for initializing the monitoring thread.

hbase.master.initializationmonitor.timeout (default value: 3,600,000 ms)

8.20.8 Why Does SocketTimeoutException Occur When a Client Queries HBase?

Question

Why does the following exception occur on the client when I use the HBase client to operate table data?

```
2015-12-15 02:41:14,054 | WARN | [task-result-getter-2] | Lost task 2.0 in stage 58.0 (TID 3288, linux-175): org.apache.hadoop.hbase.client.RetriesExhaustedException: Failed after attempts=36, exceptions: Tue Dec 15 02:41:14 CST 2015, null, java.net.SocketTimeoutException: callTimeout=60000, callDuration=60303: row 'xxxxxx' on table 'xxxxxx',\x05\x1E \x80\x00\x00\x00\x80\x00\x00\x00\x00\x00\x00\x00\x80\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00, 1449912620868.6a6b7d0c272803d8186930a3bafb10a9., hostname=xxxxxx,16020,1449941841479, seqNum=5 at org.apache.hadoop.hbase.client.RpcRetryingCallerWithReadReplicas.throwEnrichedException(RpcRetryingCallerWithReadReplicas.java:275) at org.apache.hadoop.hbase.client.ScannerCallableWithReplicas.call(ScannerCallableWithReplicas.java:223) at org.apache.hadoop.hbase.client.ScannerCallableWithReplicas.call(ScannerCallableWithReplicas.java:61) at org.apache.hadoop.hbase.client.RpcRetryingCaller.callWithoutRetries(RpcRetryingCaller.java:200) at org.apache.hadoop.hbase.client.ClientScanner.call(ClientScanner.java:323)
```

At the same time, the following log is displayed on RegionServer:

```
2015-12-15 02:45:44,551 | WARN | PriorityRpcServer.handler=7,queue=1,port=16020 | (responseTooSlow): {"call": "Scan(org.apache.hadoop.hbase.protobuf.generated.ClientProtos$ScanRequest)", "starttimems": 1450118730780, "responsesize": 416, "method": "Scan", "processingtimems": 13770, "client": "10.9
```

```
1.8.175:41182","queuetimems":0,"class":"HRegionServer"} |
org.apache.hadoop.hbase.ipc.RpcServer.logResponse(RpcServer.java:2221)
2015-12-15 02:45:57,722 | WARN | PriorityRpcServer.handler=3,queue=1,port=16020 | (responseTooSlow):
{"call":"Scan(org.apache.hadoop.hbase.protobuf.generated.ClientProtos
$ScanRequest)","starttimems":1450118746297,"responsesize":416,
"method":"Scan","processingtimems":11425,"client":"10.91.8.175:41182","queuetimems":1746,"class":"HRegi
onServer"} | org.apache.hadoop.hbase.ipc.RpcServer.logResponse(RpcServer.java:2221)
2015-12-15 02:47:21,668 | INFO | LruBlockCacheStatsExecutor | totalSize=7.54 GB, freeSize=369.52 MB,
max=7.90 GB, blockCount=406107,
accesses=35400006, hits=16803205, hitRatio=47.47%, , cachingAccesses=31864266, cachingHits=14806045,
cachingHitsRatio=46.47%,
evictions=17654, evicted=16642283, evictedPerRun=942.69189453125 |
org.apache.hadoop.hbase.io.hfile.LruBlockCache.logStats(LruBlockCache.java:858)
2015-12-15 02:52:21,668 | INFO | LruBlockCacheStatsExecutor | totalSize=7.51 GB, freeSize=395.34 MB,
max=7.90 GB, blockCount=403080,
accesses=35685793, hits=16933684, hitRatio=47.45%, , cachingAccesses=32150053, cachingHits=14936524,
cachingHitsRatio=46.46%,
evictions=17684, evicted=16800617, evictedPerRun=950.046142578125 |
org.apache.hadoop.hbase.io.hfile.LruBlockCache.logStats(LruBlockCache.java:858)
```

Answer

The memory allocated to RegionServer is too small and the number of Regions is too large. As a result, the memory is insufficient during the running, and the server responds slowly to the client. Modify the following memory allocation parameters in the **hbase-site.xml** configuration file of RegionServer:

Table 8-26 RegionServer memory allocation parameters

Parameter	Description	Default Value
GC_OPTS	Initial memory and maximum memory allocated to RegionServer in startup parameters.	-Xms8G -Xmx8G
hfile.block.cache.size	Percentage of the maximum heap (-Xmx setting) allocated to the block cache of HFiles or StoreFiles.	When offheap is disabled, the default value is 0.25 . When offheap is enabled, the default value is 0.1 .

8.20.9 Why Modified and Deleted Data Can Still Be Queried by Using the Scan Command?

Question

Why modified and deleted data can still be queried by using the **scan** command?

```
scan '<table_name>',{FILTER=>"SingleColumnValueFilter('<column_family>','column',=,'binary:<value>')"} }
```

Answer

Because of the scalability of HBase, all values specific to the versions in the queried column are all matched by default, even if the values have been modified

or deleted. For a row where column matching has failed (that is, the column does not exist in the row), the HBase also queries the row.

If you want to query only the new values and rows where column matching is successful, you can use the following statement:

```
scan '<table_name>',  
{FILTER=>"SingleColumnValueFilter('<column_family>','column','=', 'binary:<value>',true,true)"}
```

This command can filter all rows where column query has failed. It queries only the latest values of the current data in the table; that is, it does not query the values before modification or the deleted values.

NOTE

The related parameters of **SingleColumnValueFilter** are described as follows:

SingleColumnValueFilter(final byte[] family, final byte[] qualifier, final CompareOp compareOp, ByteArrayComparable comparator, final boolean filterIfMissing, final boolean latestVersionOnly)

Parameter description:

- family: family of the column to be queried.
- qualifier: column to be queried.
- compareOp: comparison operation, such as = and >.
- comparator: target value to be queried.
- filterIfMissing: whether a row is filtered out if the queried column does not exist. The default value is false.
- latestVersionOnly: whether values of the latest version are queried. The default value is false.

8.20.10 Why "java.lang.UnsatisfiedLinkError: Permission denied" exception thrown while starting HBase shell?

Question

Why "java.lang.UnsatisfiedLinkError: Permission denied" exception thrown while starting HBase shell?

Answer

During HBase shell execution JRuby create temporary files under **java.io.tmpdir** path and default value of **java.io.tmpdir** is **/tmp**. If NOEXEC permission is set to /tmp directory then HBase shell start will fail with "java.lang.UnsatisfiedLinkError: Permission denied" exception.

So "java.io.tmpdir" must be set to a different path in HBASE_OPTS/CLIENT_GC_OPTS if NOEXEC is set to /tmp directory.

8.20.11 When does the RegionServers listed under "Dead Region Servers" on HMaster WebUI gets cleared?

Question

When does the RegionServers listed under "Dead Region Servers" on HMaster WebUI gets cleared?

Answer

When an online RegionServer goes down abruptly, it is displayed under "Dead Region Servers" in the HMaster WebUI. When dead RegionServer restarts and reports back to HMaster successfully, the "Dead Region Servers" in the HMaster WebUI gets cleared.

The "Dead Region Servers" is also gets cleared, when the HMaster failover operation is performed successfully.

In cases when an Active HMaster hosting some regions is abruptly killed, Backup HMaster will become the new Active HMater and displays previous Active HMaster as dead RegionServer.

8.20.12 Why Are Different Query Results Returned After I Use Same Query Criteria to Query Data Successfully Imported by HBase bulkload?

Question

If the data to be imported by HBase bulkload has identical rowkeys, the data import is successful but identical query criteria produce different query results.

Answer

Data with an identical rowkey is loaded into HBase in the order in which data is read. The data with the latest timestamp is considered to be the latest data. By default, data is not queried by timestamp. Therefore, if you query for data with an identical rowkey, only the latest data is returned.

While data is being loaded by bulkload, the memory processes the data into HFiles quickly, leading to the possibility that data with an identical rowkey has a same timestamp. In this case, identical query criteria may produce different query results.

To avoid this problem, ensure that the same data file does not contain identical rowkeys while you are creating tables or loading data.

8.20.13 What Should I Do If I Fail to Create Tables Due to the FAILED_OPEN State of Regions?

Question

What should I do if I fail to create tables due to the FAILED_OPEN state of Regions?

Answer

If a network, HDFS, or Active HMaster fault occurs during the creation of tables, some Regions may fail to go online and therefore enter the FAILED_OPEN state. In this case, tables fail to be created.

The tables that fail to be created due to the preceding mentioned issue cannot be repaired. To solve this problem, perform the following operations to delete and re-create the tables:

1. Run the following command on the cluster client to repair the state of the tables:
hbase hbck -fixTableStates
2. Enter the HBase shell and run the following commands to delete the tables that fail to be created:
***truncate* '<table_name>'**
***disable* '<table_name>'**
***drop* '<table_name>'**
3. Create the tables using the recreation command.

8.20.14 How Do I Delete Residual Table Names in the /hbase/table-lock Directory of ZooKeeper?

Question

In security mode, names of tables that failed to be created are unnecessarily retained in the table-lock node (default directory is /hbase/table-lock) of ZooKeeper. How do I delete these residual table names?

Answer

Perform the following steps:

1. On the client, run the `kinit` command as the `hbase` user to obtain a security certificate.
2. Run the ***hbase zkcli*** command to launch the ZooKeeper Command Line Interface (zkCLI).
3. Run the ***ls /hbase/table*** command on the zkCLI to check whether the table name of the table that fails to be created exists.
 - If the table name exists, no further operation is required.
 - If the table name does not exist, run ***ls /hbase/table-lock*** to check whether the table name of the table fail to be created exist. If the table name exists, run the ***delete /hbase/table-lock/<table>*** command to delete the table name. In the ***delete /hbase/table-lock/<table>*** command, ***<table>*** indicates the residual table name.

8.20.15 Why Does HBase Become Faulty When I Set a Quota for the Directory Used by HBase in HDFS?

Question

Why does HBase become faulty when I set quota for the directory used by HBase in HDFS?

Answer

The flush operation of a table is to write memstore data to HDFS.

If the HDFS directory does not have sufficient disk space quota, the flush operation will fail and the region server will stop.

```
Caused by: org.apache.hadoop.hdfs.protocol.DSQuotaExceededException: The DiskSpace quota of /hbase/
data/<namespace>/<tableName> is exceeded: quota = 1024 B = 1 KB but diskspace consumed = 402655638
B = 384.00 MB
?at
org.apache.hadoop.hdfs.server.namenode.DirectoryWithQuotaFeature.verifyStoragespaceQuota(DirectoryWit
hQuotaFeature.java:211)
?at
org.apache.hadoop.hdfs.server.namenode.DirectoryWithQuotaFeature.verifyQuota(DirectoryWithQuotaFeatu
re.java:239)
?at org.apache.hadoop.hdfs.server.namenode.FSDirectory.verifyQuota(FSDirectory.java:882)
?at org.apache.hadoop.hdfs.server.namenode.FSDirectory.updateCount(FSDirectory.java:711)
?at org.apache.hadoop.hdfs.server.namenode.FSDirectory.updateCount(FSDirectory.java:670)
?at org.apache.hadoop.hdfs.server.namenode.FSDirectory.addBlock(FSDirectory.java:495)
```

In the preceding exception, the disk space quota of the **/hbase/data/<namespace>/<tableName>** table is 1 KB, but the memstore data is 384.00 MB. Therefore, the flush operation fails and the region server stops.

When the region server is terminated, HMaster replays the WAL file of the terminated region server to restore data. The disk space quota is limited. As a result, the replay operation of the WAL file fails, and the HMaster process exits unexpectedly.

```
2016-07-28 19:11:40,352 | FATAL | MASTER_SERVER_OPERATIONS-10-91-9-131:16000-0 | Caught throwable
while processing event M_SERVER_SHUTDOWN |
org.apache.hadoop.hbase.master.HMaster.abort(HMaster.java:2474)
java.io.IOException: failed log splitting for 10-91-9-131,16020,1469689987884, will retry
?at
org.apache.hadoop.hbase.master.handler.ServerShutdownHandler.resubmit(ServerShutdownHandler.java:365
)
?at
org.apache.hadoop.hbase.master.handler.ServerShutdownHandler.process(ServerShutdownHandler.java:220)
?at org.apache.hadoop.hbase.executor.EventHandler.run(EventHandler.java:129)
?at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
?at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
?at java.lang.Thread.run(Thread.java:745)
Caused by: java.io.IOException: error or interrupted while splitting logs in [hdfs://hacluster/hbase/WALs/<RS-
Hostname>,<RS-Port>,<startcode>-splitting] Task = installed = 6 done = 3 error = 3
?at org.apache.hadoop.hbase.master.SplitLogManager.splitLogDistributed(SplitLogManager.java:290)
?at org.apache.hadoop.hbase.master.MasterFileSystem.splitLog(MasterFileSystem.java:402)
?at org.apache.hadoop.hbase.master.MasterFileSystem.splitLog(MasterFileSystem.java:375)
```

Therefore, you cannot set the quota value for the HBase directory in HDFS. If the exception occurs, perform the following operations:

- Step 1** Run the **kinit Username** command on the client to enable the HBase user to obtain security authentication.
- Step 2** Run the **hdfs dfs -count -q /hbase/data/<namespace>/<tableName>** command to check the allocated disk space quota.
- Step 3** Run the following command to cancel the quota limit and restore HBase:


```
hdfs dfsadmin -clrSpaceQuota /hbase/data/<namespace>/<tableName>
```

----End

8.20.16 Why HMaster Times Out While Waiting for Namespace Table to be Assigned After Rebuilding Meta Using OfflineMetaRepair Tool and Startups Failed

Question

Why HMaster times out while waiting for namespace table to be assigned after rebuilding meta using OfflineMetaRepair tool and startups failed?

HMaster abort with following FATAL message,

```
2017-06-15 15:11:07,582 FATAL [Hostname:16000.activeMasterManager] master.HMaster: Unhandled  
exception. Starting shutdown.  
java.io.IOException: Timedout 120000ms waiting for namespace table to be assigned  
    at org.apache.hadoop.hbase.master.TableNamespaceManager.start(TableNamespaceManager.java:98)  
    at org.apache.hadoop.hbase.master.HMaster.initNamespace(HMaster.java:1054)  
    at org.apache.hadoop.hbase.master.HMaster.finishActiveMasterInitialization(HMaster.java:848)  
    at org.apache.hadoop.hbase.master.HMaster.access$600(HMaster.java:199)  
    at org.apache.hadoop.hbase.master.HMaster$2.run(HMaster.java:1871)  
    at java.lang.Thread.run(Thread.java:745)
```

Answer

When meta is rebuilt by OfflineMetaRepair tool then HMaster wait for all region server's WAL split during start up to avoid the data inconsistency problem. HMaster trigger user regions assignment once WAL split completes. So when the cluster is in the unusual scenario, there are chances WAL splitting may take long time which depends on multiple factors like too many WALs, slow I/O, region servers are not stable etc.

HMaster should be able to finish all region server WAL splitting successfully. Perform the following steps.

1. Make sure cluster is stable, no other problem exist. If any problem occurs, please correct them first.
2. Configure a large value to **hbase.master.initializationmonitor.timeout** parameters, default value is **3600000** milliseconds.
3. Restart HBase service.

8.20.17 Why Messages Containing FileNotFoundException and no lease Are Frequently Displayed in the HMaster Logs During the WAL Splitting Process?

Question

Why messages containing FileNotFoundException and no lease are frequently displayed in the HMaster logs during the WAL splitting process?

```
2017-06-10 09:50:27,586 | ERROR | split-log-closeStream-2 | Couldn't close log at hdfs://hacluster/hbase/  
data/default/largeT1/2b48346d087275fe751fc049334fda93/recovered.edits/00000000000000000000000000000000.temp |  
org.apache.hadoop.hbase.wal.WALSplitter$LogRecoveredEditsOutputSink$2.call(WALSplitter.java:1330)  
java.io.FileNotFoundException: No lease on /hbase/data/default/  
largeT1/2b48346d087275fe751fc049334fda93/recovered.edits/00000000000000000000000000000000.temp (inode  
1092653): File does not exist. [Lease: Holder: DFSCClient_NONMAPREDUCE_1202985678_1, pendingcreates:  
1936]
```

```
?at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.checkLease(FSNamesystem.java:3432)
?at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.analyzeFileState(FSNamesystem.java:3223)
?at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getNewBlockTargets(FSNamesystem.java:3057)
?at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getAdditionalBlock(FSNamesystem.java:3011)
?at org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.addBlock(NameNodeRpcServer.java:842)
?at
org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolServerSideTranslatorPB.addBlock(ClientNamenodeProtocolServerSideTranslatorPB.java:526)
?at org.apache.hadoop.hdfs.protocol.proto.ClientNamenodeProtocolProtos$ClientNamenodeProtocol$2.callBlockingMethod(ClientNamenodeProtocolProtos.java)
?at org.apache.hadoop.ipc.ProtobufRpcEngine$Server$ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:616)
?at org.apache.hadoop.ipc.RPC$Server.call(RPC.java:973)
?at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2260)
?at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2256)
?at java.security.AccessController.doPrivileged(Native Method)
?at javax.security.auth.Subject.doAs(Subject.java:422)
?at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1769)
?at org.apache.hadoop.ipc.Server$Handler.run(Server.java:2254)

?at sun.reflect.GeneratedConstructorAccessor40.newInstance(Unknown Source)
?at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:45)
?at java.lang.reflect.Constructor.newInstance(Constructor.java:423)
?at org.apache.hadoop.ipc.RemoteException.instantiateException(RemoteException.java:106)
?at org.apache.hadoop.ipc.RemoteException.unwrapRemoteException(RemoteException.java:73)
?at org.apache.hadoop.hdfs.DataStreamer.locateFollowingBlock(DataStreamer.java:1842)
?at org.apache.hadoop.hdfs.DataStreamer.nextBlockOutputStream(DataStreamer.java:1639)
?at org.apache.hadoop.hdfs.DataStreamer.run(DataStreamer.java:665)
```

Answer

During the WAL splitting process, the WAL splitting timeout period is specified by the **hbase.splitlog.manager.timeout** parameter. If the WAL splitting process fails to complete within the timeout period, the task is submitted again. Multiple WAL splitting tasks may be submitted during a specified period. If the **temp** file is deleted when one WAL splitting task completes, other tasks cannot find the file and the FileNotFound exception is reported. To avoid the problem, perform the following modifications:

The default value of **hbase.splitlog.manager.timeout** is 600,000 ms. The cluster specification is that each RegionServer has 2,000 to 3,000 regions. When the cluster is normal (HBase is normal and HDFS does not have a large number of read and write operations), you are advised to adjust this parameter based on the cluster specifications. If the actual specifications (the actual average number of regions on each RegionServer) are greater than the default specifications (the default average number of regions on each RegionServer, that is, 2,000), the adjustment solution is (actual specifications/default specifications) x Default time.

Set the **splitlog** parameter in the **hbase-site.xml** file on the server. [Table 8-27](#) describes the parameter.

Table 8-27 Description of the **splitlog** parameter

Parameter	Description	Default Value
hbase.splitlog.manager.timeout	Timeout period for receiving worker response by the distributed SplitLog management program.	600000

8.20.18 Why Does the ImportTsv Tool Display "Permission denied" When the Same Linux User as and a Different Kerberos User from the Region Server Are Used?

Question

When the same Linux user (for example, user **omm**) as and a different Kerberos user (for example, user **admin**) from the Region Server are used, why does the ImportTsv tool fail to be executed and the error message "Permission denied" is displayed?

```
Exception in thread "main" org.apache.hadoop.security.AccessControlException: Permission denied:
user=admin, access=WRITE, inode="/user/omm-bulkload/hbase-staging/
partitions_cab16de5-87c2-4153-9cca-a6f4ed4278a6":hbase:hadoop:drwx--x--x
    at org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.check(FSPermissionChecker.java:342)
    at org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.check(FSPermissionChecker.java:315)
    at
org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkPermission(FSPermissionChecker.java:23
1)
    at
com.xxx.hadoop.adapter.hdfs.plugin.xxAccessControlEnforce.checkPermission(xxAccessControlEnforce.java:69)
    at
org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkPermission(FSPermissionChecker.java:19
0)
    at org.apache.hadoop.hdfs.server.namenode.FSDirectory.checkPermission(FSDirectory.java:1789)
    at org.apache.hadoop.hdfs.server.namenode.FSDirectory.checkPermission(FSDirectory.java:1773)
    at org.apache.hadoop.hdfs.server.namenode.FSDirectory.checkAncestorAccess(FSDirectory.java:1756)
    at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.startFileInternal(FSNamesystem.java:2490)
    at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.startFileInt(FSNamesystem.java:2425)
    at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.startFile(FSNamesystem.java:2308)
    at
org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.create(NameNodeRpcServer.java:745)
    at
org.apache.hadoop.hdfs.protocolPB.ClientNameNodeProtocolServerSideTranslatorPB.create(ClientNameNodeP
rotocolServerSideTranslatorPB.java:434)
    at org.apache.hadoop.hdfs.protocol.proto.ClientNameNodeProtocolProtos$ClientNameNodeProtocol
$2.callBlockingMethod(ClientNameNodeProtocolProtos.java)
    at org.apache.hadoop.ipc.ProtobufRpcEngine$Server
$ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:616)
    at org.apache.hadoop.ipc.RPC$Server.call(RPC.java:973)
    at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2260)
    at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2256)
    at java.security.AccessController.doPrivileged(Native Method)
    at javax.security.auth.Subject.doAs(Subject.java:422)
    at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1781)
    at org.apache.hadoop.ipc.Server$Handler.run(Server.java:2254)
```

Answer

The ImportTsv tool creates a partition file in the HBase temporary directory specified by **hbase.fs.tmp.dir** in the *Client installation path* **/HBase/hbase/conf/hbase-site.xml** file. Therefore, the client (Kerberos user) must have the **rwX** permission on the specified temporary directory to perform the ImportTsv operation. The default value of **hbase.fs.tmp.dir** is **/user/\${user.name}/hbase-staging** (for example, **/user/omm/hbase-staging**). **\${user.name}** indicates the OS username (user **omm**). The client (Kerberos user, for example, user **admin**) does not have the **rwX** permission on the directory.

To solve the preceding problem, perform the following steps:

1. On the client, set **hbase.fs.tmp.dir** to the directory of the current Kerberos user (for example, **/user/admin/hbase-staging**), or provide the **rwX** permission required by the configured directory for the client (Kerberos user).
2. Perform the ImportTsv operation again.

8.20.19 Insufficient Rights When a Tenant Accesses Phoenix

Question

When a tenant accesses Phoenix, a message is displayed indicating that the tenant has insufficient rights.

Answer

You need to associate the HBase service and Yarn queues when creating a tenant.

The tenant must be granted additional rights to perform operations on Phoenix, that is, the RWX permission on the Phoenix system table.

Example:

Tenant **hbase** has been created. Log in to the HBase Shell as user **admin** and run the **scan 'hbase:acl'** command to query the role of the tenant. The role is **hbase_1450761169920** (in the format of tenant name_timestamp).

Run the following commands to grant rights to the tenant (if the Phoenix system table has not been generated, log in to the Phoenix client as user **admin** first and then grant rights on the HBase Shell):

```
grant '@hbase_1450761169920','RWX','SYSTEM.CATALOG'
```

```
grant '@hbase_1450761169920','RWX','SYSTEM.FUNCTION'
```

```
grant '@hbase_1450761169920','RWX','SYSTEM.SEQUENCE'
```

```
grant '@hbase_1450761169920','RWX','SYSTEM.STATS'
```

Create user **phoenix** and bind it with tenant **hbase**, so that tenant **hbase** can access the Phoenix client as user **phoenix**.

8.20.20 What Can I Do When HBase Fails to Recover a Task and a Message Is Displayed Stating "Rollback recovery failed"?

Question

The system automatically rolls back data after an HBase recovery task fails. If "Rollback recovery failed" is displayed, the rollback fails. After the rollback fails, data stops being processed and the junk data may be generated. How can I resolve this problem?


Answer

You need to manually clear the junk data before performing the backup or recovery task next time.

- Step 1** Install the cluster client in `/opt/client`.
- Step 2** Run `source /opt/client/bigdata_env` as the client installation user to configure environment variables.
- Step 3** Run the `kinit admin` command.
- Step 4** Run `zkCli.sh -server business IP address of ZooKeeper:2181` to connect to the ZooKeeper.
- Step 5** Run `deleteall /recovering` to delete the junk data. Run `quit` to disconnect ZooKeeper.

 **NOTE**

Running this command will cause data loss. Exercise caution.

- Step 6** Run `hdfs dfs -rm -f -r /user/hbase/backup` to delete temporary data.
- Step 7** Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Backup and Restoration > Restoration Management**. In the task list, locate the row that contains the target task and click **View History** in the **Operation** column. In the displayed dialog box, click  before a specified execution record to view the snapshot name.
Snapshot [*snapshot name*] is created successfully before recovery.
- Step 8** Switch to the client, run `hbase shell`, and then `delete_all_snapshot 'snapshot name.*'` to delete the temporary snapshot.

----End

8.20.21 How Do I Fix Region Overlapping?

Question

When the HBaseFsck tool is used to check the region status in MRS 3.x and later versions, if the log contains **ERROR: (regions region1 and region2) There is an overlap in the region chain** or **ERROR: (region region1) Multiple regions have the same startkey: xxx**, overlapping exists in some regions. How do I solve this problem?

Answer

To rectify the fault, perform the following steps:

- Step 1** Run the `hbase hbck -repair tableName` command to restore the table that contains overlapping.
- Step 2** Run the `hbase hbck tableName` command to check whether overlapping exists in the restored table.
 - If overlapping does not exist, go to [Step 3](#).
 - If overlapping exists, go to [Step 1](#).
- Step 3** Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > HBase > More > Perform HMaster Switchover** to complete the HMaster active/standby switchover.

- Step 4** Run the `hbase hbck tableName` command to check whether overlapping exists in the restored table.
- If overlapping does not exist, no further action is required.
 - If overlapping still exists, start from [Step 1](#) to perform the recovery again.
- End

8.20.22 Why Does RegionServer Fail to Be Started When GC Parameters Xms and Xmx of HBase RegionServer Are Set to 31 GB?

Question

(MRS 3.x and later versions) Check the `hbase-omm-*.out` log of the node where RegionServer fails to be started. It is found that the log contains **An error report file with more information is saved as: /tmp/hs_err_pid*.log**. Check the `/tmp/hs_err_pid*.log` file. It is found that the log contains **#Internal Error (vtableStubs_aarch64.cpp:213), pid=9456, tid=0x0000ffff97fdd200 and #guarantee(__ pc() <= s->code_end()) failed: overflowed buffer**, indicating that the problem is caused by JDK. How do I solve this problem?

Answer

To rectify the fault, perform the following steps:

- Step 1** Run the `su - omm` command on a node where RegionServer fails to be started to switch to user `omm`.
- Step 2** Run the `java -XX:+PrintFlagsFinal -version |grep HeapBase` command as user `omm`. Information similar to the following is displayed:
- ```
uintx HeapBaseMinAddress = 2147483648 {pd product}
```
- Step 3** Change the values of `-Xms` and `-Xmx` in `GC_OPTS` to values that are not between `32G-HeapBaseMinAddress` and `32G`, excluding the values of `32G` and `32G-HeapBaseMinAddress`.
- Step 4** Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance**, select the failed instance, and choose **More** > **Restart Instance** to restart the failed instance.
- End

## 8.20.23 Why Does the LoadIncrementalHFiles Tool Fail to Be Executed and "Permission denied" Is Displayed When Nodes in a Cluster Are Used to Import Data in Batches?

### Question

Why does the LoadIncrementalHFiles tool fail to be executed and "Permission denied" is displayed when a Linux user is manually created in a normal cluster and DataNode in the cluster is used to import data in batches?

```

2020-09-20 14:53:53,808 WARN [main] shortcircuit.DomainSocketFactory: error creating DomainSocket
java.net.ConnectException: connect(2) error: Permission denied when trying to connect to '/var/run/
FusionInsight-HDFS/dn_socket'
 at org.apache.hadoop.net.unix.DomainSocket.connect0(Native Method)
 at org.apache.hadoop.net.unix.DomainSocket.connect(DomainSocket.java:256)
 at org.apache.hadoop.hdfs.shortcircuit.DomainSocketFactory.createSocket(DomainSocketFactory.java:168)
 at org.apache.hadoop.hdfs.client.impl.BlockReaderFactory.nextDomainPeer(BlockReaderFactory.java:804)
 at
org.apache.hadoop.hdfs.client.impl.BlockReaderFactory.createShortCircuitReplicaInfo(BlockReaderFactory.java
:526)
 at org.apache.hadoop.hdfs.shortcircuit.ShortCircuitCache.create(ShortCircuitCache.java:785)
 at org.apache.hadoop.hdfs.shortcircuit.ShortCircuitCache.fetchOrCreate(ShortCircuitCache.java:722)
 at
org.apache.hadoop.hdfs.client.impl.BlockReaderFactory.getBlockReaderLocal(BlockReaderFactory.java:483)
 at org.apache.hadoop.hdfs.client.impl.BlockReaderFactory.build(BlockReaderFactory.java:360)
 at org.apache.hadoop.hdfs.DFSInputStream.getBlockReader(DFSInputStream.java:663)
 at org.apache.hadoop.hdfs.DFSInputStream.blockSeekTo(DFSInputStream.java:594)
 at org.apache.hadoop.hdfs.DFSInputStream.readWithStrategy(DFSInputStream.java:776)
 at org.apache.hadoop.hdfs.DFSInputStream.read(DFSInputStream.java:845)
 at java.io.DataInputStream.readFully(DataInputStream.java:195)
 at org.apache.hadoop.hbase.io.hfile.FixedFileTrailer.readFromStream(FixedFileTrailer.java:401)
 at org.apache.hadoop.hbase.io.hfile.HFile.isHFileFormat(HFile.java:651)
 at org.apache.hadoop.hbase.io.hfile.HFile.isHFileFormat(HFile.java:634)
 at org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.visitBulkHFiles(LoadIncrementalHFiles.java:1090)
 at
org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.discoverLoadQueue(LoadIncrementalHFiles.java:1006)
 at
org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.prepareHFileQueue(LoadIncrementalHFiles.java:257)
 at org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.doBulkLoad(LoadIncrementalHFiles.java:364)
 at org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.run(LoadIncrementalHFiles.java:1263)
 at org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.run(LoadIncrementalHFiles.java:1276)
 at org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.run(LoadIncrementalHFiles.java:1311)
 at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:76)
 at org.apache.hadoop.hbase.tool.LoadIncrementalHFiles.main(LoadIncrementalHFiles.java:1333)

```

## Answer

If the client that the LoadIncrementalHFiles tool depends on is installed in the cluster and is on the same node as DataNode, HDFS creates short-circuit read during the execution of the tool to improve performance. The short-circuit read depends on the **/var/run/FusionInsight-HDFS** directory (**dfs.domain.socket.path**). The default permission on this directory is **750**. This user does not have the permission to operate the directory.

To solve the preceding problem, perform the following operations:

Method 1: Create a user (recommended).

**Step 1** Create a user on Manager. By default, the user group contains the **ficommon** group.

```
[root@xxx-xxx-xxx-xxx ~]# id test
uid=20038(test) gid=9998(ficommon) groups=9998(ficommon)
```

**Step 2** Import data again.

----End

Method 2: Change the owner group of the current user.

**Step 1** Add the user to the **ficommon** group.

```
[root@xxx-xxx-xxx-xxx ~]# usermod -a -G ficommon test
[root@xxx-xxx-xxx-xxx ~]# id test
uid=2102(test) gid=2102(test) groups=2102(test),9998(ficommon)
```

**Step 2** Import data again.

----End

## 8.20.24 Why Is the Error Message "import argparse" Displayed When the Phoenix sqlline Script Is Used?

### Question

When the sqlline script is used on the client, the error message "import argparse" is displayed.

### Answer

**Step 1** Log in to the node where the HBase client is installed as user **root**. Perform security authentication using the **hbase** user.

**Step 2** Go to the directory where the sqlline script of the HBase client is stored and run the **python3 sqlline.py** command.

----End

## 8.20.25 How Do I Deal with the Restrictions of the Phoenix BulkLoad Tool?

### Question

When the indexed field data is updated, if a batch of data exists in the user table, the BulkLoad tool cannot update the global and partial mutable indexes.

### Answer

#### Problem Analysis

1. Create a table.

```
CREATE TABLE TEST_TABLE(
DATE varchar not null,
NUM integer not null,
SEQ_NUM integer not null,
ACCOUNT1 varchar not null,
ACCOUNTDES varchar,
FLAG varchar,
SALL double,
CONSTRAINT PK PRIMARY KEY (DATE,NUM,SEQ_NUM,ACCOUNT1)
);
```

2. Create a global index.

```
CREATE INDEX TEST_TABLE_INDEX ON
TEST_TABLE(ACCOUNT1,DATE,NUM,ACCOUNTDES,SEQ_NUM);
```

3. Insert data.

```
UPSERT INTO TEST_TABLE
(DATE,NUM,SEQ_NUM,ACCOUNT1,ACCOUNTDES,FLAG,SALL) values
('20201001',30201001,13,'367392332','sffa1','',);
```

4. Execute the BulkLoad task to update data.

**hbase org.apache.phoenix.mapreduce.CsvBulkLoadTool -t TEST\_TABLE -i /tmp/test.csv**, where the content of **test.csv** is as follows:

|         |          |    |           |         |         |    |
|---------|----------|----|-----------|---------|---------|----|
| 2020001 | 30201001 | 13 | 367392332 | sffa888 | 1231243 | 23 |
|---------|----------|----|-----------|---------|---------|----|

- Symptom: The existing index data cannot be directly updated. As a result, two pieces of index data exist.

```

+-----+-----+-----+-----+-----+
|:ACCOUNT1 | :DATE | :NUM | 0:ACCOUNTDES |:SEQ_NUM |
+-----+-----+-----+-----+
| 367392332 | 20201001 | 30201001 | sffa1 | 13 |
| 367392332 | 20201001 | 30201001 | sffa888 | 13 |
+-----+-----+-----+-----+

```

### Solution

- Step 1 Delete the old index table.

```
DROP INDEX TEST_TABLE_INDEX ON TEST_TABLE;
```

- Step 2 Create an index table in asynchronous mode.

```
CREATE INDEX TEST_TABLE_INDEX ON TEST_TABLE(ACCOUNT1,DATE,NUM,ACCOUNTDES,SEQ_NUM) ASYNC;
```

- Step 3 Recreate a index.

```
hbase org.apache.phoenix.mapreduce.index.IndexTool --data-table TEST_TABLE --index-table TEST_TABLE_INDEX --output-path /user/test_table
```

----End

## 8.20.26 Why a Message Is Displayed Indicating that the Permission is Insufficient When CTBase Connects to the Ranger Plug-ins?

### Question

When CTBase accesses the HBase service with the Ranger plug-ins enabled and you are creating a cluster table, a message is displayed indicating that the permission is insufficient.

```

ERROR: Create ClusterTable failed. Error: org.apache.hadoop.hbase.security.AccessDeniedException:
Insufficient permissions for user 'ctbase2@HADOOP.COM' (action=create)
at org.apache.ranger.authorization.hbase.AuthorizationSession.publishResults(AuthorizationSession.java:278)
at
org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.authorizeAccess(RangerAuthorizatio
nCoprocesor.java:654)
at
org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.requirePermission(RangerAuthorizatio
nCoprocesor.java:772)
at
org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preCreateTable(RangerAuthorizatio
nCoprocesor.java:943)
at
org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preCreateTable(RangerAuthorizatio
nCoprocesor.java:428)
at org.apache.hadoop.hbase.master.MasterCoprocesorHost$12.call(MasterCoprocesorHost.java:351)
at org.apache.hadoop.hbase.master.MasterCoprocesorHost$12.call(MasterCoprocesorHost.java:348)

```

```
at org.apache.hadoop.hbase.coprocessor.CoprocessorHost
$ObserverOperationWithoutResult.callObserver(CoprocessorHost.java:581)
at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.execOperation(CoprocessorHost.java:655)
at
org.apache.hadoop.hbase.master.MasterCoprocessorHost.preCreateTable(MasterCoprocessorHost.java:348)
at org.apache.hadoop.hbase.master.HMaster$5.run(HMaster.java:2192)
at
org.apache.hadoop.hbase.master.procedure.MasterProcedureUtil.submitProcedure(MasterProcedureUtil.java:1
34)
at org.apache.hadoop.hbase.master.HMaster.createTable(HMaster.java:2189)
at org.apache.hadoop.hbase.master.MasterRpcServices.createTable(MasterRpcServices.java:711)
at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService
$2.callBlockingMethod(MasterProtos.java)
at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:458)
at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:133)
at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:338)
at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:318)
```

## Answer

CTBase users can configure permission policies on the Ranger page and grant the READ, WRITE, CREATE, ADMIN, and EXECUTE permissions to the CTBase metadata table `_ctmeta_`, cluster table, and index table.



# 9 Using HDFS

## 9.1 Using Hadoop from Scratch

You can use Hadoop to submit wordcount jobs. Wordcount is the most classic Hadoop job and is used to count the number of words in massive text.

### Procedure

**Step 1** Prepare the wordcount program.

Multiple open source Hadoop sample programs are provided, including wordcount. You can download the Hadoop sample program from <https://dist.apache.org/repos/dist/release/hadoop/common/>.

For example, choose **hadoop-x.x.x**. On the page that is displayed, click **hadoop-x.x.x.tar.gz** to download it. Then, decompress it to obtain **hadoop-mapreduce-examples-x.x.x.jar** (the Hadoop sample program) from **hadoop-x.x.x\share\hadoop\mapreduce**. The **hadoop-mapreduce-examples-x.x.x.jar** package contains the wordcount program.

#### NOTE

**hadoop-x.x.x** indicates the Hadoop version. Choose a version based on your requirements.

**Step 2** Prepare data files.

There is no format requirement for data files. Prepare one or more **.txt** files. The following are examples of the **.txt** file:

```
qw sdfhoedfrffrofhunccgktpmhutopmma
jjpsffjfgjgtyiuymhombmbogohoyhm
jhheyombdhuaqqiqyebchdmamdhdemmj
doeyhjwedcrfvgtgbmojjyhhqssdddddskf
kjhjhkehdeiyrudjfhfhffooqweopuyyyy
```


**Step 3** Upload data to OBS.

1. Log in to OBS Console.
2. Click **Parallel File System** and choose **Create Parallel File System** to create a file system named **wordcount01**.

**wordcount01** is only an example. The file system name must be globally unique. Otherwise, the parallel file system fails to be created.


3. In the OBS file system list, click **wordcount01** and choose **Files > Create Folder** to create the **program** and **input** folders, as shown in **Figure 9-1**.

**Figure 9-1** Folder list of the wordcount01 file system



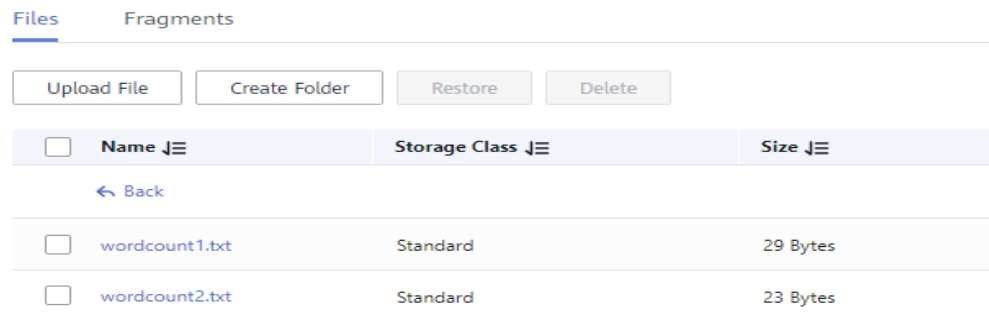
- **program**: stores user programs.
  - **input**: stores user data files.
4. Go to the **program** folder, choose **Upload File > add file**, select the program package downloaded in **Step 1** from the local host, and click **Upload**. After the upload is complete, the page shown in **Figure 9-2** is displayed.

**Figure 9-2** Program list



5. Go to the **input** folder and upload the data file prepared in **Step 2** to the **input** folder. After the upload is complete, the page shown in **Figure 9-3** is displayed.

**Figure 9-3** Data file list

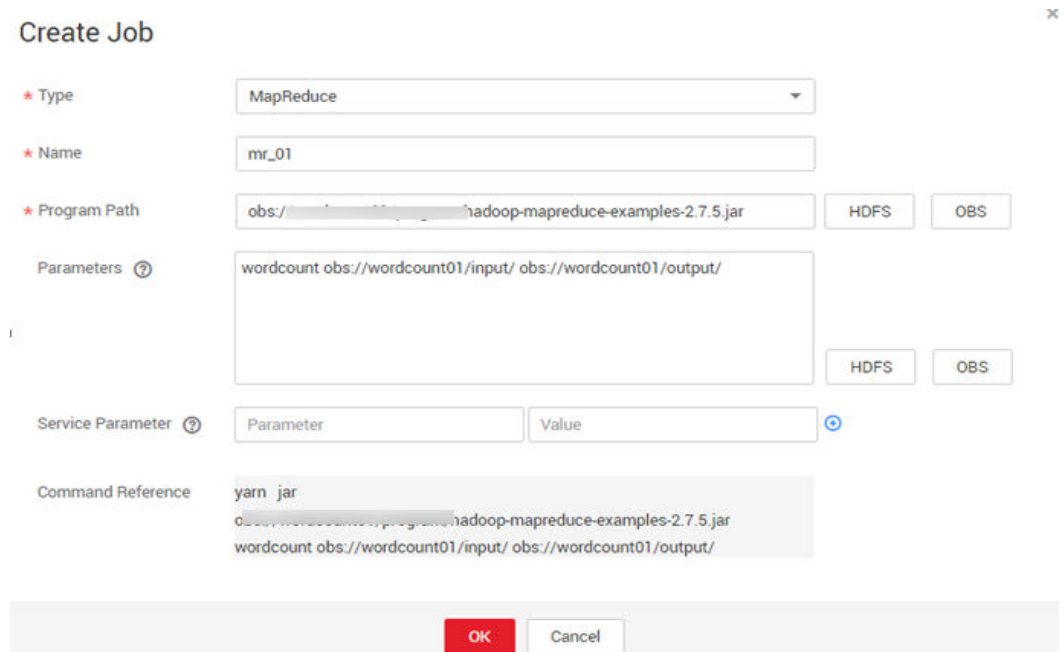


**Step 4** Log in to the MRS console. In the navigation pane on the left, click **Clusters** and choose **Active Clusters**. Click the cluster name. The cluster must contain Hadoop components.

**Step 5** Submit the wordcount job.

On the MRS console, click the **Jobs** tab and click **Create**. The **Create Job** page is displayed.

**Figure 9-4** wordcount job



- Set **Type** to **MapReduce**.
- Set **Name** to **mr\_01**.
- Set the path of the executable program to the address of the program stored on the OBS, for example, **obs://wordcount01/program/hadoop-mapreduce-examples-x.x.x.jar**.
- Enter **wordcount obs://wordcount01/input/ obs://wordcount01/output/** in the **Parameter** pane.

**NOTE**

- Replace the OBS file system name in **obs://wordcount01/input/** with the actual name of the file system created in the environment.
  - Replace the OBS file system name in **obs://wordcount01/output/** with the name of the file system created in the actual environment. Replace **output** with a directory that does not exist based on site requirements.
- **Service Parameter** can be left blank.

A job can be submitted only when the cluster is in the **Running** state.

After a job is submitted successfully, it is in the **Accepted** state by default. You do not need to manually execute the job.

**Step 6** View the job execution result.

1. Go to the **Jobs** tab page and check whether the job is successfully executed.  
It takes some time to run the job. After the job is complete, refresh the job list  
Once a job has succeeded or failed, you cannot execute it again. However, you can add or copy a job, and set job parameters to submit a job again.
2. Log in to the OBS console, go to the OBS path, and view the job output information.

You can view output files in the **output** directory created in [Step 5](#). You need to download the file to the local host and open it in text format, as shown in [Figure 9-5](#).

**Figure 9-5** Output file list

| Name       | Storage Class | Size    | Encrypted | Restoration Status | Last Modified                   | Operation           |
|------------|---------------|---------|-----------|--------------------|---------------------------------|---------------------|
| ._SUCCESS  | Standard      | 0 byte  | No        | --                 | Jan 20, 2020 10:01:21 GMT+08:00 | Download Share More |
| part-00000 | Standard      | 44 byte | No        | --                 | Jan 20, 2020 10:01:21 GMT+08:00 | Download Share More |

----End

## 9.2 Configuring Memory Management

### Scenario

In HDFS, each file object needs to register corresponding information in the NameNode and occupies certain storage space. As the number of files increases, if the original memory space cannot store the corresponding information, you need to change the memory size.

### Configuration Description

**Navigation path for setting parameters:**

Go to the **All Configurations** page of HDFS by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-1** Parameter description

| Parameter  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Default Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GC_PROFILE | <p>The NameNode memory size depends on the size of FsImage, which can be calculated based on the following formula:<br/>FsImage size = Number of files x 900 bytes. You can estimate the memory size of the NameNode of HDFS based on the calculation result.</p> <p>The value range of this parameter is as follows:</p> <ul style="list-style-type: none"> <li>• <b>high:</b> 4 GB</li> <li>• <b>medium:</b> 2 GB</li> <li>• <b>low:</b> 256 MB</li> <li>• <b>custom:</b> The memory size can be set according to the data size in GC_OPTS.</li> </ul> | custom                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| GC_OPTS    | <p>JVM parameter used for garbage collection (GC). This parameter is valid only when <b>GC_PROFILE</b> is set to <b>custom</b>. Ensure that the <b>GC_OPT</b> parameter is set correctly. Otherwise, the process will fail to be started.</p> <p><b>NOTICE</b><br/>Exercise caution when you modify the configuration. If the configuration is incorrect, the services are unavailable.</p>                                                                                                                                                              | <p>-Xms2G -Xmx4G -<br/>XX:NewSize=128M -<br/>XX:MaxNewSize=256M -<br/>XX:MetaspaceSize=128M -<br/>XX:MaxMetaspaceSize=128M -<br/>XX:+UseConcMarkSweepGC -<br/>XX:+CMSParallelRemarkEnabled<br/>-XX:CMSInitiatingOccupancy-<br/>Fraction=65 -XX:+PrintGCDetails<br/>-<br/>Dsun.rmi.dgc.client.gcInterval=0<br/>x7FFFFFFFFFFFFFFE -<br/>Dsun.rmi.dgc.server.gcInterval=0<br/>x7FFFFFFFFFFFFFFE -XX:-<br/>OmitStackTraceInFastThrow -<br/>XX:+PrintGCDateStamps -<br/>XX:+UseGCLogFileRotation -<br/>XX:NumberOfGCLogFiles=10 -<br/>XX:GCLogFileSize=1M -<br/>Djdk.tls.ephemeralDHKeySize=2<br/>048</p> |

## 9.3 Creating an HDFS Role

### Scenario

This section describes how to create and configure an HDFS role on FusionInsight Manager. The HDFS role is granted the rights to read, write, and execute HDFS directories or files.

A user has the complete permission on the created HDFS directories or files, that is, the user can directly read data from and write data to as well as authorize others to access the HDFS directories or files.

#### NOTE

- This section applies to MRS 3.x or later clusters.
- An HDFS role can be created only in security mode.
- If the current component uses Ranger for permission control, HDFS policies must be configured based on Ranger for permission management. For details, see [Adding a Ranger Access Permission Policy for HDFS](#).

### Prerequisites

The MRS cluster administrator has understood service requirements.

### Procedure

**Step 1** Log in to FusionInsight Manager, and choose **System > Permission > Role**.

**Step 2** On the displayed page, click **Create Role** and fill in **Role Name** and **Description**.

**Step 3** Configure the resource permission. For details, see [Table 9-2](#).

**File System:** HDFS directory and file permission

Common HDFS directories are as follows:

- **flume:** Flume data storage directory
- **hbase:** HBase data storage directory
- **mr-history:** MapReduce task information storage directory
- **tmp:** temporary data storage directory
- **user:** user data storage directory

**Table 9-2** Setting a role

| Task                                      | Operation                                                                                                                                                                                                                        |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Setting the HDFS administrator permission | In the <b>Configure Resource Permission</b> area, choose <i>Name of the desired cluster</i> > HDFS, and select <b>Cluster Admin Operations</b> .<br><b>NOTE</b><br>The setting takes effect after the HDFS service is restarted. |

| Task                                                                                                              | Operation                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Setting the permission for users to check and recover HDFS                                                        | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> area, choose <i>Name of the desired cluster</i> &gt; HDFS &gt; <b>File System</b>.</li> <li>2. Locate the save path of specified directories or files on HDFS.</li> <li>3. In the <b>Permission</b> column of the specified directories or files, select <b>Read</b> and <b>Execute</b>.</li> </ol> |
| Setting the permission for users to read directories or files of other users                                      | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> area, choose <i>Name of the desired cluster</i> &gt; HDFS &gt; <b>File System</b>.</li> <li>2. Locate the save path of specified directories or files on HDFS.</li> <li>3. In the <b>Permission</b> column of the specified directories or files, select <b>Read</b> and <b>Execute</b>.</li> </ol> |
| Setting the permission for users to write data to files of other users                                            | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> area, choose <i>Name of the desired cluster</i> &gt; HDFS &gt; <b>File System</b>.</li> <li>2. Locate the save path of specified files on HDFS.</li> <li>3. In the <b>Permission</b> column of the specified files, select <b>Write</b> and <b>Execute</b>.</li> </ol>                              |
| Setting the permission for users to create or delete sub-files or sub-directories in the directory of other users | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> area, choose <i>Name of the desired cluster</i> &gt; HDFS &gt; <b>File System</b>.</li> <li>2. Locate the path where the specified directory is saved in the HDFS.</li> <li>3. In the <b>Permission</b> column of the specified directories, select <b>Write</b> and <b>Execute</b>.</li> </ol>     |
| Setting the permission for users to execute directories or files of other users                                   | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> area, choose <i>Name of the desired cluster</i> &gt; HDFS &gt; <b>File System</b>.</li> <li>2. Locate the save path of specified directories or files on HDFS.</li> <li>3. In the <b>Permission</b> column of the specified directories or files, select <b>Execute</b>.</li> </ol>                 |

| Task                                                                                                      | Operation                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Setting the permission for allowing subdirectories to inherit all permissions of their parent directories | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> area, choose <i>Name of the desired cluster</i> &gt; HDFS &gt; <b>File System</b>.</li> <li>2. Locate the save path of specified directories or files on HDFS.</li> <li>3. In the <b>Permission</b> column of the specified directories or files, select <b>Recursive</b>.</li> </ol> |

**Step 4** Click **OK**, and return to the **Role** page.

----End

## 9.4 Using the HDFS Client

### Scenario

This section describes how to use the HDFS client in an O&M scenario or service scenario.

### Prerequisites

- The client has been installed.  
For example, the installation directory is **/opt/client**. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user needs to change the password upon the first login. (This operation is not required in normal mode.)

### Using the HDFS Client

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** If the cluster is in security mode, run the following command to authenticate the user. In normal mode, user authentication is not required.

```
kinit Component service user
```

**Step 5** Run the HDFS Shell command. Example:



```
hdfs dfs -ls /
----End
```

## Common HDFS Client Commands

The following table lists common HDFS client commands.

**Table 9-3** Common HDFS client commands

| Command                                                                              | Description                                                 | Example                                                                                                                                  |
|--------------------------------------------------------------------------------------|-------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>hdfs dfs -mkdir</b> <i>Folder name</i>                                            | Used to create a folder.                                    | <b>hdfs dfs -mkdir /tmp/mydir</b>                                                                                                        |
| <b>hdfs dfs -ls</b> <i>Folder name</i>                                               | Used to view a folder.                                      | <b>hdfs dfs -ls /tmp</b>                                                                                                                 |
| <b>hdfs dfs -put</b> <i>Local file on the client node Specified HDFS path</i>        | Used to upload a local file to a specified HDFS path.       | <b>hdfs dfs -put /opt/test.txt /tmp</b><br>Upload the <b>/opt/test.txt</b> file on the client node to the <b>/tmp</b> directory of HDFS. |
| <b>hdfs dfs -get</b> <i>Specified file on HDFS Specified path on the client node</i> | Used to download the HDFS file to the specified local path. | <b>hdfs dfs -get /tmp/test.txt /opt/</b><br>Download the <b>/tmp/test.txt</b> file on HDFS to the <b>/opt</b> path on the client node.   |
| <b>hdfs dfs -rm -r -f</b> <i>Specified folder on HDFS</i>                            | Used to delete a folder.                                    | <b>hdfs dfs -rm -r -f /tmp/mydir</b>                                                                                                     |
| <b>hdfs dfs -chmod</b> <i>Permission parameter File directory</i>                    | Used to configure the HDFS directory permission for a user. | <b>hdfs dfs -chmod 700 /tmp/test</b>                                                                                                     |

## Client-related FAQs

1. What do I do when the HDFS client exits abnormally and error message "java.lang.OutOfMemoryError" is displayed after the HDFS client command is running?

This problem occurs because the memory required for running the HDFS client exceeds the preset upper limit (128 MB by default). You can change the memory upper limit of the client by modifying **CLIENT\_GC\_OPTS** in *<Client installation path>/HDFS/component\_env*. For example, if you want to set the upper limit to 1 GB, run the following command:

```
CLIENT_GC_OPTS="-Xmx1G"
```

After the modification, run the following command to make the modification take effect:

```
source <Client installation path>/bigdata_env
```

2. How do I set the log level when the HDFS client is running?  
By default, the logs generated during the running of the HDFS client are printed to the console. The default log level is INFO. To enable the DEBUG log level for fault locating, run the following command to export an environment variable:  
**export HADOOP\_ROOT\_LOGGER=DEBUG,console**  
Then run the HDFS Shell command to generate the DEBUG logs.  
If you want to print INFO logs again, run the following command:  
**export HADOOP\_ROOT\_LOGGER=INFO,console**
3. How do I delete HDFS files permanently?  
HDFS provides a recycle bin mechanism. Typically, after an HDFS file is deleted, the file is moved to the recycle bin of HDFS. If the file is no longer needed and the storage space needs to be released, clear the corresponding recycle bin directory, for example, **hdfs://hacluster/user/xxx/Trash/Current/xxx**.

## 9.5 Running the DistCp Command

### Scenario

DistCp is a tool used to perform large-amount data replication between clusters or in a cluster. It uses MapReduce tasks to implement distributed copy of a large amount of data.

### Prerequisites

- The Yarn client or a client that contains Yarn has been installed. For example, the installation directory is **/opt/client**.
- Service users of each component are created by the MRS cluster administrator based on service requirements. In security mode, machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login. (Not involved in normal mode)
- To copy data between clusters, you need to enable the inter-cluster data copy function on both clusters.

### Procedure

- Step 1** Log in to the node where the client is installed.
- Step 2** Run the following command to go to the client installation directory:  
**cd /opt/client**
- Step 3** Run the following command to configure environment variables:  
**source bigdata\_env**
- Step 4** If the cluster is in security mode, the user group to which the user executing the DistCp command belongs must be **supergroup** and the user run the following command to perform user authentication. In normal mode, user authentication is not required.

**kinit** *Component service user*

**Step 5** Run the DistCp command. The following provides an example:

```
hadoop distcp hdfs://hacluster/source hdfs://hacluster/target
```

----End

## Common Usage of DistCp

1. The following is an example of the commonest usage of DistCp:

```
hadoop distcp -numListstatusThreads 40 -update -delete -prbugpaxtq hdfs://cluster1/source hdfs://cluster2/target
```

### NOTE

In the preceding command:

- **-numListstatusThreads** specifies the number of threads for creating the list of 40 copied files.
- **-update -delete** specifies that files at the source location and the target location are synchronized, and that files with excessive target locations are deleted. If you need to copy files incrementally, delete **-delete**.
- If **-prbugpaxtq** and **-update** are used, it indicates that the status information of the copied file is also updated.
- **hdfs://cluster1/source** indicates the source location, and **hdfs://cluster2/target** indicates the target location.

2. The following is an example of data copy between clusters:

```
hadoop distcp hdfs://cluster1/foo/bar hdfs://cluster2/bar/foo
```

### NOTE

The network between cluster1 and cluster2 must be reachable, and the two clusters must use the same HDFS version or compatible HDFS versions.

3. The following are multiple examples of data copy in a source directory:

```
hadoop distcp hdfs://cluster1/foo/a \
hdfs://cluster1/foo/b \
hdfs://cluster2/bar/foo
```

The preceding command is used to copy the folders a and b of cluster1 to the **/bar/foo** directory of cluster2. The effect is equivalent to that of the following commands:

```
hadoop distcp -f hdfs://cluster1/srclist \
hdfs://cluster2/bar/foo
```

The content of **srclist** is as follows. Before running the DistCp command, upload the **srclist** file to HDFS.

```
hdfs://cluster1/foo/a
hdfs://cluster1/foo/b
```

4. **-update** indicates that a to-be-copied file does not exist in the target location, or the content of the copied file in the target location is updated; and **-overwrite** is used to overwrite existing files in the target location.

The following is an example of the difference between no option and any one of the two options (either **update** or **overwrite**) that is added:

Assume that the structure of a file at the source location is as follows:

```
hdfs://cluster1/source/first/1
hdfs://cluster1/source/first/2
hdfs://cluster1/source/second/10
hdfs://cluster1/source/second/20
```

Commands without options are as follows:

```
hadoop distcp hdfs://cluster1/source/first hdfs://cluster1/source/second hdfs://cluster2/target
```

By default, the preceding command creates the **first** and **second** folders at the target location. Therefore, the copy results are as follows:

```
hdfs://cluster2/target/first/1
hdfs://cluster2/target/first/2
hdfs://cluster2/target/second/10
hdfs://cluster2/target/second/20
```

The command with any one of the two options (for example, **update**) is as follows:

```
hadoop distcp -update hdfs://cluster1/source/first hdfs://cluster1/source/second hdfs://cluster2/target
```

The preceding command copies only the content at the source location to the target location. Therefore, the copy results are as follows:

```
hdfs://cluster2/target/1
hdfs://cluster2/target/2
hdfs://cluster2/target/10
hdfs://cluster2/target/20
```

 **NOTE**

- If files with the same name exist in multiple source locations, the DistCp command fails.
  - If neither **update** nor **overwrite** is used and the file to be copied already exists in the target location, the file will be skipped.
  - When **update** is used, if the file to be copied already exists in the target location but the file content is different, the file content in the target location is updated.
  - When **overwrite** is used, if the file to be copied already exists in the target location, the file in the target location is still overwritten.
5. The following table describes other command options:

**Table 9-4** Other command options

| Option         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -p[rbugpcaxtq] | When <b>-update</b> is also used, the status information of a copied file is updated even if the content of the copied file is not updated.<br><b>r</b> : number of copies<br><b>b</b> : size of a block<br><b>u</b> : user to which the files belong<br><b>g</b> : user group to which the user belongs<br><b>p</b> : permission<br><b>c</b> : check and type<br><b>a</b> : access control<br><b>t</b> : timestamp<br><b>q</b> : quota information |
| -i             | Failures ignored during copying                                                                                                                                                                                                                                                                                                                                                                                                                     |
| -log <logdir>  | Path of the specified log                                                                                                                                                                                                                                                                                                                                                                                                                           |
| -v             | Additional information in the specified log                                                                                                                                                                                                                                                                                                                                                                                                         |

| Option                            | Description                                                                                                                                                                                                                                                                                                      |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -m <num_maps>                     | Maximum number of concurrent copy tasks that can be executed at the same time                                                                                                                                                                                                                                    |
| -numListstatusTh-reads            | Number of threads for constituting the list of copied files. This option increases the running speed of DistCp.                                                                                                                                                                                                  |
| -overwrite                        | File at the target location that is to be overwritten                                                                                                                                                                                                                                                            |
| -update                           | A file at the target location is updated if the size and check of a file at the source location are different from those of the file at the target location.                                                                                                                                                     |
| -append                           | When <b>-update</b> is also used, the content of the file at the source location is added to the file at the target location.                                                                                                                                                                                    |
| -f <urilist_uri>                  | Content of the <urilist_uri> file is used as the file list to be copied.                                                                                                                                                                                                                                         |
| -filters                          | A local file is specified whose content contains multiple regular expressions. If the file to be copied matches a regular expression, the file is not copied.                                                                                                                                                    |
| -async                            | The <b>distcp</b> command is run asynchronously.                                                                                                                                                                                                                                                                 |
| -atomic {-tmp <tmp_dir>}          | An atomic copy can be performed. You can add a temporary directory during copying.                                                                                                                                                                                                                               |
| -bandwidth                        | The transmission bandwidth of each copy task. Unit: MB/s.                                                                                                                                                                                                                                                        |
| -delete                           | The files that exist in the target location is deleted but do not exist in the source location. This option is usually used with <b>-update</b> , and indicates that files at the source location are synchronized with those at the target location and the redundant files at the target location are deleted. |
| -diff <oldSnapshot> <newSnapshot> | The differences between the old and new versions are copied to a file in the old version at the target location.                                                                                                                                                                                                 |
| -skipcrccheck                     | Whether to skip the cyclic redundancy check (CRC) between the source file and the target file.                                                                                                                                                                                                                   |
| -strategy {dynamic uniformsize}   | The policy for copying a task. The default policy is <b>uniformsize</b> , that is, each copy task copies the same number of bytes.                                                                                                                                                                               |

## FAQs of DistCp

1. When you run the DistCp command, if the content of some copied files is large, you are advised to change the timeout period of MapReduce that executes the copy task. It can be implemented by specifying the **mapreduce.task.timeout** in the DistCp command. For example, run the following command to change the timeout to 30 minutes:

```
hadoop distcp -Dmapreduce.task.timeout=1800000 hdfs://cluster1/source hdfs://cluster2/target
```

Or, you can also use **filters** to exclude the large files out of the copy process. The command example is as follows:

```
hadoop distcp -filters /opt/client/filterfile hdfs://cluster1/source hdfs://cluster2/target
```

In the preceding command, *filterfile* indicates a local file, which contains multiple expressions used to match the path of a file that is not copied. The following is an example:

```
excludeFile1.
excludeFile2.
```

2. If the DistCp command unexpectedly quits, the error message "java.lang.OutOfMemoryError" is displayed.

This is because the memory required for running the copy command exceeds the preset memory limit (default value: 128 MB). You can change the memory upper limit of the client by modifying **CLIENT\_GC\_OPTS** in *<Client installation path>/HDFS/component\_env*. For example, if you want to set the memory upper limit to 1 GB, refer to the following configuration:

```
CLIENT_GC_OPTS="-Xmx1G"
```

After the modification, run the following command to make the modification take effect:

```
source {Client installation path}/bigdata_env
```

3. When the dynamic policy is used to run the DistCp command, the command exits unexpectedly and the error message "Too many chunks created with splitRatio" is displayed.

The cause of this problem is that the value of **distcp.dynamic.max.chunks.tolerable** (default value: 20,000) is less than the value of **distcp.dynamic.split.ratio** (default value: 2) multiplied by the number of Maps. This problem occurs when the number of Maps exceeds 10,000. You can use the **-m** parameter to reduce the number of Maps to less than 10,000.

```
hadoop distcp -strategy dynamic -m 9500 hdfs://cluster1/source hdfs://cluster2/target
```

Alternatively, you can use the **-D** parameter to set **distcp.dynamic.max.chunks.tolerable** to a large value.

```
hadoop distcp -Ddistcp.dynamic.max.chunks.tolerable=30000 -strategy dynamic hdfs://cluster1/source hdfs://cluster2/target
```

## 9.6 Overview of HDFS File System Directories

This section describes the directory structure in HDFS, as shown in the following table.

**Table 9-5** HDFS directory structure (applicable to versions earlier than MRS 3.x)

| Path                                               | Type                | Function                                                                                                                              | Whether the Directory Can Be Deleted | Deletion Consequence                         |
|----------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|----------------------------------------------|
| /tmp/spark/sparkhive-scratch                       | Fixed directory     | Stores temporary files of metastore sessions in Spark JDBCServer.                                                                     | No                                   | Failed to run the task.                      |
| /tmp/sparkhive-scratch                             | Fixed directory     | Stores temporary files of metastore session that are executed using Spark CLI.                                                        | No                                   | Failed to run the task.                      |
| /tmp/carbon/                                       | Fixed directory     | Stores the abnormal data in this directory if abnormal CarbonData data exists during data import.                                     | Yes                                  | Error data is lost.                          |
| /tmp/Loader- $\{Job\ name\}_\{MR\ job\ ID\}$       | Temporary directory | Stores the region information about Loader HBase bulkload jobs. The data is automatically deleted after the job running is completed. | No                                   | Failed to run the Loader HBase Bulkload job. |
| /tmp/logs                                          | Fixed directory     | Stores the collected MR task logs.                                                                                                    | Yes                                  | MR task logs are lost.                       |
| /tmp/archived                                      | Fixed directory     | Archives the MR task logs on HDFS.                                                                                                    | Yes                                  | MR task logs are lost.                       |
| /tmp/hadoop-yarn/staging                           | Fixed directory     | Stores the run logs, summary information, and configuration attributes of ApplicationMaster running jobs.                             | No                                   | Services are running improperly.             |
| /tmp/hadoop-yarn/staging/history/done_intermediate | Fixed directory     | Stores temporary files in the <b>/tmp/hadoop-yarn/staging</b> directory after all tasks are executed.                                 | No                                   | MR task logs are lost.                       |

| Path                                  | Type            | Function                                                                                                            | Whether the Directory Can Be Deleted | Deletion Consequence                                    |
|---------------------------------------|-----------------|---------------------------------------------------------------------------------------------------------------------|--------------------------------------|---------------------------------------------------------|
| /tmp/hadoop-yarn/staging/history/done | Fixed directory | The periodic scanning thread periodically moves the <b>done_intermediate</b> log file to the <b>done</b> directory. | No                                   | MR task logs are lost.                                  |
| /tmp/mr-history                       | Fixed directory | Stores the historical record files that are pre-loaded.                                                             | No                                   | Historical MR task log data is lost.                    |
| /tmp/hive                             | Fixed directory | Stores Hive temporary files.                                                                                        | No                                   | Failed to run the Hive task.                            |
| /tmp/hive-scratch                     | Fixed directory | Stores temporary data (such as session information) generated during Hive running.                                  | No                                   | Failed to run the current task.                         |
| /user/{user}/sparkStaging             | Fixed directory | Stores temporary files of the SparkJDBCServer application.                                                          | No                                   | Failed to start the executor.                           |
| /user/spark/jars                      | Fixed directory | Stores running dependency packages of the Spark executor.                                                           | No                                   | Failed to start the executor.                           |
| /user/loader                          | Fixed directory | Stores dirty data of Loader jobs and data of HBase jobs.                                                            | No                                   | Failed to execute the HBase job. Or dirty data is lost. |
| /user/loader/etl_dirty_data_dir       |                 |                                                                                                                     |                                      |                                                         |
| /user/loader/etl_hbase_putlist_tmp    |                 |                                                                                                                     |                                      |                                                         |
| /user/loader/etl_hbase_tmp            |                 |                                                                                                                     |                                      |                                                         |
| /user/mapred                          | Fixed directory | Stores Hadoop-related files.                                                                                        | No                                   | Failed to start Yarn.                                   |



| Path               | Type                | Function                                                                                                           | Whether the Directory Can Be Deleted | Deletion Consequence                                                          |
|--------------------|---------------------|--------------------------------------------------------------------------------------------------------------------|--------------------------------------|-------------------------------------------------------------------------------|
| /user/hive         | Fixed directory     | Stores Hive-related data by default, including the depended Spark lib package and default table data storage path. | No                                   | User data is lost.                                                            |
| /user/omm-bulkload | Temporary directory | Stores HBase batch import tools temporarily.                                                                       | No                                   | Failed to import HBase tasks in batches.                                      |
| /user/hbase        | Temporary directory | Stores HBase batch import tools temporarily.                                                                       | No                                   | Failed to import HBase tasks in batches.                                      |
| /sparkJobHistory   | Fixed directory     | Stores Spark event log data.                                                                                       | No                                   | The History Server service is unavailable, and the task fails to be executed. |
| /flume             | Fixed directory     | Stores data collected by Flume from HDFS.                                                                          | No                                   | Flume runs improperly.                                                        |
| /mr-history/tmp    | Fixed directory     | Stores logs generated by MapReduce jobs.                                                                           | Yes                                  | Log information is lost.                                                      |
| /mr-history/done   | Fixed directory     | Stores logs managed by MR JobHistory Server.                                                                       | Yes                                  | Log information is lost.                                                      |

| Path             | Type                            | Function                                                                                                                                                                                                                                                                                                                                                                                              | Whether the Directory Can Be Deleted | Deletion Consequence                                            |
|------------------|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|-----------------------------------------------------------------|
| /tenant          | Created when a tenant is added. | Directory of a tenant in the HDFS. By default, the system automatically creates a folder in the <b>/tenant</b> directory based on the tenant name. For example, the default HDFS storage directory for <b>ta1</b> is <b>tenant/ta1</b> . When a tenant is created for the first time, the system creates the <b>/tenant</b> directory in the HDFS root directory. You can customize the storage path. | No                                   | The tenant account is unavailable.                              |
| /apps{1~5}/      | Fixed directory                 | Stores the Hive package used by WebHCat.                                                                                                                                                                                                                                                                                                                                                              | No                                   | Failed to run the WebHCat tasks.                                |
| /hbase           | Fixed directory                 | Stores HBase data.                                                                                                                                                                                                                                                                                                                                                                                    | No                                   | HBase user data is lost.                                        |
| /hbaseFileStream | Fixed directory                 | Stores HFS files.                                                                                                                                                                                                                                                                                                                                                                                     | No                                   | The HFS file is lost and cannot be restored.                    |
| /ats/active      | Fixed directory                 | HDFS path used to store the timeline data of running applications.                                                                                                                                                                                                                                                                                                                                    | No                                   | Failed to run the <b>tez</b> task after the directory deletion. |
| /ats/done        | Fixed directory                 | HDFS path used to store the timeline data of completed applications.                                                                                                                                                                                                                                                                                                                                  | No                                   | Automatically created after the deletion.                       |
| /flink           | Fixed directory                 | Stores the checkpoint task data.                                                                                                                                                                                                                                                                                                                                                                      | No                                   | Failed to run tasks after the deletion.                         |

**Table 9-6** Directory structure of the HDFS file system (applicable to MRS 3.x or later)

| Path                                               | Type                | Function                                                                                                                              | Whether the Directory Can Be Deleted | Deletion Consequence                                           |
|----------------------------------------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|----------------------------------------------------------------|
| /tmp/spark2x/sparkhive-scratch                     | Fixed directory     | Stores temporary files of metastore session in Spark2x JDBCServer.                                                                    | No                                   | Failed to run the task.                                        |
| /tmp/sparkhive-scratch                             | Fixed directory     | Stores temporary files of metastore sessions that are executed in CLI mode using Spark2x CLI.                                         | No                                   | Failed to run the task.                                        |
| /tmp/logs/                                         | Fixed directory     | Stores container log files.                                                                                                           | Yes                                  | Container log files cannot be viewed.                          |
| /tmp/carbon/                                       | Fixed directory     | Stores the abnormal data in this directory if abnormal CarbonData data exists during data import.                                     | Yes                                  | Error data is lost.                                            |
| /tmp/Loader- <i>{Job name}_</i> <i>{MR job ID}</i> | Temporary directory | Stores the region information about Loader HBase bulkload jobs. The data is automatically deleted after the job running is completed. | No                                   | Failed to run the Loader HBase Bulkload job.                   |
| /tmp/hadoop-omm/yarn/system/rmstore                | Fixed directory     | Stores the ResourceManager running information.                                                                                       | Yes                                  | Status information is lost after ResourceManager is restarted. |
| /tmp/archived                                      | Fixed directory     | Archives the MR task logs on HDFS.                                                                                                    | Yes                                  | MR task logs are lost.                                         |
| /tmp/hadoop-yarn/staging                           | Fixed directory     | Stores the run logs, summary information, and configuration attributes of ApplicationMaster running jobs.                             | No                                   | Services are running improperly.                               |

| Path                                               | Type            | Function                                                                                                            | Whether the Directory Can Be Deleted | Deletion Consequence                                    |
|----------------------------------------------------|-----------------|---------------------------------------------------------------------------------------------------------------------|--------------------------------------|---------------------------------------------------------|
| /tmp/hadoop-yarn/staging/history/done_intermediate | Fixed directory | Stores temporary files in the <b>/tmp/hadoop-yarn/staging</b> directory after all tasks are executed.               | No                                   | MR task logs are lost.                                  |
| /tmp/hadoop-yarn/staging/history/done              | Fixed directory | The periodic scanning thread periodically moves the <b>done_intermediate</b> log file to the <b>done</b> directory. | No                                   | MR task logs are lost.                                  |
| /tmp/mr-history                                    | Fixed directory | Stores the historical record files that are pre-loaded.                                                             | No                                   | Historical MR task log data is lost.                    |
| /tmp/hive-scratch                                  | Fixed directory | Stores temporary data (such as session information) generated during Hive running.                                  | No                                   | Failed to run the current task.                         |
| /user/{user}/.spark Staging                        | Fixed directory | Stores temporary files of the SparkJDBCServer application.                                                          | No                                   | Failed to start the executor.                           |
| /user/spark2x/jars                                 | Fixed directory | Stores running dependency packages of the Spark2x executor.                                                         | No                                   | Failed to start the executor.                           |
| /user/loader                                       | Fixed directory | Stores dirty data of Loader jobs and data of HBase jobs.                                                            | No                                   | Failed to execute the HBase job. Or dirty data is lost. |
| /user/loader/etl_dirty_data_dir                    |                 |                                                                                                                     |                                      |                                                         |
| /user/loader/etl_hbase_pu tlist_tmp                |                 |                                                                                                                     |                                      |                                                         |
| /user/loader/etl_hbase_tm p                        |                 |                                                                                                                     |                                      |                                                         |

| Path                                       | Type                | Function                                                                                                           | Whether the Directory Can Be Deleted | Deletion Consequence                                                          |
|--------------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------|--------------------------------------|-------------------------------------------------------------------------------|
| /user/oozie                                | Fixed directory     | Stores dependent libraries required for Oozie running, which needs to be manually uploaded.                        | No                                   | Failed to schedule Oozie.                                                     |
| /user/mapred/hadoop-mapreduce-3.1.1.tar.gz | Fixed files         | Stores JAR files used by the distributed MR cache.                                                                 | No                                   | The MR distributed cache function is unavailable.                             |
| /user/hive                                 | Fixed directory     | Stores Hive-related data by default, including the depended Spark lib package and default table data storage path. | No                                   | User data is lost.                                                            |
| /user/omm-bulkload                         | Temporary directory | Stores HBase batch import tools temporarily.                                                                       | No                                   | Failed to import HBase tasks in batches.                                      |
| /user/hbase                                | Temporary directory | Stores HBase batch import tools temporarily.                                                                       | No                                   | Failed to import HBase tasks in batches.                                      |
| /spark2xJobHistory2x                       | Fixed directory     | Stores Spark2x eventlog data.                                                                                      | No                                   | The History Server service is unavailable, and the task fails to be executed. |
| /flume                                     | Fixed directory     | Stores data collected by Flume from HDFS.                                                                          | No                                   | Flume runs improperly.                                                        |
| /mr-history/tmp                            | Fixed directory     | Stores logs generated by MapReduce jobs.                                                                           | Yes                                  | Log information is lost.                                                      |
| /mr-history/done                           | Fixed directory     | Stores logs managed by MR JobHistory Server.                                                                       | Yes                                  | Log information is lost.                                                      |

| Path             | Type                            | Function                                                                                                                                                                                                                                                                                                                                                                                                | Whether the Directory Can Be Deleted | Deletion Consequence                         |
|------------------|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|----------------------------------------------|
| /tenant          | Created when a tenant is added. | Directory of a tenant in the HDFS. By default, the system automatically creates a folder in the / <b>tenant</b> directory based on the tenant name. For example, the default HDFS storage directory for <b>ta1</b> is <b>tenant/ta1</b> . When a tenant is created for the first time, the system creates the / <b>tenant</b> directory in the HDFS root directory. You can customize the storage path. | No                                   | The tenant account is unavailable.           |
| /apps{1~5}/      | Fixed directory                 | Stores the Hive package used by WebHCat.                                                                                                                                                                                                                                                                                                                                                                | No                                   | Failed to run the WebHCat tasks.             |
| /hbase           | Fixed directory                 | Stores HBase data.                                                                                                                                                                                                                                                                                                                                                                                      | No                                   | HBase user data is lost.                     |
| /hbaseFileStream | Fixed directory                 | Stores HFS files.                                                                                                                                                                                                                                                                                                                                                                                       | No                                   | The HFS file is lost and cannot be restored. |

## 9.7 Changing the DataNode Storage Directory

### Scenario

 NOTE

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

If the storage directory defined by the HDFS DataNode is incorrect or the HDFS storage plan changes, the MRS cluster administrator needs to modify the DataNode storage directory on FusionInsight Manager to ensure smooth HDFS running. Changing the ZooKeeper storage directory includes the following scenarios:

- Change the storage directory of the DataNode role. In this way, the storage directories of all DataNode instances are changed.

- Change the storage directory of a single DataNode instance. In this way, only the storage directory of this instance is changed, and the storage directories of other instances remain the same.

## Impact on the System

- The HDFS service needs to be stopped and restarted during the process of changing the storage directory of the DataNode role, and the cluster cannot provide services before it is completely started.
- The DataNode instance needs to be stopped and restarted during the process of changing the storage directory of the instance, and the instance at this node cannot provide services before it is started.
- The directory for storing service parameter configurations must also be updated.

## Prerequisites

- New disks have been prepared and installed on each data node, and the disks are formatted.
- New directories have been planned for storing data in the original directories.
- The HDFS client has been installed.
- The service user **hdfs** is available.
- When changing the storage directory of a single DataNode instance, ensure that the number of active DataNode instances is greater than the value of **dfs.replication**.

## Procedure

### Check the environment.

- Step 1** Log in to the server where the HDFS client is installed as user **root**, and run the following command to configure environment variables:

**source** *Installation directory of the HDFS client*/**bigdata\_env**

- Step 2** If the cluster is in security mode, run the following command to authenticate the user:

**kinit hdfs**

- Step 3** Run the following command on the HDFS client to check whether all directories and files in the HDFS root directory are normal:

**hdfs fsck /**

Check the fsck command output.

- If the following information is displayed, no file is lost or damaged. Go to [Step 4](#).  
The filesystem under path '/' is HEALTHY
- If other information is displayed, some files are lost or damaged. Go to [Step 5](#).

- Step 4** Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether **Running Status** of HDFS is **Normal**.

- If yes, go to [Step 6](#).
- If no, the HDFS status is unhealthy. Go to [Step 5](#).

**Step 5** Rectify the HDFS fault.. The task is complete.

**Step 6** Determine whether to change the storage directory of the DataNode role or that of a single DataNode instance:

- To change the storage directory of the DataNode role, go to [Step 7](#).
- To change the storage directory of a single DataNode instance, go to [Step 12](#).

### Changing the storage directory of the DataNode role

**Step 7** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Stop Instance** to stop the HDFS service.

**Step 8** Log in to each data node where the HDFS service is installed as user **root** and perform the following operations:

1. Create a target directory (**data1** and **data2** are original directories in the cluster).

For example, to create a target directory `${BIGDATA_DATA_HOME}/hadoop/data3/dn`, run the following command:

```
mkdir -p ${BIGDATA_DATA_HOME}/hadoop/data3/dn
```

2. Mount the target directory to the new disk. For example, mount `${BIGDATA_DATA_HOME}/hadoop/data3` to the new disk.

3. Modify permissions on the new directory.

For example, to create a target directory `${BIGDATA_DATA_HOME}/hadoop/data3/dn`, run the following commands:

```
chmod 700 ${BIGDATA_DATA_HOME}/hadoop/data3/dn -R and chown omm:wheel ${BIGDATA_DATA_HOME}/hadoop/data3/dn -R
```

4. Copy the data to the target directory.

For example, if the old directory is `${BIGDATA_DATA_HOME}/hadoop/data1/dn` and the target directory is `${BIGDATA_DATA_HOME}/hadoop/data3/dn`, run the following command:

```
cp -af ${BIGDATA_DATA_HOME}/hadoop/data1/dn/* ${BIGDATA_DATA_HOME}/hadoop/data3/dn
```

**Step 9** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations** to go to the HDFS service configuration page.

Change the value of `dfs.datanode.data.dir` from the default value `%{@auto.detect.datapart.dn}` to the new target directory, for example, `${BIGDATA_DATA_HOME}/hadoop/data3/dn`.

For example, the original data storage directories are `/srv/BigData/hadoop/data1`, `/srv/BigData/hadoop/data2`. To migrate data from the `/srv/BigData/hadoop/data1` directory to the newly created `/srv/BigData/hadoop/data3` directory, replace the whole parameter with `/srv/BigData/hadoop/data2`, `/srv/BigData/hadoop/data3`. Separate multiple storage directories with commas (.). In this example, changed directories are `/srv/BigData/hadoop/data2`, `/srv/BigData/hadoop/data3`.



- Step 10** Click **Save**. Choose **Cluster** > *Name of the desired cluster* > **Services**. On the page that is displayed, start the services that have been stopped.
- Step 11** After the HDFS is started, run the following command on the HDFS client to check whether all directories and files in the HDFS root directory are correctly copied:

```
hdfs fsck /
```

Check the fsck command output.

- If the following information is displayed, no file is lost or damaged, and data replication is successful. No further action is required.  
The filesystem under path '/' is HEALTHY
- If other information is displayed, some files are lost or damaged. In this case, check whether 8.4 is correct and run the `hdfs fsck Name of the damaged file -delete` command.

### Changing the storage directory of a single DataNode instance

- Step 12** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance**. Select the HDFS instance whose storage directory needs to be modified, and choose **More** > **Stop Instance**.
- Step 13** Log in to the DataNode node as user **root**, and perform the following operations:
1. Create a target directory.  
For example, to create a target directory `${BIGDATA_DATA_HOME}/hadoop/data3/dn`, run the following command:  

```
mkdir -p ${BIGDATA_DATA_HOME}/hadoop/data3/dn
```
  2. Mount the target directory to the new disk.  
For example, mount `${BIGDATA_DATA_HOME}/hadoop/data3` to the new disk.
  3. Modify permissions on the new directory.  
For example, to create a target directory `${BIGDATA_DATA_HOME}/hadoop/data3/dn`, run the following commands:  

```
chmod 700 ${BIGDATA_DATA_HOME}/hadoop/data3/dn -R and chown omm:wheel ${BIGDATA_DATA_HOME}/hadoop/data3/dn -R
```
  4. Copy the data to the target directory.  
For example, if the old directory is `${BIGDATA_DATA_HOME}/hadoop/data1/dn` and the target directory is `${BIGDATA_DATA_HOME}/hadoop/data3/dn`, run the following command:  

```
cp -af ${BIGDATA_DATA_HOME}/hadoop/data1/dn/* ${BIGDATA_DATA_HOME}/hadoop/data3/dn
```
- Step 14** On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **HDFS** > **Instance**. Click the specified DataNode instance and go to the **Configurations** page.

Change the value of `dfs.datanode.data.dir` from the default value `%{@auto.detect.datapart.dn}` to the new target directory, for example, `${BIGDATA_DATA_HOME}/hadoop/data3/dn`.

For example, the original data storage directories are `/srv/BigData/hadoop/data1`, `/srv/BigData/hadoop/data2`. To migrate data from the `/srv/BigData/`

**hadoop/data1** directory to the newly created **/srv/BigData/hadoop/data3** directory, replace the whole parameter with **/srv/BigData/hadoop/data2,/srv/BigData/hadoop/data3**.

**Step 15** Click **Save**, and then click **OK**.

**Operation succeeded** is displayed. click **Finish**.

**Step 16** Choose **More > Restart Instance** to restart the DataNode instance.

----End

## 9.8 Configuring HDFS Directory Permission

### Scenario

The permission for some HDFS directories is **777** or **750** by default, which brings potential security risks. You are advised to modify the permission for the HDFS directories after the HDFS is installed to increase user security.

### Procedure

Log in to the HDFS client as the administrator and run the following command to modify the permission for the **/user** directory.

The permission is set to **1777**, that is, **1** is added to the original permission. This indicates that only the user who creates the directory can delete it.

```
hdfs dfs -chmod 1777 /user
```

To ensure security of the system file, you are advised to harden the security for non-temporary directories. The following directories are examples:

- `/user:777`
- `/mr-history:777`
- `/mr-history/tmp:777`
- `/mr-history/done:777`
- `/user/mapred:755`

## 9.9 Configuring NFS

### Scenario

#### NOTE

This section applies to MRS 3.x or later.

Before deploying a cluster, you can deploy a Network File System (NFS) server based on requirements to store NameNode metadata to enhance data reliability.

If the NFS server has been deployed and NFS services are configured, you can follow operations in this section to configure NFS on the cluster. These operations are optional.

## Procedure

- Step 1** Check the permission of the shared NFS directories on the NFS server to ensure that the server can access NameNode in the MRS cluster.
- Step 2** Log in to the active NameNode as user **root**.
- Step 3** Run the following commands to create a directory and assign it write permissions:

```
mkdir ${BIGDATA_DATA_HOME}/namenode-nfs
chown omm:wheel ${BIGDATA_DATA_HOME}/namenode-nfs
chmod 750 ${BIGDATA_DATA_HOME}/namenode-nfs
```

- Step 4** Run the following command to mount the NFS to the active NameNode:

```
mount -t nfs -o rsize=8192,wsiz=8192,soft,nolock,timeo=3,intr IP address of
the NFS server.Shared directory ${BIGDATA_DATA_HOME}/namenode-nfs
```

For example, if the IP address of the NFS server is **192.168.0.11** and the shared directory is **/opt/Hadoop/NameNode**, run the following command:

```
mount -t nfs -o rsize=8192,wsiz=8192,soft,nolock,timeo=3,intr
192.168.0.11:/opt/Hadoop/NameNode ${BIGDATA_DATA_HOME}/namenode-
nfs
```

- Step 5** Perform [Step 2](#) to [Step 4](#) on the standby NameNode.

 **NOTE**

The names of the shared directories (for example, **/opt/Hadoop/NameNode**) created on the NFS server by the active and standby NameNodes must be different.

- Step 6** Log in to FusionInsight Manager, and choose **Cluster > Name of the desired cluster > Service > HDFS > Configuration > All Configurations**.

- Step 7** In the search box, search for **dfs.namenode.name.dir**, add **\${BIGDATA\_DATA\_HOME}/namenode-nfs** to **Value**, and click **Save**. Separate paths with commas (,).

- Step 8** Click **OK**. On the **Dashboard** tab page, choose **More > Restart Service** to restart the service.

----End

## 9.10 Planning HDFS Capacity

In HDFS, DataNode stores user files and directories as blocks, and file objects are generated on the NameNode to map each file, directory, and block on the DataNode.

The file objects on the NameNode require certain memory capacity. The memory consumption linearly increases as more file objects generated. The number of file objects on the NameNode increases and the objects consume more memory when the files and directories stored on the DataNode increase. In this case, the existing hardware may not meet the service requirement and the cluster is difficult to be scaled out.

Capacity planning of the HDFS that stores a large number of files is to plan the capacity specifications of the NameNode and DataNode and to set parameters according to the capacity plans.

## Capacity Specifications

- NameNode capacity specifications

Each file object on the NameNode corresponds to a file, directory, or block on the DataNode.

A file uses at least one block. The default size of a block is **134,217,728**, that is, 128 MB, which can be set in the **dfs.blocksize** parameter. By default, a file whose size is less than 128 MB occupies only one block. If the file size is greater than 128 MB, the number of occupied blocks is the file size divided by 128 MB (Number of occupied blocks = File size/128). The directories do not occupy any blocks.

Based on **dfs.blocksize**, the number of file objects on the NameNode is calculated as follows:

**Table 9-7** Number of NameNode file objects

| Size of a File                 | Number of File Objects                                  |
|--------------------------------|---------------------------------------------------------|
| < 128 MB                       | 1 (File) + 1 (Block) = 2                                |
| > 128 MB (for example, 128 GB) | 1 (File) + 1,024 (128 GB/128 MB = 1,024 blocks) = 1,025 |

The maximum number of file objects supported by the active and standby NameNodes is 300,000,000 (equivalent to 150,000,000 small files).

**dfs.namenode.max.objects** specifies the number of file objects that can be generated in the system. The default value is **0**, which indicates that the number of generated file objects is not limited.

- DataNode capacity specifications

In HDFS, blocks are stored on the DataNode as copies. The default number of copies is **3**, which can be set in the **dfs.replication** parameter.

The number of blocks stored on all DataNode role instances in the cluster can be calculated based on the following formula: Number of HDFS blocks x 3  
Average number of saved blocks = Number of HDFS blocks x 3/Number of DataNodes

**Table 9-8** DataNode specifications

| Item                                                                | Specifications |
|---------------------------------------------------------------------|----------------|
| Maximum number of blocks supported by a DataNode instance           | 5,000,000      |
| Maximum number of blocks supported by a disk on a DataNode instance | 500,000        |

| Item                                                                                                            | Specifications |
|-----------------------------------------------------------------------------------------------------------------|----------------|
| Minimum number of disks required when the number of blocks supported by a DataNode instance reaches the maximum | 10             |

**Table 9-9** Number of DataNodes

| Number of HDFS Blocks | Minimum Number of DataNode Roles   |
|-----------------------|------------------------------------|
| 10,000,000            | $10,000,000 * 3 / 5,000,000 = 6$   |
| 50,000,000            | $50,000,000 * 3 / 5,000,000 = 30$  |
| 100,000,000           | $100,000,000 * 3 / 5,000,000 = 60$ |

## Setting Memory Parameters

- 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=0x7FFFFFFFFFFFFFFE -
Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFFFFFFFFFE -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 9-10** 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   |
| 50,000,000             | -Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G   |

| Number of File Objects | Reference Value                                      |
|------------------------|------------------------------------------------------|
| 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 |

- 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=0x7FFFFFFFFFFFFFFE -
Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFFFFFFFFFE -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 average number of blocks stored in each DataNode instance in the cluster is: Number of HDFS blocks x 3/Number of DataNodes. If the average number of blocks changes, you need to change **-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M** in the default value. The following table lists the reference values.

**Table 9-11** 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.

## Viewing the HDFS Capacity Status

- NameNode information

For versions earlier than MRS 3.x: Log in to the MRS console, and choose **Components > HDFS > NameNode (Active)**. Click **Overview** and check the number of file objects, files, directories, or blocks in the HDFS in **Summary**.

For MRS 3.x or later: Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > HDFS > NameNode(Active)**, and click **Overview** to view information like the number of file objects, files, directories, and blocks in HDFS in **Summary** area.

- DataNode information

For versions earlier than MRS 3.x: Log in to the MRS console and choose **Components > HDFS > NameNode (Active)**. Click **DataNodes** and check the number of blocks of all DataNodes that report alarms.

For MRS 3.x or later: Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > HDFS > NameNode(Active)**, and click **DataNodes** to view the number of blocks on all DataNodes that report alarms.

- Alarm information

Check whether the alarms whose IDs are 14007, 14008, and 14009 are generated and change the alarm thresholds as required.

## 9.11 Configuring ulimit for HBase and HDFS

### Symptom

When an HDFS file is opened and the number of handles is limited, the following error occurs:

```
IOException (Too many open files)
```

### Procedure

You can contact the cluster administrator to increase the number of handles for each user. This is a configuration on the OS instead of that for HBase or HDFS. It is recommended that the cluster administrator set the number of handles based on the service volume of HBase and HDFS and the permissions of each user. If a user needs to frequently perform many operations on the HDFS with heavy service volume, set a large number of handles for the user to avoid the preceding error.

**Step 1** Log in to operating systems of all nodes or clients in the cluster as user **root** and go to the **/etc/security** directory.

**Step 2** Run the following command to edit the **limits.conf** file:

```
vi limits.conf
```

Added the following contents:

```
hdfs - nofile 32768
hbase - nofile 32768
```

**hdfs** and **hbase** indicate the OS user names used in services.

 NOTE

- Only user **root** has the permission to edit the **limits.conf** file.
- If the modification does not take effect, check whether there are other **nofile** values for the OS user in the **/etc/security/limits.d** directory. Such values may overwrite the value configured in **/etc/security/limits.conf**.
- If a user needs to perform operations on HBase, it is recommended that the handle count of the user be set to more than 10,000. If a user needs to perform operations on HDFS, it is recommended that the handle count be set based on the service volume. (A small value is not recommended.) If a user needs to perform operations on HBase and HDFS, it is recommended that the handle count be set to a large value, for example, 32,768.

**Step 3** You can run the following command to view the maximum number of handles of a user:

```
su - user_name
```

```
ulimit -n
```

The command output displays the maximum number of handles of the user. The following is an example.

```
8194
```

```
----End
```

## 9.12 Balancing DataNode Capacity

### Scenario

 NOTE


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

In the HDFS cluster, unbalanced disk usage among DataNodes may occur, for example, when new DataNodes are added to the cluster. Unbalanced disk usage may result in multiple problems. For example, MapReduce applications cannot make full use of local computing advantages, network bandwidth usage between data nodes cannot be optimal, or node disks cannot be used. Therefore, the MRS cluster administrator needs to periodically check and maintain DataNode data balance.


HDFS provides a capacity balancing program Balancer. By running Balancer, you can balance the HDFS cluster and ensure that the difference between the disk usage of each DataNode and that of the HDFS cluster does not exceed the threshold. DataNode disk usage before and after balancing is shown in [Figure 9-6](#) and [Figure 9-7](#), respectively.



**Figure 9-6** DataNode disk usage before balancing



**Figure 9-7** DataNode disk usage after balancing



The time of the balancing operation is affected by the following two factors:

1. Total amount of data to be migrated:  
The data volume of each DataNode must be greater than  $(\text{Average usage} - \text{Threshold}) \times \text{Average data volume}$  and less than  $(\text{Average usage} + \text{Threshold}) \times \text{Average data volume}$ . If the actual data volume is less than the minimum value or greater than the maximum value, imbalance occurs. The system sets the largest deviation volume on all DataNodes as the total data volume to be migrated.
2. Balancer migration is performed in sequence in iteration mode. The amount of data to be migrated in each iteration does not exceed 10 GB, and the usage of each iteration is recalculated.

Therefore, for a cluster, you can estimate the time consumed by each iteration (by observing the time consumed by each iteration recorded in balancer logs) and divide the total data volume by 10 GB to estimate the task execution time.

The balancer can be started or stopped at any time.

## Impact on the System

- The balance operation occupies network bandwidth resources of DataNodes. Perform the operation during maintenance based on service requirements.
- The balance operation may affect the running services if the bandwidth traffic (the default bandwidth control is 20 MB/s) is reset or the data volume is increased.

## Prerequisites

The client has been installed.

## Procedure

- Step 1** Log in to the node where the client is installed as a client installation user. Run the following command to switch to the client installation directory, for example, `/opt/client`:

```
cd /opt/client
```

 **NOTE**

If the cluster is in normal mode, run the `su - omm` command to switch to user `omm`.

- Step 2** Run the following command to configure environment variables:

```
source bigdata_env
```

- Step 3** If the cluster is in security mode, run the following command to authenticate the HDFS identity:

```
kinit hdfs
```

- Step 4** Determine whether to adjust the bandwidth control.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

- Step 5** Run the following command to change the maximum bandwidth of Balancer, and then go to [Step 6](#).

```
hdfs dfsadmin -setBalancerBandwidth <bandwidth in bytes per second>
```

*<bandwidth in bytes per second>* indicates the bandwidth control value, in bytes. For example, to set the bandwidth control to 20 MB/s (the corresponding value is 20971520), run the following command:

```
hdfs dfsadmin -setBalancerBandwidth 20971520
```

 **NOTE**

- The default bandwidth control is 20 MB/s. This value is applicable to the scenario where the current cluster uses the 10GE network and services are being executed. If the service idle time window is insufficient for balance maintenance, you can increase the value of this parameter to shorten the balance time, for example, to 209715200 (200 MB/s).
- The value of this parameter depends on the networking. If the cluster load is high, you can change the value to 209715200 (200 MB/s). If the cluster is idle, you can change the value to 1073741824 (1 GB/s).
- If the bandwidth of the DataNodes cannot reach the specified maximum bandwidth, modify the HDFS parameter `dfs.datanode.balance.max.concurrent.moves` on FusionInsight Manager, and change the number of threads for balancing on each DataNode to **32** and restart the HDFS service.

- Step 6** Run the following command to start the balance task:

```
bash /opt/client/HDFS/hadoop/sbin/start-balancer.sh -threshold <threshold of balancer>
```

**-threshold** specifies the deviation value of the DataNode disk usage, which is used for determining whether the HDFS data is balanced. When the difference between the disk usage of each DataNode and the average disk usage of the entire HDFS cluster is less than this threshold, the system considers that the HDFS cluster has been balanced and ends the balance task.

For example, to set deviation rate to 5%, run the following command:

```
bash /opt/client/HDFS/hadoop/sbin/start-balancer.sh -threshold 5
```

 **NOTE**

- The preceding command executes the task in the background. You can query related logs in the **hadoop-root-balancer-Host name.out** log file in the **/opt/client/HDFS/hadoop/logs** directory of the host.
- To stop the balance task, run the following command:  
**bash /opt/client/HDFS/hadoop/sbin/stop-balancer.sh**
- If only data on some nodes needs to be balanced, you can add the **-include** parameter in the script to specify the nodes to be migrated. You can run commands to view the usage of different parameters.  
For example, run the **bash /opt/client/HDFS/hadoop/sbin/start-balancer.sh -threshold 5 include IP address 1,IP address 2,IP address 3** command.
- **/opt/client** is the client installation directory. If the directory is inconsistent, replace it.
- If the command fails to be executed and the error information **Failed to APPEND\_FILE /system/balancer.id** is displayed in the log, run the following command to forcibly delete **/system/balancer.id** and run the **start-balancer.sh** script again:  
**hdfs dfs -rm -f /system/balancer.id**

**Step 7** After you run the script in [Step 6](#), the **hadoop-root-balancer-Host name.out** log file is generated in **/opt/client/HDFS/hadoop/logs**, the client installation directory. You can view the following information in the log:

- Time Stamp
- Bytes Already Moved
- Bytes Left To Move
- Bytes Being Moved

If message "Balance took xxx seconds" is displayed in the log, the balancing operation is complete.

----End

## Related Tasks

### Enable automatic execution of the balance task

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

**Step 2** Choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations**, select **All Configurations**, search for the following parameters, and change the parameter values.

- **dfs.balancer.auto.enable** indicates whether to enable automatic balance task execution. The default value **false** indicates that automatic balance task execution is disabled. The value **true** indicates that automatic execution is enabled.

- **dfs.balancer.auto.cron.expression** indicates the task execution time. The default value **0 1 \* \* 6** indicates that the task is executed at 01:00 every Saturday. This parameter is valid only when the automatic execution is enabled.

**Table 9-12** describes the expression for modifying this parameter. \* indicates consecutive time segments.

**Table 9-12** Parameters in the execution expression

| Column | Description                                                    |
|--------|----------------------------------------------------------------|
| 1      | Minute. The value ranges from 0 to 59.                         |
| 2      | Hour. The value ranges from 0 to 23.                           |
| 3      | Date. The value ranges from 1 to 31.                           |
| 4      | Month. The value ranges from 1 to 12.                          |
| 5      | Week. The value ranges from 0 to 6. <b>0</b> indicates Sunday. |

- **dfs.balancer.auto.stop.cron.expression** indicates the task ending time. The default value is empty, indicating that the running balance task is not automatically stopped. For example, **0 5 \* \* 6** indicates that the balance task is stopped at 05:00 every Saturday. This parameter is valid only when the automatic execution is enabled.

**Table 9-12** describes the expression for modifying this parameter. \* indicates consecutive time segments.

**Step 3** Running parameters of the balance task that is automatically executed are shown in **Table 9-13**.

**Table 9-13** Running parameters of the automatic balancer

| Parameter                           | Parameter description                                                                                                                                                        | Default Value                       |
|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| dfs.balancer.auto.threshold         | Specifies the balancing threshold of the disk capacity percentage. This parameter is valid only when <b>dfs.balancer.auto.enable</b> is set to <b>true</b> .                 | 10                                  |
| dfs.balancer.auto.exclude.datanodes | Specifies the list of DataNodes on which automatic disk balancing is not required. This parameter is valid only when <b>dfs.balancer.auto.enable</b> is set to <b>true</b> . | The value is left blank by default. |
| dfs.balancer.auto.bandwidthPerSec   | Specifies the maximum bandwidth (MB/s) of each DataNode for load balancing.                                                                                                  | 20                                  |

| Parameter                           | Parameter description                                                                                                                                                                                                                                                                                                                                                      | Default Value |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| dfs.balancer.auto.maxIdleIterations | Specifies the maximum number of consecutive idle iterations of Balancer. An idle iteration is an iteration without moving blocks. When the number of consecutive idle iterations reaches the maximum number, the balance task ends. The value <b>-1</b> indicates infinity.                                                                                                | 5             |
| dfs.balancer.auto.maxDataNodesNum   | Controls the number of DataNodes that perform automatic balance tasks. Assume that the value of this parameter is $N$ . If $N$ is greater than 0, data is balanced between $N$ DataNodes with the highest percentage of remaining space and $N$ DataNodes with the lowest percentage of remaining space. If $N$ is 0, data is balanced among all DataNodes in the cluster. | 5             |

**Step 4** Click **Save** to make configurations take effect. You do not need to restart the HDFS service.

Go to the `/var/log/Bigdata/hdfs/nn/hadoop-omm-balancer-Host name.log` file to view the task execution logs saved in the active NameNode.

----End

## 9.13 Configuring Replica Replacement Policy for Heterogeneous Capacity Among DataNodes

### Scenario

By default, NameNode randomly selects a DataNode to write files. If the disk capacity of some DataNodes in a cluster is inconsistent (the total disk capacity of some nodes is large and of some nodes is small), the nodes with small disk capacity will be fully written. To resolve this problem, change the default disk selection policy for data written to DataNode to the available space block policy. This policy increases the probability of writing data blocks to the node with large available disk space. This ensures that the node usage is balanced when disk capacity of DataNodes is inconsistent.

### Impact on the System

The disk selection policy is changed to `org.apache.hadoop.hdfs.server.blockmanagement.AvailableSpaceBlockPlacem`

**entPolicy.** It is proven that the HDFS file write performance optimizes by 3% after the modification.

 NOTE

**The default replica storage policy of the NameNode is as follows:**

1. First replica: stored on the node where the client resides.
2. Second replica: stored on DataNodes of the remote rack.
3. Third replica: stored on different nodes of the same rack for the node where the client resides.

If there are more replicas, randomly store them on other DataNodes.

The replica selection mechanism

(**org.apache.hadoop.hdfs.server.blockmanagement.AvailableSpaceBlockPlacementPolicy**) is as follows:

1. First replica: stored on the DataNode where the client resides (the same as the default storage policy).
2. Second replica:
  - When selecting a storage node, select two data nodes that meet the requirements.
  - Compare the disk usages of the two DataNodes. If the difference is smaller than 5%, store the replicas to the first node.
  - If the difference exceeds 5%, there is a 60% probability (specified by **dfs.namenode.available-space-block-placement-policy.balanced-space-preference-fraction** and default value is **0.6**) that the replica is written to the node whose disk space usage is low.
3. As for the storage of the third replica and subsequent replicas, refer to that of the second replica.

## Prerequisites

The total disk capacity deviation of DataNodes in the cluster cannot exceed 100%.

## Procedure

- Step 1** Go to the **All Configurations** page of HDFS by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** Modify the disk selection policy parameters when HDFS writes data. Search for the **dfs.block.replicator.classname** parameter and change its value to **org.apache.hadoop.hdfs.server.blockmanagement.AvailableSpaceBlockPlacementPolicy**.
- Step 3** Save the modified configuration. Restart the expired service or instance for the configuration to take effect.

----End

## 9.14 Configuring the Number of Files in a Single HDFS Directory

### Scenario

Generally, multiple services are deployed in a cluster, and the storage of most services depends on the HDFS file system. Different components such as Spark

and Yarn or clients are constantly writing files to the same HDFS directory when the cluster is running. However, the number of files in a single directory in HDFS is limited. Users must plan to prevent excessive files in a single directory and task failure.

You can set the number of files in a single directory using the **dfs.namenode.fs-limits.max-directory-items** parameter in HDFS.

## Procedure

- Step 1** Go to the **All Configurations** page of HDFS by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** Search for the configuration item **dfs.namenode.fs-limits.max-directory-items**.

**Table 9-14** Parameter description

| Parameter                                  | Description                                                           | Default Value |
|--------------------------------------------|-----------------------------------------------------------------------|---------------|
| dfs.namenode.fs-limits.max-directory-items | Maximum number of items in a directory<br>Value range: 1 to 6,400,000 | 1048576       |

- Step 3** Set the maximum number of files that can be stored in a single HDFS directory. Save the modified configuration. Restart the expired service or instance for the configuration to take effect.

 **NOTE**

Plan data storage in advance based on time and service type categories to prevent excessive files in a single directory. You are advised to use the default value, which is about 1 million pieces of data in a single directory.

----End

## 9.15 Configuring the Recycle Bin Mechanism

### Scenario

On HDFS, deleted files are moved to the recycle bin (trash can) so that the data deleted by mistake can be restored.

You can set the time threshold for storing files in the recycle bin. Once the file storage duration exceeds the threshold, it is permanently deleted from the recycle bin. If the recycle bin is cleared, all files in the recycle bin are permanently deleted.

### Configuration Description

If a file is deleted from HDFS, the file is saved in the trash space rather than cleared immediately. After the aging time is due, the deleted file becomes an aging file and will be cleared based on the system mechanism or manually cleared by users.

**Parameter portal:**

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-15** Parameter description

| Parameter                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Default Value |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| fs.trash.interval            | Trash collection time, in minutes. If data in the trash station exceeds the time, the data will be deleted. Value range: 1440 to 259200                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 1440          |
| fs.trash.checkpoint.interval | Interval between trash checkpoints, in minutes. The value must be less than or equal to the value of <b>fs.trash.interval</b> . The checkpoint program creates a checkpoint every time it runs and removes the checkpoint created <b>fs.trash.interval</b> minutes ago. For example, the system checks whether aging files exist every 10 minutes and deletes aging files if any. Files that are not aging are stored in the checkpoint list waiting for the next check.<br><br>If this parameter is set to 0, the system does not check aging files and all aging files are saved in the system.<br><br>Value range: 0 to <i>fs.trash.interval</i><br><br><b>NOTE</b><br>It is not recommended to set this parameter to 0 because aging files will use up the disk space of the cluster. | 60            |

## 9.16 Setting Permissions on Files and Directories

### Scenario

HDFS allows users to modify the default permissions of files and directories. The default mask provided by the HDFS for creating file and directory permissions is **022**. If you have special requirements for the default permissions, you can set configuration items to change the default permissions.

### Configuration Description

**Parameter portal:**

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).



**Table 9-16** Parameter description

| Parameter                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Default Value |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| fs.permissions.umask-mode | <p>This <b>umask</b> value (user mask) is used when the user creates files and directories in the HDFS on the clients. This parameter is similar to the file permission mask on Linux.</p> <p>The parameter value can be in octal or in symbolic, for example, <b>022</b> (octal, the same as <b>u=rwx,g=r-x,o=r-x</b> in symbolic), or <b>u=rwx,g=rwx,o=</b> (symbolic, the same as <b>007</b> in octal).</p> <p><b>NOTE</b><br/>The octal mask is opposite to the actual permission value. You are advised to use the symbol notation to make the description clearer.</p> | 022           |

## 9.17 Setting the Maximum Lifetime and Renewal Interval of a Token

### Scenario

In security mode, users can flexibly set the maximum token lifetime and token renewal interval in HDFS based on cluster requirements.

### Configuration Description

#### Navigation path for setting parameters:

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-17** Parameter description

| Parameter                                    | Description                                                                                                                                             | Default Value |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| dfs.namenode.delegation.token.max-lifetime   | This parameter is a server parameter. It specifies the maximum lifetime of a token. Unit: milliseconds. Value range: 10,000 to 10,000,000,000,000       | 604,800,000   |
| dfs.namenode.delegation.token.renew-interval | This parameter is a server parameter. It specifies the maximum lifetime to renew a token. Unit: milliseconds. Value range: 10,000 to 10,000,000,000,000 | 86,400,000    |

## 9.18 Configuring the Damaged Disk Volume

### Scenario

In the open source version, if multiple data storage volumes are configured for a DataNode, the DataNode stops providing services by default if one of the volumes is damaged. You can change the value of `dfs.datanode.failed.volumes.tolerated` to specify the number of damaged disk volumes that are allowed. If the number of damaged volumes does not exceed the threshold, DataNode continues to provide services.

### Configuration Description

**Navigation path for setting parameters:**

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-18** Parameter description

| Parameter                                          | Description                                                                                                                                                                                                                                                                                                                                            | Default Value                                            |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|
| <code>dfs.datanode.failed.volumes.tolerated</code> | Specifies the number of damaged volumes that are allowed before the DataNode stops providing services. By default, there must be at least one valid volume. The value <b>-1</b> indicates that the minimum value of a valid volume is <b>1</b> . The value greater than or equal to <b>0</b> indicates the number of damaged volumes that are allowed. | Versions earlier than MRS 3.x: 0<br>MRS 3.x or later: -1 |

## 9.19 Configuring Encrypted Channels

### Scenario

Encrypted channel is an encryption protocol of remote procedure call (RPC) in HDFS. When a user invokes RPC, the user's login name will be transmitted to RPC through RPC head. Then RPC uses Simple Authentication and Security Layer (SASL) to determine an authorization protocol (Kerberos and DIGEST-MD5) to complete RPC authorization. When you deploy a security cluster, use a secure encrypted channel and configure the following parameters:

### Configuration Description

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-19** Parameter description

| Parameter             | 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 again. Otherwise, the HDFS cannot provide the read and write services.</li> </ul> <p>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, RPC channels 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></p> <p>You can set this parameter on the HDFS component configuration page. The parameter setting takes effect globally, that is, the setting of whether the RPC channel is encrypted takes effect on all modules in Hadoop.</p> <p>There are three encryption modes.</p> <ul style="list-style-type: none"> <li><b>authentication:</b> This is the default value in normal mode. In this mode, data is directly transmitted without encryption after being authenticated. This mode ensures performance but has security risks.</li> <li><b>integrity:</b> Data is transmitted without encryption or authentication. To ensure data security, exercise caution when using this mode.</li> <li><b>privacy:</b> This is the default value in security mode, indicating that data is transmitted after authentication and encryption. This mode reduces the performance.</li> </ul> | <ul style="list-style-type: none"> <li>Security mode: privacy</li> <li>Normal mode: authentication</li> </ul> |

## 9.20 Reducing the Probability of Abnormal Client Application Operation When the Network Is Not Stable

### Scenario

Clients probably encounter running errors when the network is not stable. Users can adjust the following parameter values to improve the running efficiency.

### Configuration Description

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-20** Parameter description

| Parameter                                  | Description                                                                                                                                                                                                                                                  | Default Value |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| ha.health-monitor.rpc-timeout.ms           | Timeout interval during the NameNode health check performed by ZKFC. Increasing this value can prevent dual active NameNodes and reduce the probability of application running exceptions on clients.<br>Unit: millisecond. Value range: 30,000 to 3,600,000 | 180,000       |
| ipc.client.connect.max.retries.on.timeouts | Number of retry times when the socket connection between a server and a client times out.<br>Value range: 1 to 256                                                                                                                                           | 45            |
| ipc.client.connect.timeout                 | Timeout interval of the socket connection between a client and a server. Increasing the value of this parameter increases the timeout interval for setting up a connection.<br>Unit: millisecond. Value range: 1 to 3,600,000                                | 20,000        |

## 9.21 Configuring the NameNode Blacklist

### Scenario

 **NOTE**

This section applies to MRS 3.x or later.

In the existing default DFSclient failover proxy provider, if a NameNode in a process is faulty, all HDFS client instances in the same process attempt to connect

to the NameNode again. As a result, the application waits for a long time and timeout occurs.

When clients in the same JVM process connect to the NameNode that cannot be accessed, the system is overloaded. The NameNode blacklist is equipped with the MRS cluster to avoid this problem.

In the new Blacklisting DFSClient failover provider, the faulty NameNode is recorded in a list. The DFSClient then uses the information to prevent the client from connecting to such NameNodes again. This function is called NameNode blacklisting.

For example, there is a cluster with the following configurations:

```
namenode: nn1, nn2
```

```
dfs.client.failover.connection.retries: 20
```

Processes in a single JVM: 10 clients


In the preceding cluster, if the active **nn1** cannot be accessed, client1 will retry the connection for 20 times. Then, a failover occurs, and client1 will connect to **nn2**. In the same way, other clients also connect to **nn2** when the failover occurs after retrying the connection to **nn1** for 20 times. Such process prolongs the fault recovery of NameNode.

In this case, the NameNode blacklisting adds **nn1** to the blacklist when client1 attempts to connect to the active **nn1** which is already faulty. Therefore, other clients will avoid trying to connect to **nn1** but choose **nn2** directly.

#### NOTE

If, at any time, all NameNodes are added to the blacklist, the content in the blacklist will be cleared, and the client attempts to connect to the NameNodes based on the initial NameNode list. If any fault occurs again, the NameNode is still added to the blacklist.

**Figure 9-8** NameNode blacklisting working principle



## Configuration Description

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-21** NameNode blacklisting parameters

| Parameter                                                      | Description                                                                                                                                                                                                                                                                                | Default Value                                                           |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| dfs.client.failover.proxy.provider.<br><i>[nameservice ID]</i> | Client Failover proxy provider class which creates the NameNode proxy using the authenticated protocol.<br>Set this parameter to <b>org.apache.hadoop.hdfs.server.namenode.ha.BlackListingFailoverProxyProvider</b> .<br>You can configure the observer NameNode to process read requests. | org.apache.hadoop.hdfs.server.namenode.ha.AdaptiveFailoverProxyProvider |

## 9.22 Optimizing HDFS NameNode RPC QoS

### Scenarios


 NOTE

This section applies to MRS 3.x or later.

Several finished Hadoop clusters are faulty because the NameNode is overloaded and unresponsive.


Such problem is caused by the initial design of Hadoop: In Hadoop, the NameNode functions as an independent part and in its namespace coordinates various HDFS operations, including obtaining the data block location, listing directories, and creating files. The NameNode receives HDFS operations, regards them as RPC calls, and places them in the FIFO call queue for read threads to process. Requests in FIFO call queue are served first-in first-out. However, users who perform more I/O operations are served more time than those performing fewer I/O operations. In this case, the FIFO is unfair and causes the delay.

**Figure 9-9** NameNode request processing based on the FIFO call queue



The unfair problem and delaying mentioned before can be improved by replacing the FIFO queue with a new type of queue called FairCallQueue. In this way, FAIR queues assign incoming RPC calls to multiple queues based on the scale of the caller's call. The scheduling module tracks the latest calls and assigns a higher priority to users with a smaller number of calls.

**Figure 9-10** NameNode request processing based on FAIRCallQueue



### Configuration Description

- FairCallQueue ensures quality of service (QoS) by internally adjusting the order in which RPCs are invoked.

This queue consists of the following parts:

- a. DecayRpcScheduler: used to provide priority values from 0 to N (the value 0 indicates the highest priority).
- b. Multi-level queues (located in the FairCallQueue): used to ensure that queues are invoked in order of priority.
- c. Multi-channel converters (provided with Weighted Round Robin Multiplexer): used to provide logic control for queue selection.

After the FairCallQueue is configured, the control module determines the sub-queue to which the received invoking is allocated. The current scheduling module is DecayRpcScheduler, which only continuously tracks the priority numbers of various calls and periodically reduces these numbers.

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-22** FairCallQueue parameters

| Parameter                                    | Description                                                                                                                                 | Default Value                                         |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|
| <code>ipc.&lt;port&gt;.callqueue.impl</code> | Specifies the queue implementation class. You need to run the <b>org.apache.hadoop.ipc.FairCallQueue</b> command to enable the QoS feature. | <code>java.util.concurrent.LinkedBlockingQueue</code> |

- RPC BackOff

Backoff is one of the FairCallQueue functions. It requires the client to retry operations (such as creating, deleting, and opening a file) after a period of time. When the backoff occurs, the RCP server throws RetriableException. The FairCallQueue performs backoff in either of the following cases:

- The queue is full, that is, there are many client calls in the queue.
- The queue response time is longer than the threshold time (specified by the **ipc.<port>.decay-scheduler.backoff.responsetime.thresholds** parameter).



**Table 9-23** RPC Backoff configuration

| Parameter                                                                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Default Value           |
|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------|
| <code>ipc.&lt;port&gt;.backoff.enable</code>                                   | Specifies whether to enable the backoff. When the current application contains a large number of user callings, the RPC request is blocked if the connection limit of the operating system is not reached. Alternatively, when the RPC or NameNode is heavily loaded, some explicit exceptions can be thrown back to the client based on certain policies. The client can understand these exceptions and perform exponential rollback, which is another implementation of the <code>RetryInvocationHandler</code> class. | false                   |
| <code>ipc.&lt;port&gt;.decay-scheduler.backoff.response-time.enable</code>     | Indicate whether to enable the backoff based on the average queue response time.                                                                                                                                                                                                                                                                                                                                                                                                                                          | false                   |
| <code>ipc.&lt;port&gt;.decay-scheduler.backoff.response-time.thresholds</code> | Configure the response time threshold for each queue. The response time threshold must match the number of priorities (the value of <code>ipc.&lt;port&gt;.faircallqueue.priority-levels</code> ). Unit: millisecond                                                                                                                                                                                                                                                                                                      | 10000,20000,30000,40000 |

 **NOTE**

- `<port>` indicates the RPC port configured on the NameNode.
- The backoff function based on the response time takes effect only when `ipc.<port>.backoff.enable` is set to **true**.

## 9.23 Optimizing HDFS DataNode RPC QoS

### Scenario

When the speed at which the client writes data to the HDFS is greater than the disk bandwidth of the DataNode, the disk bandwidth is fully occupied. As a result, the DataNode does not respond. The client can back off only by canceling or restoring the channel, which results in write failures and unnecessary channel recovery operations.

#### NOTE

This section applies to MRS 3.x or later.

### Configuration

The new configuration parameter **dfs.pipeline.ecn** is introduced. When this configuration is enabled, the DataNode sends a signal from the write channel when the write channel is overloaded. The client may perform backoff based on the blocking signal to prevent the system from being overloaded. This configuration parameter is introduced to make the channel more stable and reduce unnecessary cancellation or recovery operations. After receiving the signal, the client backs off for a period of time (5,000 ms), and then adjusts the backoff time based on the related filter (the maximum backoff time is 50,000 ms).

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-24** DN ECN configuration

| Parameter        | Description                                                                      | Default Value |
|------------------|----------------------------------------------------------------------------------|---------------|
| dfs.pipeline.ecn | After configuration, the DataNode can send blocking notifications to the client. | false         |

## 9.24 Configuring Reserved Percentage of Disk Usage on DataNodes

### Scenario

When the Yarn local directory and DataNode directory are on the same disk, the disk with larger capacity can run more tasks. Therefore, more intermediate data is stored in the Yarn local directory.

Currently, you can set **dfs.datanode.du.reserved** to configure the absolute value of the reserved disk space on DataNodes. A small value cannot meet the

requirements of a disk with large capacity. However, configuring a large value for a disk with same capacity wastes a lot of disk space.

To avoid this problem, a new parameter **dfs.datanode.du.reserved.percentage** is introduced to configure the reserved percentage of the disk space.

 **NOTE**

- If **dfs.datanode.du.reserved.percentage** and **dfs.datanode.du.reserved** are configured at the same time, the larger value of the reserved disk space calculated using the two parameters is used as the reserved space of the data nodes.
- You are advised to set **dfs.datanode.du.reserved** or **dfs.datanode.du.reserved.percentage** based on the actual disk space.

## Configuration Description

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-25** Parameter description

| Parameter                           | Description                                                                                                                                                                                            | Default Value |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| dfs.datanode.du.reserved.percentage | Indicates the percentage of the reserved disk space on DataNodes. The DataNode permanently reserves the disk space calculated using this percentage.<br>The value is an integer ranging from 0 to 100. | 10            |

## 9.25 Configuring HDFS NodeLabel

### Scenario

You need to configure the nodes for storing HDFS file data blocks based on data features. You can configure a label expression to an HDFS directory or file and assign one or more labels to a DataNode so that file data blocks can be stored on specified DataNodes.

If the label-based data block placement policy is used for selecting DataNodes to store the specified files, the DataNode range is specified based on the label expression. Then proper nodes are selected from the specified range.

 **NOTE**

This section applies to MRS 3.x or later.

After cross-AZ HA is enabled for a single cluster, the HDFS NodeLabel function cannot be configured.

- Scenario 1: DataNodes partitioning scenario  
Scenario description:

When different application data is required to run on different nodes for separate management, label expressions can be used to achieve separation of different services, storing specified services on corresponding nodes.

By configuring the NodeLabel feature, you can perform the following operations:

- Store data in **/Hbase** to DN1, DN2, DN3, and DN4.
- Store data in **/Spark** to DN5, DN6, DN7, and DN8.

**Figure 9-11** DataNode partitioning scenario



**NOTE**


- Run the **hdfs nodelabel -setLabelExpression -expression 'LabelA[fallback=NONE]' -path /Hbase** command to set an expression for the **Hbase** directory. As shown in **Figure 9-11**, the data block replicas of files in the **/Hbase** directory are placed on the nodes labeled with the **LabelA**, that is, DN1, DN2, DN3, and DN4. Similarly, run the **hdfs nodelabel -setLabelExpression -expression 'LabelB[fallback=NONE]' -path /Spark** command to set an expression for the **Spark** directory. Data block replicas of files in the **/Spark** directory can be placed only on nodes labeled with **LabelB**, that is, DN5, DN6, DN7, and DN8.
  - For details about how to set labels for a data node, see [Configuration Description](#).
  - If multiple racks are available in one cluster, it is recommended that DataNodes of these racks should be available under each label, to ensure reliability of data block placement.
- Scenario 2: Specifying replica location when there are multiple racks

Scenario description:

In a heterogeneous cluster, customers need to allocate certain nodes with high availability to store important commercial data. Label expressions can be used to specify replica location so that the replica can be placed on a high reliable node.

Data blocks in the **/data** directory have three replicas by default. In this case, at least one replica is stored on a node of RACK1 or RACK2 (nodes of RACK1 and RACK2 are high reliable), and the other two are stored separately on the nodes of RACK3 and RACK4.

Figure 9-12 Scenario example



**NOTE**

Run the `hdfs nodelabel -setLabelExpression -expression 'LabelA|| LabelB[fallback=NONE],LabelC,LabelD' -path /data` command to set an expression for the `/data` directory.

When data is to be written to the `/data` directory, at least one data block replica is stored on a node labeled with the LabelA or LabelB, and the other two data block replicas are stored separately on the nodes labeled with the LabelC and LabelD.

### Configuration Description

- DataNode label configuration

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

Table 9-26 Parameter description

| Parameter                      | Description                                                                                                                                                                                     | Default Value                                                                    |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| dfs.block.replicator.classname | Used to configure the DataNode policy of HDFS. To enable the NodeLabel function, set this parameter to <b>org.apache.hadoop.hdfs.server.blockmanagement.BlockPlacementPolicyWithNodeLabel</b> . | org.apache.hadoop.hdfs.server.blockmanagement.AvailableSpaceBlockPlacementPolicy |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Default Value |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| host2tags | Used to configure a mapping between a DataNode host and a label.<br><br>The host name can be configured with an IP address extension expression (for example, <b>192.168.1.[1-128]</b> or <b>192.168.[2-3].[1-128]</b> ) or a regular expression (for example, <b>/datanode-[123]/</b> or <b>/datanode-\d{2}/</b> ) starting and ending with a slash (/). The label configuration name cannot contain the following characters: = / \ <b>Note:</b> The IP address must be a service IP address. | -             |

 NOTE

- The **host2tags** configuration item is described as follows:  
Assume there are 20 DataNodes which range from dn-1 to dn-20 in a cluster and the IP addresses of clusters range from 10.1.120.1 to 10.1.120.20. The value of **host2tags** can be represented in either of the following methods:  
**Regular expression of the host name**  
**/dn-\d/ = label-1** indicates that the labels corresponding to dn-1 to dn-9 are label-1, that is, dn-1 = label-1, dn-2 = label-1, ..., dn-9 = label-1.  
**/dn-((1[0-9]\$)|(20\$))/ = label-2** indicates that the labels corresponding to dn-10 to dn-20 are label-2, that is, dn-10 = label-2, dn-11 = label-2, ...dn-20 = label-2.  
**IP address range expression**  
**10.1.120.[1-9] = label-1** indicates that the labels corresponding to 10.1.120.1 to 10.1.120.9 are label-1, that is, 10.1.120.1 = label-1, 10.1.120.2 = label-1, ..., and 10.1.120.9 = label-1.  
**10.1.120.[10-20] = label-2** indicates that the labels corresponding to 10.1.120.10 to 10.1.120.20 are label-2, that is, 10.1.120.10 = label-2, 10.1.120.11 = label-2, ..., and 10.1.120.20 = label-2.
- Label-based data block placement policies are applicable to capacity expansion and reduction scenarios.  
A newly added DataNode will be assigned a label if the IP address of the DataNode is within the IP address range in the **host2tags** configuration item or the host name of the DataNode matches the host name regular expression in the **host2tags** configuration item.  
For example, the value of **host2tags** is **10.1.120.[1-9] = label-1**, but the current cluster has only three DataNodes: 10.1.120.1 to 10.1.120.3. If DataNode 10.1.120.4 is added for capacity expansion, the DataNode is labeled as label-1. If the 10.1.120.3 DataNode is deleted or out of the service, no data block will be allocated to the node.
- Set label expressions for directories or files.
  - On the HDFS parameter configuration page, configure **path2expression** to configure the mapping between HDFS directories and labels. If the configured HDFS directory does not exist, the configuration can succeed. When a directory with the same name as the HDFS directory is created manually, the configured label mapping relationship will be inherited by the directory within 30 minutes. After a labeled directory is deleted, a

- new directory with the same name as the deleted one will inherit its mapping within 30 minutes.
- For details about configuring items using commands, see the **hdfs nodelabel -setLabelExpression** command.
- To set label expressions using the Java API, invoke the **setLabelExpression(String src, String labelExpression)** method using the instantiated object `NodeLabelFileSystem`. *src* indicates a directory or file path on HDFS, and **labelExpression** indicates the label expression.
- After the `NodeLabel` is enabled, you can run the **hdfs nodelabel -listNodeLabels** command to view the label information of each `DataNode`.

## Block Replica Location Selection

`Nodelabel` supports different placement policies for replicas. The expression **label-1,label-2,label-3** indicates that three replicas are respectively placed in `DataNodes` containing label-1, label-2, and label-3. Different replica policies are separated by commas (,).

If you want to place two replicas in `DataNode` with label-1, set the expression as follows: **label-1 [replica=2],label-2,label-3**. In this case, if the default number of replicas is 3, two nodes with label-1 and one node with label-2 are selected. If the default number of replicas is 4, two nodes with label-1, one node with label-2, and one node with label-3 are selected. Note that the number of replicas is the same as that of each replica policy from left to right. However, the number of replicas sometimes exceeds the expressions. If the default number of replicas is 5, the extra replica is placed on the last node, that is, the node labeled with label-3.

When the ACLs function is enabled and the user does not have the permission to access the labels used in the expression, the `DataNode` with the label is not selected for the replica.

## Deletion of Redundant Block Replicas

If the number of block replicas exceeds the value of **dfs.replication** (number of file replicas specified by the user), HDFS will delete redundant block replicas to ensure cluster resource usage.

The deletion rules are as follows:

- Preferentially delete replicas that do not meet any expression.

For example: The default number of file replicas is **3**.

The label expression of `/test` is **LA[replica=1],LB[replica=1],LC[replica=1]**.

The file replicas of `/test` are distributed on four nodes (D1 to D4), corresponding to labels (LA to LD).

```
D1:LA
D2:LB
D3:LC
D4:LD
```

Then, block replicas on node D4 will be deleted.

- If all replicas meet the expressions, delete the redundant replicas which are beyond the number specified by the expression.
- For example: The default number of file replicas is **3**.

The label expression of **/test** is **LA[replica=1],LB[replica=1],LC[replica=1]**.

The file replicas of **/test** are distributed on the following four nodes, corresponding to the following labels.

```
D1:LA
D2:LA
D3:LB
D4:LC
```

Then, block replicas on node D1 or D2 will be deleted.

- If a file owner or group of a file owner cannot access a label, preferentially delete the replica from the DataNode mapped to the label.

## Example of label-based block placement policy

Assume that there are six DataNodes, namely, dn-1, dn-2, dn-3, dn-4, dn-5, and dn-6 in a cluster and the corresponding IP address range is 10.1.120.[1-6]. Six directories must be configured with label expressions. The default number of block replicas is **3**.

- The following provides three expressions of the DataNode label in **host2labels** file. The three expressions have the same function.
  - Regular expression of the host name
 

```
/dn-[1456]/ = label-1,label-2
/dn-[26]/ = label-1,label-3
/dn-[3456]/ = label-1,label-4
/dn-5/ = label-5
```
  - IP address range expression
 

```
10.1.120.[1-6] = label-1
10.1.120.1 = label-2
10.1.120.2 = label-3
10.1.120.[3-6] = label-4
10.1.120.[4-6] = label-2
10.1.120.5 = label-5
10.1.120.6 = label-3
```
  - Common host name expression
 

```
/dn-1/ = label-1, label-2
/dn-2/ = label-1, label-3
/dn-3/ = label-1, label-4
/dn-4/ = label-1, label-2, label-4
/dn-5/ = label-1, label-2, label-4, label-5
/dn-6/ = label-1, label-2, label-3, label-4
```
- The label expressions of the directories are set as follows:
 

```
/dir1 = label-1
/dir2 = label-1 && label-3
/dir3 = label-2 || label-4[replica=2]
/dir4 = (label-2 || label-3) && label-4
/dir5 = !label-1
/sdir2.txt = label-1 && label-3[replica=3,fallback=NONE]
/dir6 = label-4[replica=2],label-2
```

### NOTE

For details about the label expression configuration, see the **hdfs nodelabel - setLabelExpression** command.

The file data block storage locations are as follows:

- Data blocks of files in the **/dir1** directory can be stored on any of the following nodes: dn-1, dn-2, dn-3, dn-4, dn-5, and dn-6.
- Data blocks of files in the **/dir2** directory can be stored on the dn-2 and dn-6 nodes. The default number of block replicas is **3**. The expression



matches only two DataNodes. The third replica will be stored on one of the remaining nodes in the cluster.

- Data blocks of files in the **/dir3** directory can be stored on any three of the following nodes: dn-1, dn-3, dn-4, dn-5, and dn-6.
- Data blocks of files in the **/dir4** directory can be stored on the dn-4, dn-5, and dn-6 nodes.
- Data blocks of files in the **/dir5** directory do not match any DataNode and will be stored on any three nodes in the cluster, which is the same as the default block selection policy.
- For the data blocks of the **/sdir2.txt** file, two replicas are stored on the dn-2 and dn-6 nodes. The left one is not stored in the node because **fallback=NONE** is enabled.
- Data blocks of the files in the **/dir6** directory are stored on the two nodes with label-4 selected from dn-3, dn-4, dn-5, and dn-6 and another node with label-2. If the specified number of file replicas in the **/dir6** directory is more than 3, the extra replicas will be stored on a node with label-2.

## Restrictions

In configuration files, **key** and **value** are separated by equation signs (=), colons (:), and whitespace. Therefore, the host name of the **key** cannot contain these characters because these characters may be considered as separators.

## 9.26 Configuring HDFS Mover

### Scenario

Mover is a new data migration tool whose working mode is similar to that of the HDFS Balancer. Mover can redistribute data in the cluster based on the configured data storage policy.

Use Mover to periodically check whether the specified HDFS file or directory in the HDFS file system meets the preset storage policy. If not, migrate data to make them meet the policy.

#### NOTE

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

### Configuration Description

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-27** Parameter description

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                    | Default Value |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| dfs.mover.auto.enable            | Specifies whether to enable the data replica migration function. This function supports multiple modes. The default value is <b>false</b> , indicating that this function is disabled.                                                                                                                                                                                         | false         |
| dfs.mover.auto.cron.expression   | Specifies the CRON expression for HDFS automatic data migration, and is used to control the start time of data migration. This parameter is valid only when <b>dfs.mover.auto.enable</b> is set to <b>true</b> . The default value is <b>0 * * * *</b> , indicating that the task is executed on the hour. For details about CRON expression, see <a href="#">Table 9-28</a> . | 0 * * * *     |
| dfs.mover.auto.hdfsfiles_or_dirs | Specifies HDFS file and directory lists that implement automatic replica migration in specified clusters. Multiple values are separated by space. This parameter is valid only when <b>dfs.mover.auto.enable</b> is set to <b>true</b> .                                                                                                                                       | -             |

**Table 9-28** CRON expressions

| Column | Description                                                    |
|--------|----------------------------------------------------------------|
| 1      | Minute. The value ranges from 0 to 59.                         |
| 2      | Hour. The value ranges from 0 to 23.                           |
| 3      | Date. The value ranges from 1 to 31.                           |
| 4      | Month. The value ranges from 1 to 12.                          |
| 5      | Week. The value ranges from 0 to 6. <b>0</b> indicates Sunday. |

## Use Restrictions

Run the command on the HDFS client to enable the mover function. The command format is as follows:

**hdfs mover -p** *<Full path or directory path of an HDFS file >*

### NOTE

Users running this command on the client must have the **supergroup** permission. You can use the system user **hdfs** of the HDFS service. Alternatively, you can create a user with the **supergroup** permission in the cluster and then run the command.

## 9.27 Using HDFS AZ Mover

### Scenario

AZ Mover is a copy migration tool used to move copies to meet the new AZ policies set on the directory. It can be used to migrate copies from one AZ policy to another. AZ Mover instructs NameNode to move copies based on a new AZ policy. If the NameNode refuses to delete the old copies, the new policy may not be met. For example, the copies are marked as outdated.

### Restrictions

- Changing the policy name to **LOCAL\_AZ** is the same as that to **ONE\_AZ** because the client location cannot be determined when the uploaded file is written.
- Mover cannot determine the AZ status. As a result, the copy may be moved to the abnormal AZ and depends on NameNode for further processing.
- Mover depends on whether the number of DataNodes in each AZ meets the minimum requirement. If the AZ Mover is executed in an AZ with a small number of DataNodes, the result may be different from the expected result.
- Mover only meets the AZ-level policies and does not guarantee to meet the basic block placement policy (BPP).
- Mover does not support the change of replication factors. If the number of copies in the new AZ is different from that in the old AZ, an exception occurs.

### Procedure

**Step 1** Run the following command to go to the client installation directory. For example, if the client installation directory is **/opt/client**, run the following command:

```
cd /opt/client
```

**Step 2** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 3** If the cluster is in security mode, the user must have the read permission on the source directory or file and the write permission on the destination directory, and run the following command to authenticate the user: In normal mode, skip user authentication.

```
kinit Component service user
```

**Step 4** Create a directory and set an AZ policy.

Run the following command to create a directory.

```
hdfs dfs -mkdir <path>
```

Run the following command to set the AZ policy (**azexpression** indicates the AZ policy):

```
hdfs dfsadmin -setAZExpression <path> <azexpression>
```

Run the following command to view the AZ policy:

```
hdfs dfsadmin -getAZExpression <path>
```

**Step 5** Upload files to the directory.

```
hdfs dfs -put <localfile> <hdfs-path>
```

**Step 6** Delete the old policy from the directory and set a new policy.

Run the following command to clear the old policy:

```
hdfs dfsadmin -clearAZExpression <path>
```

Run the following command to configure a new policy:

```
hdfs dfsadmin -setAZExpression <path> <azexpression>
```

**Step 7** Run the **azmover** command to make the copy distribution meet the new AZ policy.

```
hdfs azmover -p /targetDirecotry
```

----End

## 9.28 Configuring HDFS DiskBalancer

### Scenario

DiskBalancer is an online disk balancer that balances disk data on running DataNodes based on various indicators. It works in the similar way of the HDFS Balancer. The difference is that HDFS Balancer balances data between DataNodes, while HDFS DiskBalancer balances data among disks on a single DataNode.

Data among disks may be unevenly distributed if a large number of files have been deleted from a cluster running for a long time, or disk capacity expansion is performed on a node in the cluster. Uneven data distribution may deteriorate the concurrent read/write performance of the HDFS, or cause service failure due to inappropriate HDFS write policies. In this case, the data density among disks on a node needs to be balanced to prevent heterogeneous small disks from becoming the performance bottleneck of the node.

#### NOTE

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

### Configuration Description

Go to the **All Configurations** page of HDFS and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 9-29** Parameter description

| Parameter                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Default Value |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| dfs.disk.balancer.auto.enabled                  | Indicates whether to enable the HDFS DiskBalancer function. The default value is <b>false</b> , indicating that this function is disabled.                                                                                                                                                                                                                                                                                                                                                     | false         |
| dfs.disk.balancer.auto.cron.expression          | CRON expression of the HDFS disk balancing operation, which is used to control the start time of the balancing operation. This parameter is valid only when <b>dfs.disk.balancer.auto.enabled</b> is set to <b>true</b> . The default value is <b>0 1 * * 6</b> , indicating that tasks are executed at 01:00 every Saturday. For details about cron expression, see <a href="#">Table 9-30</a> . The default value indicates that the DiskBalancer check is executed at 01:00 every Saturday. | 0 1 * * 6     |
| dfs.disk.balancer.max.disk.throughputInMBperSec | Specifies the maximum disk bandwidth that can be used for disk data balancing. The unit is MB/s, and the default value is <b>10</b> . Set this parameter based on the actual disk conditions of the cluster.                                                                                                                                                                                                                                                                                   | 10            |
| dfs.disk.balancer.max.disk.errors               | Specifies the maximum number of errors that are allowed in a specified movement process. If the value exceeds this threshold, the movement fails.                                                                                                                                                                                                                                                                                                                                              | 5             |
| dfs.disk.balancer.block.tolerance.percent       | Specifies the difference threshold between the data storage capacity and perfect status of each disk during data balancing among disks. For example, the ideal data storage capacity of each disk is 1 TB, and this parameter is set to <b>10</b> . When the data storage capacity of the target disk reaches 900 GB, the storage status of the disk is considered as perfect. Value range: 1 to 100.                                                                                          | 10            |
| dfs.disk.balancer.plan.threshold.percent        | Specifies the data density difference that is allowed between two disks during disk data balancing. If the absolute value of the data density difference between any two disks exceeds the threshold, data balancing is required. Value range: 1 to 100.                                                                                                                                                                                                                                       | 10            |
| dfs.disk.balancer.top.nodes.number              | Specifies the top <i>N</i> nodes whose disk data needs to be balanced in the cluster.                                                                                                                                                                                                                                                                                                                                                                                                          | 5             |

To use this function, set **dfs.disk.balancer.auto.enabled** to **true** and configure a proper CRON expression. Set other parameters based on the cluster status.

**Table 9-30** CRON expressions

| Column | Description                                                    |
|--------|----------------------------------------------------------------|
| 1      | Minute. The value ranges from 0 to 59.                         |
| 2      | Hour. The value ranges from 0 to 23.                           |
| 3      | Date. The value ranges from 1 to 31.                           |
| 4      | Month. The value ranges from 1 to 12.                          |
| 5      | Week. The value ranges from 0 to 6. <b>0</b> indicates Sunday. |

## Use Restrictions

1. Data can only be moved between disks of the same type. For example, data can only be moved between SSDs or between DISKS.
2. Enabling this function occupies disk I/O resources and network bandwidth resources of involved nodes. Enable this function in off-peak hours.
3. The DataNodes specified by the **dfs.disk.balancer.top.nodes.number** parameter are frequently calculated. Therefore, set the parameter to a small value.
4. Commands for using the DiskBalancer function on the HDFS client are as follows:

**Table 9-31** DiskBalancer commands

| Syntax                                                                         | Description                                                                                                                                                                                                                                                        |
|--------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>hdfs diskbalancer -report -top &lt;N&gt;</code>                          | Set <i>N</i> to an integer greater than 0. This command can be used to query the top <i>N</i> nodes that require disk data balancing in the cluster.                                                                                                               |
| <code>hdfs diskbalancer -plan &lt;Hostname  IP Address&gt;</code>              | This command can be used to generate a JSON file based on the DataNode. The file contains information about the source disk, target disk, and blocks to be moved. In addition, this command can be used to specify other parameters such as the network bandwidth. |
| <code>hdfs diskbalancer -query &lt;Hostname:\$dfs.datanode.ipc.port&gt;</code> | The default port number of the cluster is 9867. This command is used to query the running status of the DiskBalancer task on the current node.                                                                                                                     |

| Syntax                                                   | Description                                                                                                      |
|----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <code>hdfs diskbalancer -execute &lt;planfile&gt;</code> | In this command, <b>planfile</b> indicates the JSON file generated in the second command. Use the absolute path. |
| <code>hdfs diskbalancer -cancel &lt;planfile&gt;</code>  | This command is used to cancel the running planfile. Use the absolute path.                                      |

 **NOTE**

- Users running this command on the client must have the **supergroup** permission. You can use the system user **hdfs** of the HDFS service. Alternatively, you can create a user with the **supergroup** permission in the cluster and then run the command.
- Only formats and usage of commands are provided in [Table 9-31](#). For more parameters to be configured for each command, run the **hdfs diskbalancer -help <command>** command to view detailed information.
- When you troubleshoot performance problems during the cluster O&M, check whether the HDFS disk balancing occurs in the event information of the cluster. If yes, check whether DiskBalancer is enabled in the cluster.
- After the automatic DiskBalancer function is enabled, the ongoing task stops only after the current data balancing is complete. The task cannot be canceled during the balancing.
- You can manually specify certain nodes for data balancing on the client.

## 9.29 Configuring the Observer NameNode to Process Read Requests

### Scenario

In an HDFS cluster configured with HA, the active NameNode processes all client requests, and the standby NameNode reserves the latest metadata and block location information. However, in this architecture, the active NameNode is the bottleneck of client request processing. This bottleneck is more obvious in clusters with a large number of requests.

To address this issue, a new NameNode is introduced: an observer NameNode. Similar to the standby NameNode, the observer NameNode also reserves the latest metadata information and block location information. In addition, the observer NameNode can process read requests from clients in the same way as the active NameNode. In typical HDFS clusters with many read requests, the observer NameNode can be used to process read requests, reducing the active NameNode load and improving the cluster capability of processing requests.

 **NOTE**

This section applies to MRS 3.x or later.

## Impact on the System

- The active NameNode load can be reduced and the capability of HDFS cluster processing requests can be improved, which is especially obvious for large clusters.
- The client application configuration needs to be updated.

## Prerequisites

- The HDFS cluster has been installed, the active and standby NameNodes are running properly, and the HDFS service is normal.
- The `/${BIGDATA_DATA_HOME}/namenode` partition has been created on the node where the observer NameNode is to be installed.

## Procedure

The following steps describe how to configure the observer NameNode of a hacluster and enable it to process read requests. If there are multiple pairs of NameServices in the cluster and they are all in use, perform the following steps to configure the observer NameNode for each pair.

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

**Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **NameService Management**.

**Step 3** Click **Add** next to **hacluster**.

**Step 4** On the **Add NameNode** page, set **NameNode type** to **Observer** and click **Next**.

**Step 5** On the **Assign Role** page, select the planned host, add the observer NameNode, and click **Next**.

### NOTE

A maximum of five observer NameNodes can be added to each pair of NameServices.

**Step 6** On the configuration page, configure the storage directory and port number of the NameNode as planned and click **Next**.

**Step 7** Confirm the information, click **Submit**, and wait until the installation of the observer NameNode is complete.

**Step 8** Restart the upper-layer components that depend on HDFS, update the client application configuration, and restart the client application.

----End

## 9.30 Performing Concurrent Operations on HDFS Files

### Scenario

Performing this operation can concurrently modify file and directory permissions and access control tools in a cluster.



 **NOTE**

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

## Impact on the System

Performing concurrent file modification operations in a cluster has adverse impacts on the cluster performance. Therefore, you are advised to do so when the cluster is idle.

## Prerequisites

- The HDFS client or clients including HDFS has been installed. For example, the installation directory is **/opt/client**.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user needs to change the password upon the first login. (This operation is not required in normal mode.)

## Procedure

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** If the cluster is in security mode, the user executing the DistCp command must belong to the **supergroup** group and run the following command to perform user authentication. In normal mode, user authentication is not required.

```
kinit Component service user
```

**Step 5** Increase the JVM size of the client to prevent out of memory (OOM). (32 GB is recommended for 100 million files.)

 **NOTE**

The HDFS client exits abnormally and the error message "java.lang.OutOfMemoryError" is displayed after the HDFS client command is executed.

This problem occurs because the memory required for running the HDFS client exceeds the preset upper limit (128 MB by default). You can change the memory upper limit of the client by modifying **CLIENT\_GC\_OPTS** in *<Client installation path>/HDFS/component\_env*. For example, if you want to set the upper limit to 1 GB, run the following command:

```
CLIENT_GC_OPTS="-Xmx1G"
```

After the modification, run the following command to make the modification take effect:

```
source <Client installation path>/bigdata_env
```

**Step 6** Run the concurrent commands shown in the following table.

| Command                                                                                                                                      | Description                                                                                                                                                                                                                                                                                                                                                                                           | Function                                                                   |
|----------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| hdfs quickcmds [-t threadsNumber] [-p principal] [-k keytab] -setrep <rep> <path> ...                                                        | <p><b>threadsNumber</b> indicates the number of concurrent threads. The default value is the number of vCPUs of the local host.</p> <p><b>principal</b> indicates the Kerberos user.</p> <p><b>keytab</b> indicates the Keytab file.</p> <p><b>rep</b> indicates the number of replicas.</p> <p><b>path</b> indicates the HDFS directory.</p>                                                         | Used to concurrently set the number of copies of all files in a directory. |
| hdfs quickcmds [-t threadsNumber] [-p principal] [-k keytab] -chown [owner][: [group]] <path> ...                                            | <p><b>threadsNumber</b> indicates the number of concurrent threads. The default value is the number of vCPUs of the local host.</p> <p><b>principal</b> indicates the Kerberos user.</p> <p><b>keytab</b> indicates the Keytab file.</p> <p><b>owner</b> indicates the owner.</p> <p><b>group</b> indicates the group to which the user belongs.</p> <p><b>path</b> indicates the HDFS directory.</p> | Used to concurrently set the owner group of all files in the directory.    |
| hdfs quickcmds [-t threadsNumber] [-p principal] [-k keytab] -chmod <mode> <path> ...                                                        | <p><b>threadsNumber</b> indicates the number of concurrent threads. The default value is the number of vCPUs of the local host.</p> <p><b>principal</b> indicates the Kerberos user.</p> <p><b>keytab</b> indicates the Keytab file.</p> <p><b>mode</b> indicates the permission (for example, 754).</p> <p><b>path</b> indicates the HDFS directory.</p>                                             | Used to concurrently set permissions for all files in a directory.         |
| hdfs quickcmds [-t threadsNumber] [-p principal] [-k keytab] -setfacl [{-b -k} {-m -x} <acl_spec>] <path> ...] [--set <acl_spec> <path> ...] | <p><b>threadsNumber</b> indicates the number of concurrent threads. The default value is the number of vCPUs of the local host.</p> <p><b>principal</b> indicates the Kerberos user.</p> <p><b>keytab</b> indicates the Keytab file.</p> <p><b>acl_spec</b> indicates the ACL list separated by commas (,).</p> <p><b>path</b> indicates the HDFS directory.</p>                                      | Used to concurrently set ACL information for all files in a directory.     |

----End

## 9.31 Introduction to HDFS Logs

### Log Description

**Log path:** The default path of HDFS logs is `/var/log/Bigdata/hdfs/Role name`.

- NameNode: `/var/log/Bigdata/hdfs/nn` (run logs) and `/var/log/Bigdata/audit/hdfs/nn` (audit logs)
- DataNode: `/var/log/Bigdata/hdfs/dn` (run logs) and `/var/log/Bigdata/audit/hdfs/dn` (audit logs)
- ZKFC: `/var/log/Bigdata/hdfs/zkfc` (run logs) and `/var/log/Bigdata/audit/hdfs/zkfc` (audit logs)
- JournalNode: `/var/log/Bigdata/hdfs/jn` (run logs) and `/var/log/Bigdata/audit/hdfs/jn` (audit logs)
- Router: `/var/log/Bigdata/hdfs/router` (run logs) and `/var/log/Bigdata/audit/hdfs/router` (audit logs)
- HttpFS: `/var/log/Bigdata/hdfs/httpfs` (run logs) and `/var/log/Bigdata/audit/hdfs/httpfs` (audit logs)

**Log archive rule:** The automatic HDFS log compression function is enabled. By default, when the size of logs exceeds 100 MB, logs are automatically compressed into a log file named in the following format: `<Original log file name>-<yyyy-mm-dd_hh-mm-ss.[ID].log.zip`. A maximum of 100 latest compressed files are reserved. The number of compressed files can be configured on Manager.

**Table 9-32** HDFS log list

| Type    | Name                                            | Description                                                                                |
|---------|-------------------------------------------------|--------------------------------------------------------------------------------------------|
| Run log | hadoop-<SSH_USER>-<process_name>-<hostname>.log | HDFS system log, which records most of the logs generated when the HDFS system is running. |
|         | hadoop-<SSH_USER>-<process_name>-<hostname>.out | Log that records the HDFS running environment information.                                 |
|         | hadoop.log                                      | Log that records the operation of the Hadoop client.                                       |

| Type | Name                                          | Description                                                                                                                                             |
|------|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
|      | hdfs-period-check.log                         | Log that records scripts that are executed periodically, including automatic balancing, data migration, and JournalNode data synchronization detection. |
|      | <process_name>-<SSH_USER>-<DATE>-<PID>-gc.log | Garbage collection log file                                                                                                                             |
|      | postinstallDetail.log                         | Work log before the HDFS service startup and after the installation.                                                                                    |
|      | hdfs-service-check.log                        | Log that records whether the HDFS service starts successfully.                                                                                          |
|      | hdfs-set-storage-policy.log                   | Log that records the HDFS data storage policies.                                                                                                        |
|      | cleanupDetail.log                             | Log that records the cleanup logs about the uninstallation of the HDFS service.                                                                         |
|      | prestartDetail.log                            | Log that records cluster operations before the HDFS service startup.                                                                                    |
|      | hdfs-recover-fsimage.log                      | Recovery log of the NameNode metadata.                                                                                                                  |
|      | datanode-disk-check.log                       | Log that records the disk status check during the cluster installation and use.                                                                         |
|      | hdfs-availability-check.log                   | Log that check whether the HDFS service is available.                                                                                                   |
|      | hdfs-backup-fsimage.log                       | Backup log of the NameNode metadata.                                                                                                                    |
|      | startDetail.log                               | Detailed log that records the HDFS service startup.                                                                                                     |

| Type        | Name                                                                                                                                                               | Description                                                                                            |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
|             | hdfs-blockplacement.log                                                                                                                                            | Log that records the placement policy of HDFS blocks.                                                  |
|             | upgradeDetail.log                                                                                                                                                  | Upgrade logs.                                                                                          |
|             | hdfs-clean-acls-java.log                                                                                                                                           | Log that records the clearing of deleted roles' ACL information by HDFS.                               |
|             | hdfs-haCheck.log                                                                                                                                                   | Run log that checks whether the NameNode in active or standby state has obtained scripts.              |
|             | <process_name>-jvmpause.log                                                                                                                                        | Log that records JVM pauses during process running.                                                    |
|             | hadoop-<SSH_USER>-balancer-<hostname>.log                                                                                                                          | Run log of HDFS automatic balancing.                                                                   |
|             | hadoop-<SSH_USER>-balancer-<hostname>.out                                                                                                                          | Log that records information of the environment where HDFS executes automatic balancing.               |
|             | hdfs-switch-namenode.log                                                                                                                                           | Run log that records the HDFS active/standby switchover.                                               |
|             | hdfs-router-admin.log                                                                                                                                              | Run log of the mount table management operation                                                        |
|             | threadDump-<DATE>.log                                                                                                                                              | Instance process stack log                                                                             |
| Tomcat logs | hadoop-omm-host1.out, https-catalina.<DATE>.log, https-host-manager.<DATE>.log, https-localhost.<DATE>.log, https-manager.<DATE>.log, localhost_access_web_log.log | Tomcat run log                                                                                         |
| Audit log   | hdfs-audit-<process_name>.log<br>ranger-plugin-audit.log                                                                                                           | Audit log that records the HDFS operations (such as creating, deleting, modifying and querying files). |

| Type | Name               | Description              |
|------|--------------------|--------------------------|
|      | SecurityAuth.audit | HDFS security audit log. |

## Log Level

**Table 9-33** lists the log levels supported by HDFS. The log levels include FATAL, ERROR, WARN, INFO, and DEBUG. Logs of which the levels are higher than or equal to the set level will be printed by programs. The higher the log level is set, the fewer the logs are recorded.

**Table 9-33** Log levels

| Level | Description                                                    |
|-------|----------------------------------------------------------------|
| FATAL | Indicates the critical error information about system running. |
| ERROR | Indicates the error information about system running.          |
| WARN  | Indicates that the current event processing exists exceptions. |
| INFO  | Indicates that the system and events are running properly.     |
| DEBUG | Indicates the system and system debugging information.         |

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of HDFS by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the left menu bar, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.

 **NOTE**

The configurations take effect immediately without restarting the service.

----End

## Log Formats

The following table lists the HDFS log formats.

**Table 9-34** Log formats

| Type      | Format                                                                                                                                       | Example                                                                                                                                                                                                                                                                                                                                                                             |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Run log   | <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> | 2015-01-26 18:43:42,840   INFO   IPC Server handler 40 on 8020   Rolling edit logs   org.apache.hadoop.hdfs.server.namenode.FSEditLog.rollEditLog(FSEditLog.java:1096)                                                                                                                                                                                                              |
| Audit log | <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> | 2015-01-26 18:44:42,607   INFO   IPC Server handler 32 on 8020   allowed=true ugi=hbase (auth:SIMPLE) ip=/10.177.112.145 cmd=getfileinfo src=/hbase/WALs/hghoulaslx410,16020,1421743096083/hghoulaslx410%2C16020%2C1421743096083.1422268722795 dst=null perm=null   org.apache.hadoop.hdfs.server.namenode.FSNameSystem\$DefaultAuditLogger.logAuditMessage(FSNamesystem.java:7950) |

## 9.32 HDFS Performance Tuning

### 9.32.1 Improving Write Performance

#### Scenario

Improve the HDFS write performance by modifying the HDFS attributes.

 **NOTE**

This section applies to MRS 3.x or later.

#### Procedure

Navigation path for setting parameters:

On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** and select **All Configurations**. Enter a parameter name in the search box.

**Table 9-35** Parameters for improving HDFS write performance

| Parameter                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Default Value |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| dfs.datanode.drop.cache.behind.reads | <p>Specifies whether to enable a DataNode to automatically clear all data in the cache after the data in the cache is transferred to the client.</p> <ul style="list-style-type: none"> <li><b>true:</b> The cached data is discarded. This parameter needs to be configured on the DataNode. You are advised to set it to <b>true</b> if data is repeatedly read only a few times, so that the cache can be used by other operations.</li> <li><b>false:</b> You are advised to set it to <b>false</b> if data is read repeatedly for many times to improve the read speed.</li> </ul> <p><b>NOTE</b><br/>This parameter is optional for improving write performance. You can configure it as needed.</p> | false         |
| dfs.client-write-packet-size         | <p>Specifies the size of the client write packet. When the HDFS client writes data to the DataNode, the data will be accumulated until a packet is generated. Then, the packet is transmitted over the network. This parameter specifies the size (unit: byte) of the data packet to be transmitted, which can be specified by each job.</p> <p>In the 10-Gigabit network, you can increase the value of this parameter to enhance the transmission throughput.</p>                                                                                                                                                                                                                                        | 262144        |

## 9.32.2 Improving Read Performance Using Client Metadata Cache

### Scenario

Improve the HDFS read performance by using the client to cache the metadata for block locations.

 **NOTE**

This function is recommended only for reading files that are not modified frequently. Because the data modification done on the server side by some other client is invisible to the cache client, which may cause the metadata obtained from the cache to be outdated.

This section applies to MRS 3.x or later.



## Procedure

### Navigation path for setting parameters:

On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations**, select **All Configurations**, and enter the parameter name in the search box.

**Table 9-36** Parameter configuration

| Parameter                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Default Value |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| dfs.client.metadata.cache.enabled     | Enables or disables the client to cache the metadata for block locations. Set this parameter to <b>true</b> and use it along with the <b>dfs.client.metadata.cache.pattern</b> parameter to enable the cache.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | false         |
| dfs.client.metadata.cache.pattern     | Indicates the regular expression pattern of the path of the file to be cached. Only the metadata for block locations of these files is cached until the metadata expires. This parameter is valid only when <b>dfs.client.metadata.cache.enabled</b> is set to <b>true</b> .<br>Example: <b>/test.*</b> indicates that all files whose paths start with <b>/test</b> are read.<br><b>NOTE</b> <ul style="list-style-type: none"> <li>To ensure consistency, configure a specific mode to cache only files that are not frequently modified by other clients.</li> <li>The regular expression pattern verifies only the path of the URI, but not the schema and authority in the case of the Fully Qualified path.</li> </ul> | -             |
| dfs.client.metadata.cache.expiry.sec  | Indicates the duration for caching metadata. The cache entry becomes invalid after its caching time exceeds this duration. Even metadata that is frequently used during the caching process can become invalid.<br>Time suffixes <b>s/m/h</b> can be used to indicate second, minute, and hour, respectively.<br><b>NOTE</b><br>If this parameter is set to <b>0s</b> , the cache function is disabled.                                                                                                                                                                                                                                                                                                                      | 60s           |
| dfs.client.metadata.cache.max.entries | Indicates the maximum number of non-expired data items that can be cached at a time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 65536         |

 NOTE

Call `DFSClient#clearLocatedBlockCache()` to completely clear the client cache before it expires.

The sample usage is as follows:

```
FileSystem fs = FileSystem.get(conf);
DistributedFileSystem dfs = (DistributedFileSystem) fs;
DFSClient dfsClient = dfs.getClient();
dfsClient.clearLocatedBlockCache();
```

## 9.32.3 Improving the Connection Between the Client and NameNode Using Current Active Cache

### Scenario

When HDFS is deployed in high availability (HA) mode with multiple NameNode instances, the HDFS client needs to connect to each NameNode in sequence to determine which is the active NameNode and use it for client operations.

Once the active NameNode is identified, its details can be cached and shared to all clients running on the client host. In this way, each new client first tries to load the details of the active NameNode from the cache and save the RPC call to the standby NameNode, which can help a lot in abnormal scenarios, for example, when the standby NameNode cannot be connected for a long time.

When a fault occurs and the other NameNode is switched to the active state, the cached details are updated to the information about the current active NameNode.

 NOTE

This section applies to MRS 3.x or later.

### Procedure

Navigation path for setting parameters:

On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations**, select **All Configurations**, and enter the parameter name in the search box.

**Table 9-37** Configuration parameters

| Parameter                                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Default Value                                                           |
|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| dfs.client.failover.proxy.provider<br>[nameservice ID] | Client Failover proxy provider class which creates the NameNode proxy using the authenticated protocol. If this parameter is set to <b>org.apache.hadoop.hdfs.server.namenode.ha.BlackListingFailoverProxyProvider</b> , you can use the NameNode blacklist feature on the HDFS client. If this parameter is set to <b>org.apache.hadoop.hdfs.server.namenode.ha.ObserverReadProxyProvider</b> , you can configure the observer NameNode to process read requests. | org.apache.hadoop.hdfs.server.namenode.ha.AdaptiveFailoverProxyProvider |
| dfs.client.failover.activeinfo.share.flag              | Specifies whether to enable the cache function and share the detailed information about the current active NameNode with other clients. Set it to <b>true</b> to enable the cache function.                                                                                                                                                                                                                                                                        | false                                                                   |
| dfs.client.failover.activeinfo.share.path              | Specifies the local directory for storing the shared files created by all clients in the host. If a cache area is to be shared by different users, the directory must have required permissions (for example, creating, reading, and writing cache files in the specified directory).                                                                                                                                                                              | /tmp                                                                    |
| dfs.client.failover.activeinfo.share.io.timeout.sec    | (Optional) Used to control timeout. The cache file is locked when it is being read or written, and if the file cannot be locked within the specified time, the attempt to read or update the caches will be abandoned. The unit is second.                                                                                                                                                                                                                         | 5                                                                       |

 **NOTE**

The cache files created by the HDFS client are reused by other clients, and thus these files will not be deleted from the local system. If this function is disabled, you may need to manually clear the data.

## 9.33 FAQ

## 9.33.1 NameNode Startup Is Slow

### Question

The NameNode startup is slow when it is restarted immediately after a large number of files (for example, 1 million files) are deleted.

### Answer

It takes time for the DataNode to delete the corresponding blocks after files are deleted. When the NameNode is restarted immediately, it checks the block information reported by all DataNodes. If a deleted block is found, the NameNode generates the corresponding INFO log information, as shown below:

```
2015-06-10 19:25:50,215 | INFO | IPC Server handler 36 on 25000 | BLOCK* processReport: blk_1075861877_2121067 on node 10.91.8.218:9866 size 10249 does not belong to any file | org.apache.hadoop.hdfs.server.blockmanagement.BlockManager.processReport(BlockManager.java:1854)
```

A log is generated for each deleted block. A file may contain one or more blocks. Therefore, after startup, the NameNode spends a large amount of time printing logs when a large number of files are deleted. As a result, the NameNode startup becomes slow.

To address this issue, the following operations can be performed to speed up the startup:

1. After a large number of files are deleted, wait until the DataNode deletes the corresponding blocks and then restart the NameNode.  
You can run the `hdfs dfsadmin -report` command to check the disk space and check whether the files have been deleted.
2. If a large number of the preceding logs are generated, you can change the NameNode log level to **ERROR** so that the NameNode stops printing such logs.

After the NameNode is restarted, change the log level back to **INFO**. You do not need to restart the service after changing the log level.

## 9.33.2 DataNode Is Normal but Cannot Report Data Blocks

### Question

The DataNode is normal, but cannot report data blocks. As a result, the existing data blocks cannot be used.

### Answer

This error may occur when the number of data blocks in a data directory exceeds four times the upper limit (4 x 1 MB). And the DataNode generates the following error logs:

```
2015-11-05 10:26:32,936 | ERROR | DataNode:[[[DISK]file:/srv/BigData/hadoop/data1/dn/]] heartbeating to vm-210/10.91.8.210:8020 | Exception in BPOfferService for Block pool BP-805114975-10.91.8.210-1446519981645 (Datanode Uuid bcada350-0231-413b-bac0-8c65e906c1bb) service to vm-210/10.91.8.210:8020 | BPServiceActor.java:824 java.lang.IllegalStateException:com.google.protobuf.InvalidProtocolBufferException:Protocol message was
```

```

too large.May
be malicious.Use CodedInputStream.setSizeLimit() to increase the size limit. at
org.apache.hadoop.hdfs.protocol.BlockListAsLongs$BufferDecoder$1.next(BlockListAsLongs.java:369)
at org.apache.hadoop.hdfs.protocol.BlockListAsLongs$BufferDecoder$1.next(BlockListAsLongs.java:347) at
org.apache.hadoop.hdfs.
protocol.BlockListAsLongs$BufferDecoder.getBlockListAsLongs(BlockListAsLongs.java:325) at
org.apache.hadoop.hdfs.protocolPB.DatanodeProtocolClientSideTranslatorPB.
blockReport(DatanodeProtocolClientSideTranslatorPB.java:190) at
org.apache.hadoop.hdfs.server.datanode.BPServiceActor.blockReport(BPServiceActor.java:473)
at org.apache.hadoop.hdfs.server.datanode.BPServiceActor.offerService(BPServiceActor.java:685) at
org.apache.hadoop.hdfs.server.datanode.BPServiceActor.run(BPServiceActor.java:822)
at java.lang.Thread.run(Thread.java:745) Caused
by:com.google.protobuf.InvalidProtocolBufferException:Protocol message was too large.May be
malicious.Use CodedInputStream.setSizeLimit()
to increase the size limit. at
com.google.protobuf.InvalidProtocolBufferException.sizeLimitExceeded(InvalidProtocolBufferException.java:1
10) at com.google.protobuf.CodedInputStream.refillBuffer(CodedInputStream.java:755)
at com.google.protobuf.CodedInputStream.readRawByte(CodedInputStream.java:769) at
com.google.protobuf.CodedInputStream.readRawVarint64(CodedInputStream.java:462) at
com.google.protobuf.
CodedInputStream.readInt64(CodedInputStream.java:363) at
org.apache.hadoop.hdfs.protocol.BlockListAsLongs$BufferDecoder$1.next(BlockListAsLongs.java:363)

```

The number of data blocks in the data directory is displayed as **Metric**. You can monitor its value through **http://<datanode-ip>:<http-port>/jmx**. If the value is greater than four times the upper limit (4 x 1 MB), you are advised to configure multiple drives and restart HDFS.

#### Recovery procedure:

1. Configure multiple data directories on the DataNode.

For example, configure multiple directories on the DataNode where only the / **data1/datadir** directory is configured:

```
<property> <name>dfs.datanode.data.dir</name> <value>/data1/datadir</value> </property>
```

Configure as follows:

```
<property> <name>dfs.datanode.data.dir</name> <value>/data1/datadir/,/data2/datadir,/data3/
datadir</value> </property>
```

#### NOTE

You are advised to configure multiple data directories on multiple disks. Otherwise, performance may be affected.

2. Restart the HDFS.
3. Perform the following operation to move the data to the new data directory:
 

```
mv /data1/datadir/current/finalized/subdir1 /data2/datadir/current/finalized/
subdir1
```
4. Restart the HDFS.

### 9.33.3 HDFS WebUI Cannot Properly Update Information About Damaged Data

#### Question

1. When errors occur in the **dfs.datanode.data.dir** directory of DataNode due to the permission or disk damage, HDFS WebUI does not display information about damaged data.
2. After errors are restored, HDFS WebUI does not timely remove related information about damaged data.

## Answer

1. DataNode checks whether the disk is normal only when errors occur in file operations. Therefore, only when a data damage is detected and the error is reported to NameNode, NameNode displays information about the damaged data on HDFS WebUI.
2. After errors are fixed, you need to restart DataNode. During restarting DataNode, all data states are checked and damaged data information is uploaded to NameNode. Therefore, after errors are fixed, damaged data information is not displayed on the HDFS WebUI only by restarting DataNode.

### 9.33.4 Why Does the Distcp Command Fail in the Secure Cluster, Causing an Exception?

#### Question

Why distcp command fails in the secure cluster with the following error displayed?

Client side exception

```
Invalid arguments: Unexpected end of file from server
```

Server side exception

```
javax.net.ssl.SSLException: Unrecognized SSL message, plaintext connection?
```

#### Answer

The preceding error may occur if **webhdfs://** is used in the distcp command. The reason is that the big data cluster uses the HTTPS mechanism, that is, **dfs.http.policy** is set to **HTTPS\_ONLY** in **core-site.xml** file. To avoid the error, replace **webhdfs://** with **swebhdfs://** in the file.

For example:

```
./hadoop distcp swwebhdfs://IP:PORT/testfile hdfs://IP:PORT/testfile1
```

### 9.33.5 Why Does DataNode Fail to Start When the Number of Disks Specified by dfs.datanode.data.dir Equals dfs.datanode.failed.volumes.tolerated?

#### Question

If the number of disks specified by **dfs.datanode.data.dir** is equal to the value of **dfs.datanode.failed.volumes.tolerated**, DataNode startup will fail.

#### Answer

By default, the failure of a single disk will cause the HDFS DataNode process to shut down, which results in the NameNode scheduling additional replicas for each block that is present on the DataNode. This causes needless replications of blocks that reside on disks that have not failed.

To prevent this, you can configure DataNodes to tolerate the failure of **dfs.data.dir** directories; use the **dfs.datanode.failed.volumes.tolerated** parameter in **hdfs-**

**site.xml**. For example, if the value for this parameter is 3, the DataNode will only shut down after four or more data directories have failed. This value is respected on DataNode startup.

When we are configuring tolerate volumes which should be always less than the configured volumes or else we can keep this as -1 which is equal to n-1 (where n is number of disks) then DataNode will not be shut down.

### 9.33.6 Failed to Calculate the Capacity of a DataNode when Multiple data.dir Directories Are Configured in a Disk Partition

#### Question

The capacity of a DataNode fails to calculate when multiple data.dir directories are configured in a disk partition.

#### Answer

Currently, the capacity is calculated based on disks, which is similar to the **df** command in Linux. Ideally, users do not configure multiple data.dir directories in a disk partition. Otherwise, all data will be written to the same disk, greatly deteriorating the performance.

You are advised to configure them as below.

For example, if a node contains the following disks:

```
host-4:~ # df -h
Filesystem Size Used Avail Use% Mounted on
/dev/sda1 352G 11G 324G 4% /
udev 190G 252K 190G 1% /dev
tmpfs 190G 72K 190G 1% /dev/shm
/dev/sdb1 2.7T 74G 2.5T 3% /data1
/dev/sdc1 2.7T 75G 2.5T 3% /data2
/dev/sdd1 2.7T 73G 2.5T 3% /da
```

Recommended configuration:

```
<property>
<name>dfs.datanode.data.dir</name>
<value>/data1/datadir/,/data2/datadir,/data3/datadir</value>
</property>
```

Unrecommended configuration:

```
<property>
<name>dfs.datanode.data.dir</name>
<value>/data1/datadir1/,/data2/datadir1,/data3/datadir1,/data1/datadir2,data1/datadir3,/data2/datadir2,/
data2/datadir3,/data3/datadir2,/data3/datadir3</value>
</property>
```

## 9.33.7 Standby NameNode Fails to Be Restarted When the System Is Powered off During Metadata (Namespace) Storage

### Question

When the standby NameNode is powered off during metadata (namespace) storage, it fails to be started and the following error information is displayed.

```
2015-12-04 11:49:12,121 | ERROR | main | Failed to load image from FS
ImageFile(file=/srv/BigData/namenode/current/fsimage_000000000000096
080,
cpktTxId=0000000000000096080) | FSImage.java:685
java.io.IOException: Invalid MD5 file /srv/BigData/namenode/current/f
simage_0000000000000096080.md5:
the content " " does not match the expecte
d pattern.
at org.apache.hadoop.hdfs.util.MD5FileUtils.readStoredMd5(MD5FileUtil
s.java:92)
at org.apache.hadoop.hdfs.util.MD5FileUtils.readStoredMd5ForFile(MD5F
ileUtils.java:109)
at org.apache.hadoop.hdfs.server.namenode.FSImage.loadFSImage(FSImage
.java:975)
at org.apache.hadoop.hdfs.server.namenode.FSImage.loadFSImageFile(FSI
mage.java:744)
at org.apache.hadoop.hdfs.server.namenode.FSImage.loadFSImage(FSImage
.java:682)
at org.apache.hadoop.hdfs.server.namenode.FSImage.recoverTransitionRe
ad(FSImage.java:300)
at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.loadFSImage(FS
Namesystem.java:968)
at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.loadFromDisk(F
SNamesystem.java:675)
at org.apache.hadoop.hdfs.server.namenode.NameNode.loadNamesystem(Nam
eNode.java:625)
at org.apache.hadoop.hdfs.server.namenode.NameNode.initialize(NameNod
e.java:685)
at org.apache.hadoop.hdfs.server.namenode.NameNode.<init>(NameNode.ja
va:889)
at org.apache.hadoop.hdfs.server.namenode.NameNode.<init>(NameNode.ja
va:872)
at org.apache.hadoop.hdfs.server.namenode.NameNode.createNameNode(Nam
eNode.java:1580)
at org.apache.hadoop.hdfs.server.namenode.NameNode.main(NameNode.java
:1654)
```

### Answer

When the standby NameNode is powered off during metadata (namespace) storage, it fails to be started and the MD5 file is damaged. Remove the damaged fsimage and start the standby NameNode to rectify the fault. After the rectification, the standby NameNode loads the previous fsimage and reproduces all edits.

Recovery procedure:

1. Run the following command to remove the damaged fsimage:  
**rm -rf \${BIGDATA\_DATA\_HOME}/namenode/current/  
fsimage\_000000000000096**
2. Start the standby NameNode.



## 9.33.8 Why Data in the Buffer Is Lost If a Power Outage Occurs During Storage of Small Files

### Question

Why data in the buffer is lost if a power outage occurs during storage of small files?

### Answer

Because of a power outage, the blocks in the buffer are not written to the disk immediately after the write operation is completed. To enable synchronization of blocks to the disk, set **dfs.datanode.synconclose** to **true** in the **hdfs-site.xml** file.

By default, **dfs.datanode.synconclose** is set to **false**. This improves the performance but can cause a buffer data loss in the case of a power outage, and therefore, it is recommended that **dfs.datanode.synconclose** be set to **true** even if this may affect the performance. You can determine whether to enable the synchronization function based on your actual situation.

## 9.33.9 Why Does Array Border-crossing Occur During FileInputFormat Split?

### Question

When HDFS calls the FileInputFormat getSplit method, the ArrayIndexOutOfBoundsException: 0 appears in the following log:

```
java.lang.ArrayIndexOutOfBoundsException: 0
at org.apache.hadoop.mapred.FileInputFormat.identifyHosts(FileInputFormat.java:708)
at org.apache.hadoop.mapred.FileInputFormat.getSplitHostsAndCachedHosts(FileInputFormat.java:675)
at org.apache.hadoop.mapred.FileInputFormat.getSplits(FileInputFormat.java:359)
at org.apache.spark.rdd.HadoopRDD.getPartitions(HadoopRDD.scala:210)
at org.apache.spark.rdd.RDD$$anonfun$partitions$2.apply(RDD.scala:239)
at org.apache.spark.rdd.RDD$$anonfun$partitions$2.apply(RDD.scala:237)
at scala.Option.getOrElse(Option.scala:120)
at org.apache.spark.rdd.RDD.partitions(RDD.scala:237)
at org.apache.spark.rdd.MapPartitionsRDD.getPartitions(MapPartitionsRDD.scala:35)
```

### Answer

The elements of each block correspondent frame are as below: `/default/rack0/;/default/rack0/datanodeip:port`.

The problem is due to a block damage or loss, making the block correspondent machine ip and port become null. Use **hdfs fsck** to check the file blocks health state when this problem occurs, and remove damaged block or restore the missing block to re-computing the task.

## 9.33.10 Why Is the Storage Type of File Copies DISK When the Tiered Storage Policy Is LAZY\_PERSIST?

### Question

When the storage policy of the file is set to **LAZY\_PERSIST**, the storage type of the first replica should be **RAM\_DISK**, and the storage type of other replicas should be **DISK**.

But why is the storage type of all copies shown as **DISK** actually?

### Answer

When a user writes into a file whose storage policy is **LAZY\_PERSIST**, three replicas are written one by one. The first replica is preferentially written into the DataNode where the client is located. The storage type of all replicas is **DISK** in the following scenarios:

- If the DataNode where the client is located does not have the RAM disk, the first replica is written into the disk of the DataNode where the client is located, and other replicas are written into the disks of other nodes.
- If the DataNode where the client is located has the RAM disk, and the value of **dfs.datanode.max.locked.memory** is not specified or smaller than the value of **dfs.blocksize**, the first replica is written into the disk of the DataNode where the client is located, and other replicas are written into the disks of other nodes.

## 9.33.11 The HDFS Client Is Unresponsive When the NameNode Is Overloaded for a Long Time

### Question

When the NameNode node is overloaded (100% of the CPU is occupied), the NameNode is unresponsive. The HDFS clients that are connected to the overloaded NameNode fail to run properly. However, the HDFS clients that are newly connected to the NameNode will be switched to a backup NameNode and run properly.

### Answer

The default configuration must be used (as described in [Table 9-38](#)) when the error preceding described occurs: the **keep alive** mechanism is enabled for the RPC connection between the HDFS client and the NameNode. The **keep alive** mechanism will keep the HDFS client waiting for the response from server and prevent the connection from being out timed, causing the unresponsiveness of the HDFS client.

Perform the following operations to the unresponsive HDFS client:

- Leave the HDFS client waiting. Once the CPU usage of the node where NameNode locates drops, the NameNode will obtain CPU resources and the HDFS client will receive a response.

- If you do not want to leave the HDFS client running, restart the application where the HDFS client locates to reconnect the HDFS client to another idle NameNode.

Procedure:

Configure the following parameters in the **core-site.xml** file on the client.

**Table 9-38** Parameter description

Parameter	Description	Default Value
ipc.client.ping	<p>If the <b>ipc.client.ping</b> parameter is configured to <b>true</b>, the HDFS client will wait for the response from the server and periodically send the <b>ping</b> message to avoid disconnection caused by <b>tcp timeout</b>.</p> <p>If the <b>ipc.client.ping</b> parameter is configured to <b>false</b>, the HDFS client will set the value of <b>ipc.ping.interval</b> as the timeout time. If no response is received within that time, timeout occurs.</p> <p>To avoid the unresponsiveness of HDFS when the NameNode is overloaded for a long time, you are advised to set the parameter to <b>false</b>.</p>	true
ipc.ping.interval	<p>If the value of <b>ipc.client.ping</b> is <b>true</b>, <b>ipc.ping.interval</b> indicates the interval between sending the ping messages.</p> <p>If the value of <b>ipc.client.ping</b> is <b>false</b>, <b>ipc.ping.interval</b> indicates the timeout time for connection.</p> <p>To avoid the unresponsiveness of HDFS when the NameNode is overloaded for a long time, you are advised to set the parameter to a large value, for example <b>900000</b> (unit ms) to avoid timeout when the server is busy.</p>	60000

### 9.33.12 Can I Delete or Modify the Data Storage Directory in DataNode?

#### Question

- In DataNode, the storage directory of data blocks is specified by **dfs.datanode.data.dir**. Can I modify **dfs.datanode.data.dir** to modify the data storage directory?
- Can I modify files under the data storage directory?

#### Answer

During the system installation, you need to configure the **dfs.datanode.data.dir** parameter to specify one or more root directories.

- During the system installation, you need to configure the `dfs.datanode.data.dir` parameter to specify one or more root directories.
- Exercise caution when modifying `dfs.datanode.data.dir`. You can configure this parameter to add a new data root directory.
- Do not modify or delete data blocks in the storage directory. Otherwise, the data blocks will lose.

 **NOTE**

Similarly, do not delete the storage directory, or modify or delete data blocks under the directory using the following parameters:

- `dfs.namenode.edits.dir`
- `dfs.namenode.name.dir`
- `dfs.journalnode.edits.dir`

## 9.33.13 Blocks Miss on the NameNode UI After the Successful Rollback

### Question

Why are some blocks missing on the NameNode UI after the rollback is successful?

### Answer

This problem occurs because blocks with new IDs or genstamps may exist on the DataNode. The block files in the DataNode may have different generation flags and lengths from those in the rollback images of the NameNode. Therefore, the NameNode rejects these blocks in the DataNode and marks the files as damaged.

#### Scenarios:

1. Before an upgrade:  
Client A writes some data to file X. (Assume A bytes are written.)
2. During an upgrade:  
Client A still writes data to file X. (The data in the file is A + B bytes.)
3. After an upgrade:  
Client A completes the file writing. The final data is A + B bytes.
4. Rollback started:  
The status will be rolled back to the status before the upgrade. That is, file X in NameNode will have A bytes, but block files in DataNode will have A + B bytes.

#### Recovery procedure:

1. Obtain the list of damaged files from NameNode web UI or run the following command to obtain:  
**`hdfs fsck <filepath> -list-corruptfileblocks`**
2. Run the following command to delete unnecessary files:

## hdfs fsck <corrupt file path> - delete

### NOTE

Deleting a file is a high-risk operation. Ensure that the files are no longer needed before performing this operation.

3. For the required files, run the **fsck** command to obtain the block list and block sequence.
  - In the block sequence table provided, use the block ID to search for the data directory in the DataNode and download the corresponding block from the DataNode.
  - Write all such block files in appending mode based on the sequence to construct the original file.

Example:

File 1--> blk\_1, blk\_2, blk\_3

Create a file by combining the contents of all three block files from the same sequence.

- Delete the old file from HDFS and rewrite the new file.

## 9.33.14 Why Is "java.net.SocketException: No buffer space available" Reported When Data Is Written to HDFS

### Question

Why is an "java.net.SocketException: No buffer space available" exception reported when data is written to HDFS?

This problem occurs when files are written to the HDFS. Check the error logs of the client and DataNode.

The client logs are as follows:

**Figure 9-13** Client logs

```

2017-07-05 21:58:06.459 INFO [htable-pool3-t1] ipc.AbstractRpcClient: RPC Server Kerberos principal name for service=ClientService is hbase/hadoop.hadoop123.com@HADOOP12
2017-07-05 21:58:06.893 WARN [main] mapreduce.LoadIncrementalHFiles: Skipping non-directory hdfs://hacluster/HBaseTest/bulkload_output/_SUCCESS
2017-07-05 21:59:13.211 WARN [main] hdfs.BlockReaderFactory: I/O error constructing remote block reader.
java.net.SocketException: No buffer space available
 at sun.nio.ch.Net.connect0(Native Method)
 at sun.nio.ch.Net.connect(Net.java:454)
 at sun.nio.ch.SocketChannelImpl.connect(SocketChannelImpl.java:446)
 at sun.nio.ch.SocketChannelImpl.connect(SocketChannelImpl.java:648)
 at org.apache.hadoop.net.SocketIOWithTimeout.connect(SocketIOWithTimeout.java:192)
 at org.apache.hadoop.net.NetUtils.connect(NetUtils.java:531)
 at org.apache.hadoop.hdfs.DFSClient.newConnectedPeer(DFSClient.java:3345)
 at org.apache.hadoop.hdfs.BlockReaderFactory.nextTcpPeer(BlockReaderFactory.java:789)
 at org.apache.hadoop.hdfs.BlockReaderFactory.getRemoteBlockReaderFromTcp(BlockReaderFactory.java:706)
 at org.apache.hadoop.hdfs.BlockReaderFactory.build(BlockReaderFactory.java:369)
 at org.apache.hadoop.hdfs.DFSInputStream.getBlockReader(DFSInputStream.java:713)
 at org.apache.hadoop.hdfs.DFSInputStream.blockSeekTo(DFSInputStream.java:663)
 at org.apache.hadoop.hdfs.DFSInputStream.readWithStrategy(DFSInputStream.java:919)
 at org.apache.hadoop.hdfs.DFSInputStream.read(DFSInputStream.java:973)
 at java.io.DataInputStream.readFully(DataInputStream.java:195)
 at org.apache.hadoop.hbase.io.hfile.FixedFileTrailer.readFromStream(FixedFileTrailer.java:391)
 at org.apache.hadoop.hbase.io.hfile.HFile.isHFileFormat(HFile.java:578)
 at org.apache.hadoop.hbase.io.hfile.HFile.isHFileFormat(HFile.java:560)
 at org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles.visitBulkHFiles(LoadIncrementalHFiles.java:229)
 at org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles.discoverLoadQueue(LoadIncrementalHFiles.java:281)
 at org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles.prepareFileQueue(LoadIncrementalHFiles.java:452)
 at org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles.doBulkLoad(LoadIncrementalHFiles.java:365)
 at org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles.doBulkLoad(LoadIncrementalHFiles.java:331)
 at org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles.run(LoadIncrementalHFiles.java:1197)
 at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:70)
 at org.apache.hadoop.hbase.mapreduce.LoadIncrementalHFiles.main(LoadIncrementalHFiles.java:1114)
2017-07-05 21:59:13.215 WARN [main] hdfs.DFSClient: Failed to connect to /192.168.152.128:25009 for block BP-1969348819-192.168.199.5-1497961637591:blk_1107301222_335745
ffer space available
java.net.SocketException: No buffer space available
 at sun.nio.ch.Net.connect0(Native Method)
 at sun.nio.ch.Net.connect(Net.java:454)
 at sun.nio.ch.SocketChannelImpl.connect(SocketChannelImpl.java:648)
 at org.apache.hadoop.net.SocketIOWithTimeout.connect(SocketIOWithTimeout.java:192)
 at org.apache.hadoop.net.NetUtils.connect(NetUtils.java:531)
 at org.apache.hadoop.hdfs.DFSClient.newConnectedPeer(DFSClient.java:3345)

```

DataNode logs are as follows:

```

2017-07-24 20:43:39,269 | ERROR | DataXceiver for client DFSClient_NONMAPREDUCE_996005058_86
at /192.168.164.155:40214 [Receiving block
BP-1287143557-192.168.199.6-1500707719940:blk_1074269754_528941 with io weight 10] |
DataNode{data=FSDataset{dirpath='[/srv/BigData/hadoop/data1/dn/current, /srv/BigData/hadoop/
data2/dn/current, /srv/BigData/hadoop/data3/dn/current, /srv/BigData/hadoop/data4/dn/current, /srv/
BigData/hadoop/data5/dn/current, /srv/BigData/hadoop/data6/dn/current, /srv/BigData/hadoop/data7/dn/
current]'}, localName='192-168-164-155:9866', datanodeUuid='a013e29c-4e72-400c-bc7b-bbbf0799604c',
xmitsInProgress=0}:Exception transferring block
BP-1287143557-192.168.199.6-1500707719940:blk_1074269754_528941 to mirror 192.168.202.99:9866:
java.net.SocketException: No buffer space available | DataXceiver.java:870
2017-07-24 20:43:39,269 | INFO | DataXceiver for client DFSClient_NONMAPREDUCE_996005058_86
at /192.168.164.155:40214 [Receiving block
BP-1287143557-192.168.199.6-1500707719940:blk_1074269754_528941 with io weight 10] | opWriteBlock
BP-1287143557-192.168.199.6-1500707719940:blk_1074269754_528941 received exception
java.net.SocketException: No buffer space available | DataXceiver.java:933
2017-07-24 20:43:39,270 | ERROR | DataXceiver for client DFSClient_NONMAPREDUCE_996005058_86
at /192.168.164.155:40214 [Receiving block
BP-1287143557-192.168.199.6-1500707719940:blk_1074269754_528941 with io weight 10] |
192-168-164-155:9866:DataXceiver error processing WRITE_BLOCK operation src: /192.168.164.155:40214
dst: /192.168.164.155:9866 | DataXceiver.java:304 java.net.SocketException: No buffer space available
at sun.nio.ch.Net.connect0(Native Method)
at sun.nio.ch.Net.connect(Net.java:454)
at sun.nio.ch.Net.connect(Net.java:446)
at sun.nio.ch.SocketChannelImpl.connect(SocketChannelImpl.java:648)
at org.apache.hadoop.net.SocketIOWithTimeout.connect(SocketIOWithTimeout.java:192)
at org.apache.hadoop.net.NetUtils.connect(NetUtils.java:531)
at org.apache.hadoop.net.NetUtils.connect(NetUtils.java:495)
at org.apache.hadoop.hdfs.server.datanode.DataXceiver.writeBlock(DataXceiver.java:800)
at org.apache.hadoop.hdfs.protocol.datatransfer.Receiver.opWriteBlock(Receiver.java:138)
at org.apache.hadoop.hdfs.protocol.datatransfer.Receiver.processOp(Receiver.java:74)
at org.apache.hadoop.hdfs.server.datanode.DataXceiver.run(DataXceiver.java:265)
at java.lang.Thread.run(Thread.java:748)

```

## Answer

The preceding problem may be caused by network memory exhaustion.

You can increase the threshold of the network device based on the actual scenario.

Example:

```

[root@xxxx ~]# cat /proc/sys/net/ipv4/neigh/default/gc_thresh*
128
512
1024
[root@xxxx ~]# echo 512 > /proc/sys/net/ipv4/neigh/default/gc_thresh1
[root@xxxx ~]# echo 2048 > /proc/sys/net/ipv4/neigh/default/gc_thresh2
[root@xxxx ~]# echo 4096 > /proc/sys/net/ipv4/neigh/default/gc_thresh3
[root@xxxx ~]# cat /proc/sys/net/ipv4/neigh/default/gc_thresh*
512
2048
4096

```

You can also add the following parameters to the `/etc/sysctl.conf` file. The configuration takes effect even if the host is restarted.

```

net.ipv4.neigh.default.gc_thresh1 = 512
net.ipv4.neigh.default.gc_thresh2 = 2048
net.ipv4.neigh.default.gc_thresh3 = 4096

```

## 9.33.15 Why are There Two Standby NameNodes After the active NameNode Is Restarted?

### Question

Why are there two standby NameNodes after the active NameNode is restarted?

When this problem occurs, check the ZooKeeper and ZooKeeper FC logs. You can find that the sessions used for the communication between the ZooKeeper server and client (ZKFC) are inconsistent. The session ID of the ZooKeeper server is **0x164cb2b3e4b36ae4**, and the session ID of the ZooKeeper FC is **0x144cb2b3e4b36ae4**. Such inconsistency means that the data interaction between the ZooKeeper server and ZKFC fails.

Content of the ZooKeeper log is as follows:

```
2015-04-15 21:24:54,257 | INFO | CommitProcessor:22 | Established session 0x164cb2b3e4b36ae4 with negotiated timeout 45000 for client /192.168.0.117:44586 | org.apache.zookeeper.server.ZooKeeperServer.finishSessionInit(ZooKeeperServer.java:623)
2015-04-15 21:24:54,261 | INFO | NIOServerCxn.Factory:192-168-0-114/192.168.0.114:2181 | Successfully authenticated client: authenticationID=hdfs/hadoop@<System domain name>; authorizationID=hdfs/hadoop@<System domain name>. | org.apache.zookeeper.server.auth.SaslServerCallbackHandler.handleAuthorizeCallback(SaslServerCallbackHandler.java:118)
2015-04-15 21:24:54,261 | INFO | NIOServerCxn.Factory:192-168-0-114/192.168.0.114:2181 | Setting authorizedID: hdfs/hadoop@<System domain name> | org.apache.zookeeper.server.auth.SaslServerCallbackHandler.handleAuthorizeCallback(SaslServerCallbackHandler.java:134)
2015-04-15 21:24:54,261 | INFO | NIOServerCxn.Factory:192-168-0-114/192.168.0.114:2181 | adding SASL authorization for authorizationID: hdfs/hadoop@<System domain name> | org.apache.zookeeper.server.ZooKeeperServer.processSasl(ZooKeeperServer.java:1009)
2015-04-15 21:24:54,262 | INFO | ProcessThread(sid:22 cport:-1): | Got user-level KeeperException when processing sessionid:0x164cb2b3e4b36ae4 type:create cxid:0x3 zxid:0x20009fafc txntype:-1 reqpath:n/a Error Path:/hadoop-ha/hacluster/ActiveStandbyElectorLock Error:KeeperErrorCode = NodeExists for /hadoop-ha/hacluster/ActiveStandbyElectorLock | org.apache.zookeeper.server.PrepareRequestProcessor.pRequest(PrepareRequestProcessor.java:648)
```

Content of the ZKFC log is as follows:

```
2015-04-15 21:24:54,237 | INFO | main-SendThread(192-168-0-114:2181) | Socket connection established to 192-168-0-114/192.168.0.114:2181, initiating session | org.apache.zookeeper.ClientCnxn.$SendThread.primeConnection(ClientCnxn.java:854)
2015-04-15 21:24:54,257 | INFO | main-SendThread(192-168-0-114:2181) | Session establishment complete on server 192-168-0-114/192.168.0.114:2181, sessionid = 0x144cb2b3e4b36ae4, negotiated timeout = 45000 | org.apache.zookeeper.ClientCnxn.$SendThread.onConnected(ClientCnxn.java:1259)
2015-04-15 21:24:54,260 | INFO | main-EventThread | EventThread shut down | org.apache.zookeeper.ClientCnxn.$EventThread.run(ClientCnxn.java:512)
2015-04-15 21:24:54,262 | INFO | main-EventThread | Session connected. | org.apache.hadoop.ha.ActiveStandbyElector.processWatchEvent(ActiveStandbyElector.java:547)
2015-04-15 21:24:54,264 | INFO | main-EventThread | Successfully authenticated to ZooKeeper using SASL. | org.apache.hadoop.ha.ActiveStandbyElector.processWatchEvent(ActiveStandbyElector.java:573)
```

## Answer

- Cause Analysis  
After the active NameNode restarts, the temporary node **/hadoop-ha/hacluster/ActiveStandbyElectorLock** created on ZooKeeper is deleted. After the standby NameNode receives that information that the **/hadoop-ha/hacluster/ActiveStandbyElectorLock** node is deleted, the standby NameNode creates the **/hadoop-ha/hacluster/ActiveStandbyElectorLock** node in ZooKeeper in order to switch to the active NameNode. However, when the standby NameNode connects with ZooKeeper through the client ZKFC, the session ID of ZKFC differs from that of ZooKeeper due to network issues, overload CPU, or overload clusters. In this case, the watcher of the standby NameNode fails to detect that the temporary node has been successfully created, and fails to consider the standby NameNode as the active NameNode. After the original active NameNode restarts, it detects that the **/hadoop-ha/hacluster/ActiveStandbyElectorLock** already exists and becomes the standby NameNode. Therefore, both NameNodes are standby NameNodes.

- Solution  
You are advised to restart two ZKFCs of HDFS on FusionInsight Manager.

### 9.33.16 When Does a Balance Process in HDFS, Shut Down and Fail to be Executed Again?

#### Question

After I start a Balance process in HDFS, the process is shut down abnormally. If I attempt to execute the Balance process again, it fails again.

#### Answer

After a Balance process is executed in HDFS, another Balance process can be executed only after the **/system/balancer.id** file is automatically released.

However, if a Balance process is shut down abnormally, the **/system/balancer.id** has not been released when the Balance is executed again, which triggers the **append /system/balancer.id** operation.

- If the time spent on releasing the **/system/balancer.id** file exceeds the soft-limit lease period 60 seconds, executing the Balance process again triggers the append operation, which preempts the lease. The last block is in construction or under recovery status, which triggers the block recovery operation. The **/system/balancer.id** file cannot be closed until the block recovery completes. Therefore, the append operation fails.

After the **append /system/balancer.id** operation fails, the exception message **RecoveryInProgressException** is displayed.

```
org.apache.hadoop.ipc.RemoteException(org.apache.hadoop.protocol.RecoveryInProgressException):
Failed to APPEND_FILE /system/balancer.id for DFSClient because lease recovery is in progress. Try
again later.
```

- If the time spent on releasing the **/system/balancer.id** file is within 60 seconds, the original client continues to own the lease and the exception **AlreadyBeingCreatedException** occurs and null is returned to the client. The following exception message is displayed on the client:  

```
java.io.IOException: Cannot create any NameNode Connectors.. Exiting...
```

Either of the following methods can be used to solve the problem:

- Execute the Balance process again after the hard-limit lease period expires for 1 hour, when the original client has released the lease.
- Delete the **/system/balancer.id** file before executing the Balance process again.

### 9.33.17 "This page can't be displayed" Is Displayed When Internet Explorer Fails to Access the Native HDFS UI

#### Question

Occasionally, internet Explorer 9, Explorer 10, or Explorer 11 fails to access the native HDFS UI.



## Symptom

Internet Explorer 9, Explorer 10, or Explorer 11 fails to access the native HDFS UI, as shown in the following figure.

Turn on TLS 1.0, TLS 1.1, and TLS 1.2 in Advanced settings and try connecting to

## Cause

Some Internet Explorer 9, Explorer 10, or Explorer 11 versions fail to handle SSL handshake issues, causing access failure.

## Solution

Refresh the page.

## 9.33.18 NameNode Fails to Be Restarted Due to EditLog Discontinuity

### Question

If a JournalNode server is powered off, the data directory disk is fully occupied, and the network is abnormal, the EditLog sequence number on the JournalNode is inconsecutive. In this case, the NameNode restart may fail.

### Symptom

The NameNode fails to be restarted. The following error information is reported in the NameNode run logs:

```
2019-11-08 16:30:28,399 | ERROR | main | Failed to start namenode. | NameNode.java:1732
java.io.IOException: There appears to be a gap in the edit log. We expected txid 13698019, but got txid 13698088.
 at org.apache.hadoop.hdfs.server.namenode.MetaRecoveryContext.editLogLoaderPrompt(MetaRecoveryContext.java:94)
 at org.apache.hadoop.hdfs.server.namenode.FSEditLogLoader.loadEditRecords(FSEditLogLoader.java:278)
 at org.apache.hadoop.hdfs.server.namenode.FSEditLogLoader.loadFSEdits(FSEditLogLoader.java:188)
 at org.apache.hadoop.hdfs.server.namenode.FSImage.loadEdits(FSImage.java:924)
 at org.apache.hadoop.hdfs.server.namenode.FSImage.loadFSImage(FSImage.java:771)
 at org.apache.hadoop.hdfs.server.namenode.FSImage.recoverTransitionRead(FSImage.java:331)
 at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.loadFSImage(FSNamesystem.java:1108)
 at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.loadFromDisk(FSNamesystem.java:727)
 at org.apache.hadoop.hdfs.server.namenode.NameNode.loadNamesystem(NameNode.java:638)
 at org.apache.hadoop.hdfs.server.namenode.NameNode.initialize(NameNode.java:700)
 at org.apache.hadoop.hdfs.server.namenode.NameNode.<init>(NameNode.java:943)
 at org.apache.hadoop.hdfs.server.namenode.NameNode.<init>(NameNode.java:916)
 at org.apache.hadoop.hdfs.server.namenode.NameNode.createNameNode(NameNode.java:1655)
 at org.apache.hadoop.hdfs.server.namenode.NameNode.main(NameNode.java:1725)
```

## Solution

1. Find the active NameNode before the restart, go to its data directory (you can obtain the directory, such as **/srv/BigData/namenode/current** by checking the configuration item **dfs.namenode.name.dir**), and obtain the sequence number of the latest FsImage file, as shown in the following figure:

```

-rw-----, 1 omm wheel 574 Oct 2 01:12 edits_0000000000013259401-0000000000:
-rw-----, 1 omm wheel 575 Oct 2 01:13 edits_0000000000013259409-0000000000:
-rw-----, 1 omm wheel 42 Oct 2 01:13 edits_0000000000013259417-0000000000:
-rw-----, 1 omm wheel 1048576 Nov 8 16:01 edits_inprogress_0000000000013698088
-rw-----, 1 omm wheel 314803 Nov 8 15:53 fsimage_0000000000013698018
-rw-----, 1 omm wheel 62 Nov 8 15:53 fsimage_0000000000013698018.md5
-rw-----, 1 omm wheel 314803 Nov 8 15:56 fsimage_0000000000013698050
-rw-----, 1 omm wheel 62 Nov 8 15:56 fsimage_0000000000013698050.md5
-rw-----, 1 omm wheel 314803 Nov 8 15:59 fsimage_0000000000013698066
-rw-----, 1 omm wheel 62 Nov 8 15:59 fsimage_0000000000013698066.md5
-rw-----, 1 omm wheel 9 Oct 2 01:13 seen_txid
-rw-----, 1 omm wheel 187 Nov 8 15:59 VERSION

```

2. Check the data directory of each JournalNode (you can obtain the directory such as `/srv/BigData/journalnode/hacluster/current` by checking the value of the configuration item `dfs.journalnode.edits.dir`), and check whether the sequence number starting from that obtained in step 1 is consecutive in edits files. That is, you need to check whether the last sequence number of the previous edits file is consecutive with the first sequence number of the next edits file. (As shown in the following figure, `edits_0000000000013259231-0000000000013259237` and `edits_0000000000013259239-0000000000013259246` are not consecutive.)

```

-rw-----, 1 omm wheel 576 Oct 2 00:41 edits_0000000000013259151-0000000000013259158
-rw-----, 1 omm wheel 575 Oct 2 00:43 edits_0000000000013259159-0000000000013259166
-rw-----, 1 omm wheel 576 Oct 2 00:43 edits_0000000000013259167-0000000000013259174
-rw-----, 1 omm wheel 575 Oct 2 00:45 edits_0000000000013259175-0000000000013259182
-rw-----, 1 omm wheel 575 Oct 2 00:45 edits_0000000000013259183-0000000000013259190
-rw-----, 1 omm wheel 576 Oct 2 00:47 edits_0000000000013259191-0000000000013259198
-rw-----, 1 omm wheel 575 Oct 2 00:48 edits_0000000000013259199-0000000000013259206
-rw-----, 1 omm wheel 575 Oct 2 00:49 edits_0000000000013259207-0000000000013259214
-rw-----, 1 omm wheel 575 Oct 2 00:50 edits_0000000000013259215-0000000000013259222
-rw-----, 1 omm wheel 573 Oct 2 00:51 edits_0000000000013259223-0000000000013259230
-rw-----, 1 omm wheel 571 Oct 2 00:52 edits_0000000000013259231-0000000000013259237
-rw-----, 1 omm wheel 576 Oct 2 00:53 edits_0000000000013259239-0000000000013259246
-rw-----, 1 omm wheel 575 Oct 2 00:54 edits_0000000000013259247-0000000000013259254
-rw-----, 1 omm wheel 576 Oct 2 00:55 edits_0000000000013259255-0000000000013259262
-rw-----, 1 omm wheel 42 Oct 2 00:56 edits_0000000000013259263-0000000000013259264
-rw-----, 1 omm wheel 1107 Oct 2 00:57 edits_0000000000013259265-0000000000013259278
-rw-----, 1 omm wheel 42 Oct 2 00:58 edits_0000000000013259279-0000000000013259280
-rw-----, 1 omm wheel 1109 Oct 2 00:59 edits_0000000000013259281-0000000000013259294
-rw-----, 1 omm wheel 42 Oct 2 01:00 edits_0000000000013259295-0000000000013259296
-rw-----, 1 omm wheel 1299 Oct 2 01:01 edits_0000000000013259297-0000000000013259312
-rw-----, 1 omm wheel 260 Oct 2 01:02 edits_0000000000013259313-0000000000013259316
-rw-----, 1 omm wheel 984 Oct 2 01:03 edits_0000000000013259317-0000000000013259328
-rw-----, 1 omm wheel 572 Oct 2 01:04 edits_0000000000013259329-0000000000013259336
-rw-----, 1 omm wheel 575 Oct 2 01:05 edits_0000000000013259337-0000000000013259344
-rw-----, 1 omm wheel 983 Oct 2 01:06 edits_0000000000013259345-0000000000013259356

```

3. If the edits files are not consecutive, check whether the edits files with the related sequence number exist in the data directories of other JournalNodes or NameNode. If the edits files can be found, copy a consecutive segment to the JournalNode.
4. In this way, all inconsecutive edits files are restored.
5. Restart the NameNode and check whether the restart is successful. If the fault persists, contact technical support.

# 10 Using Hive

## 10.1 Using Hive from Scratch

Hive is a data warehouse framework built on Hadoop. It maps structured data files to a database table and provides SQL-like functions to analyze and process data. It also allows you to quickly perform simple MapReduce statistics using SQL-like statements without the need of developing a specific MapReduce application. It is suitable for statistical analysis of data warehouses.

### Background

Suppose a user develops an application to manage users who use service A in an enterprise. The procedure of operating service A on the Hive client is as follows:

#### Operations on common tables:

- Create the **user\_info** table.
- Add users' educational backgrounds and professional titles to the table.
- Query user names and addresses by user ID.
- Delete the user information table after service A ends.

**Table 10-1** User information

ID	Name	Gender	Age	Address
12005000201	A	Male	19	City A
12005000202	B	Female	23	City B
12005000203	C	Male	26	City C
12005000204	D	Male	18	City D
12005000205	E	Female	21	City E
12005000206	F	Male	32	City F
12005000207	G	Female	29	City G

ID	Name	Gender	Age	Address
12005000208	H	Female	30	City H
12005000209	I	Male	26	City I
12005000210	J	Female	25	City J

## Procedure

**Step 1** Download the client configuration file.

**For versions earlier than MRS 3.x, perform the following operations:**

1. Log in to MRS Manager. For details, see [Accessing Manager](#). Then, choose **Services**.
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.

**For MRS 3.x or later, perform the following operations:**

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Dashboard** > **More** > **Download Client**.
3. Download the cluster client.  
Set **Select Client Type** to **Configuration Files Only**, select a platform type, and click **OK** to generate the client configuration file which is then saved in the **/tmp/FusionInsight-Client/** directory on the active management node by default.

**Step 2** Log in to the active management node of Manager.

**For versions earlier than MRS 3.x, perform the following operations:**

1. On the MRS console, click **Clusters**, choose **Active Clusters**, and click a cluster name. On the **Nodes** tab, view the node names. The node whose name contains **master1** is the Master1 node, and the node whose name contains **master2** is the Master2 node.  
The active and standby management nodes of MRS Manager are installed on Master nodes by default. Because Master1 and Master2 are switched over in active and standby mode, Master1 is not always the active management node of MRS Manager. Run a command in Master1 to check whether Master1 is active management node of MRS Manager. For details about the command, see [Step 2.4](#).
2. Log in to the Master1 node using the password as user **root**.
3. Run the following commands to switch to user **omm**:  

```
sudo su - root
su - omm
```

- Run the following command to check the active management node of MRS Manager:

```
sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh
```

In the command output, the node whose **HAActive** is **active** is the active management node, and the node whose **HAActive** is **standby** is the standby management node. In the following example, **mgtomsdat-sh-3-01-1** is the active management node, and **mgtomsdat-sh-3-01-2** is the standby management node.

```
Ha mode
double
NodeName HostName HAVersion StartTime HAActive
HAAllResOK HARunPhase
192-168-0-30 mgtomsdat-sh-3-01-1 V100R001C01 2014-11-18 23:43:02
active normal Activated
192-168-0-24 mgtomsdat-sh-3-01-2 V100R001C01 2014-11-21 07:14:02
standby normal Deactivated
```

- Log in to the active management node as user **root**, for example, node **192-168-0-30**.

**For MRS 3.x or later, perform the following operations:**

- Log in to any node where Manager is deployed as user **root**.
- Run the following command to identify the active and standby nodes:

```
sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh
```

In the command output, the value of **HAActive** for the active management node is **active**, and that for the standby management node is **standby**. In the following example, **node-master1** is the active management node, and **node-master2** is the standby management node.

```
HAMode
double
NodeName HostName HAVersion StartTime HAActive
HAAllResOK HARunPhase
192-168-0-30 node-master1 V100R001C01 2020-05-01 23:43:02 active
normal Activated
192-168-0-24 node-master2 V100R001C01 2020-05-01 07:14:02 standby
normal Deactivated
```

- Log in to the primary management node as user **root** and run the following command to switch to user **omm**:

```
sudo su - omm
```

**Step 3** Run the following command to go to the client installation directory:

```
cd /opt/client
```

The cluster client has been installed in advance. The following client installation directory is used as an example. Change it based on the site requirements.

**Step 4** Run the following command to update the client configuration for the active management node.

```
sh refreshConfig.sh /opt/client Full path of the client configuration file package
```

For example, run the following command:

```
sh refreshConfig.sh /opt/client /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_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.
```

**Step 5** Use the client on a Master node.

1. On the active management node, for example, **192-168-0-30**, run the following command to switch to the client directory, for example, **/opt/client**.  
**cd /opt/client**
2. Run the following command to configure environment variables:  
**source bigdata\_env**
3. If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user:  
**kinit MRS cluster user**  
Example: user **kinit hiveuser**  
The current user must have the permission to create Hive tables. If Kerberos authentication is disabled, skip this step.
4. Run the client command of the Hive component directly.  
**beeline**


**Step 6** Run the Hive client command to implement service A.

**Operations on internal tables:**

1. Create the **user\_info** user information table according to [Table 10-1](#) and add data to it.  
**create table user\_info(id string,name string,gender string,age int,addr string);**  
For MRS 1.x, MRS 3.x, or later, perform the following operations:  
**insert into table user\_info(id,name,gender,age,addr)  
values("12005000201","A","Male",19,"City A");**  
For MRS 2.x, perform the following operations:  
**insert into table user\_info values("12005000201","A","Male",19,"City A");**
2. Add users' educational backgrounds and professional titles to the **user\_info** table.  
For example, to add educational background and title information about user 12005000201, run the following command:  
**alter table user\_info add columns(education string,technical string);**
3. Query user names and addresses by user ID.  
For example, to query the name and address of user 12005000201, run the following command:  
**select name,addr from user\_info where id='12005000201';**
4. Delete the user information table.  
**drop table user\_info;**

**Operations on external partition tables:**

Create an external partition table and import data.

1. Create a path for storing external table data.  
**hdfs dfs -mkdir /hive/**  
**hdfs dfs -mkdir /hive/user\_info**
2. Create a table.  
**create external table user\_info(id string,name string,gender string,age int,addr string) partitioned by(year string) row format delimited fields terminated by ' ' lines terminated by '\n' stored as textfile location '/hive/user\_info';**  
  
 **NOTE**  
  
**fields terminated** indicates delimiters, for example, spaces.  
**lines terminated** indicates line breaks, for example, \n.  
**/hive/user\_info** indicates the path of the data file.
3. Import data.
  - a. Execute the insert statement to insert data.  
**insert into user\_info partition(year="2018") values ("12005000201","A","Male",19,"City A");**
  - b. Run the **load data** command to import file data.
    - i. Create a file based on the data in [Table 10-1](#). For example, the file name is **txt.log**. Fields are separated by space, and the line feed characters are used as the line breaks.
    - ii. Upload the file to HDFS.  
**hdfs dfs -put txt.log /tmp**
    - iii. Load data to the table.  
**load data inpath '/tmp/txt.log' into table user\_info partition (year='2011');**
4. Query the imported data.  
**select \* from user\_info;**
5. Delete the user information table.  
**drop table user\_info;**
6. Run the following command to exit:  
**!q**  
  
----End

## 10.2 Configuring Hive Parameters

### Navigation Path

Go to the Hive configurations page by referring to [Modifying Cluster Service Configuration Parameters](#).

## Parameter Description

**Table 10-2** Hive parameter description

Parameter	Description	Default Value
hive.auto.convert.join	Whether Hive converts common <b>join</b> to <b>mapjoin</b> based on the input file size. <b>NOTE</b> When Hive is used to query a join table, whatever the table size is (if the data in the join table is less than 24 MB, it is a small one), set this parameter to <b>false</b> . If this parameter is set to <b>true</b> , new <b>mapjoin</b> cannot be generated when you query a join table.	Possible values are as follows: <ul style="list-style-type: none"> <li>• true</li> <li>• false</li> </ul> The default value is <b>true</b> .
hive.default.fileformat	Indicates the default file format used by Hive.	Versions earlier than MRS 3.x: TextFile MRS 3.x or later: RCFile
hive.exec.reducers.max	Indicates the maximum number of reducers in a MapReduce job submitted by Hive.	999
hive.server2.thrift.max.worker.threads	Indicates the maximum number of threads that can be started in the HiveServer internal thread pool.	1,000
hive.server2.thrift.min.worker.threads	Indicates the number of threads started during initialization in the HiveServer internal thread pool.	5
hive.hbase.delete.mode.enabled	Indicates whether to enable the function of deleting HBase records from Hive. If this function is enabled, you can use <b>remove table xx where xxx</b> to delete HBase records from Hive. <b>NOTE</b> This parameter applies to MRS 3.x or later.	true



Parameter	Description	Default Value
hive.metastore.server.min.threads	Indicates the number of threads started by MetaStore for processing connections. If the number of threads is more than the set value, MetaStore always maintains a number of threads that is not lower than the set value, that is, the number of resident threads in the MetaStore thread pool is always higher than the set value.	200
hive.server2.enable.doAs	Indicates whether to simulate client users during sessions between HiveServer2 and other services (such as Yarn and HDFS). If you change the configuration item from <b>false</b> to <b>true</b> , users with only the column permission lose the permissions to access corresponding tables. <b>NOTE</b> This parameter applies to MRS 3.x or later.	true

## 10.3 Hive SQL

Hive SQL supports all features of Hive-3.1.0..

[Table 10-3](#) describes the extended Hive statements provided by .

**Table 10-3** Extended Hive statements

Extended Syntax	Syntax Description	Syntax Example	Example Description
<pre>CREATE [TEMPORARY] [EXTERNAL] TABLE [IF NOT EXISTS] [db_name.]table_ name (col_name data_type [COMMENT col_comment], ...) [ROW FORMAT row_format] [STORED AS file_format]   STORED BY 'storage.handler.cl ass.name' [WITH SERDEPROPERTIE S (...) ] ..... [TBLPROPERTIES ("groupId"=" group1 ","locatorId"="loc ator1")] ...;</pre>	<p>The statement is used to create a Hive table and specify locators on which table data files locate. For details, see <a href="#">Using HDFS Colocation to Store Hive Tables</a>.</p>	<pre>CREATE TABLE tab1 (id INT, name STRING) row format delimited fields terminated by '\t' stored as RCFILE TBLPROPERTIES(" groupId"=" group1 ","locatorId"="loc ator1");</pre>	<p>The statement is used to create table <b>tab1</b> and specify locator1 on which the table data of <b>tab1</b> locates.</p>

Extended Syntax	Syntax Description	Syntax Example	Example Description
<pre>CREATE [TEMPORARY] [EXTERNAL] TABLE [IF NOT EXISTS] [db_name.]table_name (col_name data_type [COMMENT col_comment], ...) [ROW FORMAT row_format] [STORED AS file_format]   STORED BY 'storage.handler.class.name' [WITH SERDEPROPERTIES (...)] ... [TBLPROPERTIES ('column.encode.columns'='col_name1,col_name2'  'column.encode.indices'='col_id1,col_id2', 'column.encode.classname'='encode_classname')]...;</pre>	<p>The statement is used to create a hive table and specify the table encryption column and encryption algorithm. For details, see <a href="#">Using the Hive Column Encryption Function</a>.</p>	<pre>create table encode_test(id INT, name STRING, phone STRING, address STRING) ROW FORMAT SERDE 'org.apache.hadoop p.hive.serde2.lazy. LazySimpleSerDe' WITH SERDEPROPERTIES S ('column.encode.i ndices'='2,3', 'column.encode.cl assname'='org.apa che.hadoop.hive.s erde2.SMS4Rewrit er') STORED AS TEXTFILE;</pre>	<p>The statement is used to create table <b>encode_test</b> and specify that column 2 and column 3 will be encrypted using the <b>org.apache.hadoop.hive.serde2.SMS4Rewriter</b> encryption algorithm class during data insertion.</p>
<pre>REMOVE TABLE hbase_tablename [WHERE where_condition];</pre>	<p>The statement is used to delete data that meets criteria from the Hive on HBase table. For details, see <a href="#">Deleting Single-Row Records from Hive on HBase</a>.</p>	<pre>remove table hbase_table1 where id = 1;</pre>	<p>The statement is used to delete data that meets the criterion of "id = 1" from the table.</p>

Extended Syntax	Syntax Description	Syntax Example	Example Description
<pre>CREATE [TEMPORARY] [EXTERNAL] TABLE [IF NOT EXISTS] [db_name.]table_ name (col_name data_type [COMMENT col_comment], ...) [ROW FORMAT row_format] <b>STORED AS inputformat 'org.apache.hado op.hive.contrib.fil eformat.Specifie dDelimiterInput- Format'</b> outputformat 'org.apache.hadoo p.hive ql.io.HiveIg noreKeyTextOutpu tFormat';</pre>	<p>The statement is used to create a hive table and specify that the table supports customized row delimiters. For details, see <a href="#">Customizing Row Separators</a>.</p>	<pre>create table blu(time string, num string, msg string) row format delimited fields terminated by ',' <b>stored as inputformat 'org.apache.hado op.hive.contrib.fil eformat.Specifie dDelimiterInput- Format'</b> outputformat 'org.apache.hadoo p.hive ql.io.HiveIg noreKeyTextOutpu tFormat';</pre>	<p>The statement is used to create table <b>blu</b> and set <b>inputformat</b> to <b>SpecifiedDelimiterInputFormat</b> so that the query row delimiter can be specified during the query.</p>

## 10.4 Permission Management

### 10.4.1 Hive Permission

Hive is a data warehouse framework built on Hadoop. It provides basic data analysis services using the Hive query language (HQL), a language like the structured query language (SQL).

MRS supports users, user groups, and roles. Permissions must be assigned to roles and then roles are bound to users or user groups. Users can obtain permissions only by binding a role or joining a group that is bound with a role.

#### NOTE

- Hive permissions in security mode need to be managed whereas those in normal mode do not.
- MRS 3.x or later supports Ranger. If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for Hive](#).

### Hive Permission Model

To use the Hive component, users must have permissions on Hive databases and tables (including external tables and views). In MRS, the complete Hive permission

model is composed of Hive metadata permission and HDFS file permission. The Hive permission model also includes the permission to use databases or tables.

- Hive metadata permission

Similar to traditional relational databases, the Hive database of MRS supports the **CREATE** and **SELECT** permission, and the Hive tables and columns support the **SELECT**, **INSERT**, and **DELETE** permissions. Hive also supports the permissions of **OWNERSHIP** and **Hive Admin Privilege**.

 **NOTE**

The **UPDATE** and **DELETE** operations on Hive tables and columns can be performed only when **ACID** is enabled.

- Hive data file permission, also known as HDFS file permission

Hive database and table files are stored in the HDFS. The created databases or tables are saved in the **/user/hive/warehouse** directory of the HDFS by default. The system automatically creates subdirectories named after database names and database table names. To access a database or a table, the corresponding file permissions (read, write, and execute) on the HDFS are required.

To perform various operations on Hive databases or tables, you need to associate the metadata permission with the HDFS file permission. For example, to query Hive data tables, you need to associate the metadata permission **SELECT** and the HDFS file permissions **Read** and **Write**.

To use the role management function of Manager GUI to manage the permissions of Hive databases and tables, you only need to configure the metadata permission, and the system will automatically associate and configure the HDFS file permission. In this way, operations on the interface are simplified, and the efficiency is improved.

## Hive Users

MRS provides users and roles to use Hive, such as creating tables, inserting data into tables, and querying tables. Hive defines the **USER** class, corresponding to user instances. Hive defines the **GROUP** class, corresponding to role instances.

You can use Manager to set permissions for Hive users. This method only supports permission setting in roles. A user or user group can obtain the permissions only after a role is bound to the user or user group. Hive users can be granted Hive administrator permissions and permissions to access databases, tables, and columns.

## Hive Usage Scenarios and Related Permissions

Creating a database with Hive requires users to join in the **hive** group, without granting a role. Users have all permissions on the databases or tables created by themselves in Hive or HDFS. They can create tables, select, delete, insert, or update data, and grant permissions to other users to allow them to access the tables and corresponding HDFS directories and files.

A user can access the tables or database only with permissions. The permission required by users varies according to Hive usage scenarios.

**Table 10-4** Hive usage scenarios

Typical Scenario	Permission
Using Hive tables, columns, or databases	<p>Permissions required in different scenarios are as follows:</p> <ul style="list-style-type: none"> <li>• To create tables, the <b>CREATE</b> permission is required.</li> <li>• To query data, the <b>SELECT</b> permission is required.</li> <li>• To insert data, the <b>INSERT</b> permission is required.</li> <li>• To delete data, the <b>DELETE</b> permission is required.</li> </ul>
Associating and using other components	<p>In addition to Hive permissions, permissions of other components are required in some scenarios, for example:</p> <ul style="list-style-type: none"> <li>• Yarn permissions are required when some HQL statements, such as <b>insert, count, distinct, group by, order by, sort by,</b> and <b>join,</b> are run. You are advised to grant Yarn permissions to the role of each Hive user.</li> <li>• HBase permission is required when Hive over HBase is used, for example, querying HBase table data in Hive.</li> </ul>

In some special Hive usage scenarios, you need to configure other types of permission.

**Table 10-5** Hive authorization precautions

Scenario	Permission
Creating Hive databases, tables, and external tables, or adding partitions to created Hive tables or external tables when data files specified by Hive users are saved to other HDFS directories except <b>/user/hive/warehouse</b>	<p>The directory must already exist, the Hive user must be the owner of the directory, and the Hive user must have the read, write, and execute permissions on the directory. The user must have the <b>read</b> and <b>write</b> permissions of all the upper-layer directories of the directory. After an administrator grants the Hive permission to the role, the HDFS permission is automatically granted.</p>

Scenario	Permission
<p>Using <b>load</b> to load data from all the files or specified files in a specified directory to Hive tables as a Hive user</p>	<ul style="list-style-type: none"> <li>• The data source is a Linux local disk, the specified directory exists, and the system user <b>omm</b> has read and execute permission of the directory and all its upper-layer directories. The specified file exists, and user <b>omm</b> has read permission of the file and has the read and execute permission of all the upper-layer directories of the file.</li> <li>• The data source is HDFS, the specified directory exists, and the Hive user is the owner of the directory and has read, write, and execute permission on the directory and its subdirectories, and has read and write permission on all its upper-layer directories. The specified file exists, and the Hive user is the owner of the file and has read, write, and execute permission, and has read and execute permission on the file and all its upper-layer directories.</li> </ul> <p><b>NOTE</b> When <b>load</b> is used to import data to a Linux local disk, files must be loaded to the HiveServer on which the command is run and the permission must be modified. You are advised to run the command on a client. The HiveServer to which the client is connected can be found. For example, if the Hive client displays <b>0:</b> <b>jdbc:hive2://10.172.0.43:21066/</b>&gt;, the IP address of the connected HiveServer is 10.172.0.43.</p>
<p>Creating or deleting functions or modifying any database</p>	<p>The <b>Hive Admin Privilege</b> is required.</p>
<p>Performing operations on all databases and tables in Hive</p>	<p>The user must be added to the <b>supergroup</b> user group and granted <b>Hive Admin Privilege</b>.</p>

## 10.4.2 Creating a Hive Role

### Scenario

Create and configure a Hive role on Manager as an MRS cluster administrator. The Hive role can be granted the permissions of the Hive administrator and the permissions to operate Hive table data.

Creating a database with Hive requires users to join in the **hive** group, without granting a role. Users have all permissions on the databases or tables created by themselves in Hive or HDFS. They can create tables, select, delete, insert, or update data, and grant permissions to other users to allow them to access the tables and corresponding HDFS directories and files. The created databases or tables are saved in the **/user/hive/warehouse** directory of the HDFS by default.

 **NOTE**

- A Hive role can be created only in security mode.
- If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for Hive](#) for MRS 3.x or later that supports Ranger.

## Prerequisites

- The MRS cluster administrator has understood service requirements.
- Log in to FusionInsight Manager.
- The Hive client has been installed.

## Procedure

For versions earlier than MRS 3.x, perform the following operations to create a Hive role:

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

**Step 2** Choose **System > Permission > Manage Role**.

**Step 3** Click **Create Role**, and set **Role Name** and **Description**.

**Step 4** Set permissions. For details, see [Table 10-6](#).

- **Hive Admin Privilege:** Hive administrator permissions. If you want to use this permission, run the **set role admin** command to set the permission before running SQL statements.
- **Hive Read Write Privileges:** Hive data table management permission, which is the operation permission to set and manage the data of created tables. Select the permissions of a database as required. To specify permissions on tables, click the database name and select the permissions of the tables.

 **NOTE**

- Hive role management supports Hive administrator permissions and the permissions to access tables and views, but does not support granting permissions on databases.
- The permissions of the Hive administrator do not include the permission to manage HDFS.
- If there are too many tables in the database or too many files in tables, the permission granting may last a while. For example, if a table contains 10,000 files, the permission granting lasts about 2 minutes.



**Table 10-6** Setting a role

Scenario	Role Authorization
Setting the Hive administrator permission	In the <b>Permission</b> table, click <b>Hive</b> and select <b>Hive Admin Privilege</b> . <b>NOTE</b> After being bound to the Hive administrator role, perform the following operations during each maintenance operation: <ol style="list-style-type: none"> <li>1. Log in to the node where the client is installed.</li> <li>2. Run the following command to configure environment variables: For example, if the Hive client installation directory is <b>/opt/hiveclient</b>, run <b>source /opt/hiveclient/bigdata_env</b>.</li> <li>3. Run the following command to authenticate the user: <b>kinit Hive service user</b></li> <li>4. Run the following command to log in to the client tool: <b>beeline</b></li> <li>5. Run the following command to update the administrator permissions of the Hive user: <b>set role admin;</b></li> </ol>
Setting the permission to query a table of another user in the default database	<ol style="list-style-type: none"> <li>1. In the <b>Permission</b> table, choose <b>Hive &gt; Hive Read Write Privileges</b>.</li> <li>2. In the <b>Permission</b> column of the specified table, select <b>SELECT</b>.</li> </ol>
Setting the permission to query a table of another user in the default database	<ol style="list-style-type: none"> <li>1. In the <b>Permission</b> table, choose <b>Hive &gt; Hive Read Write Privileges</b>.</li> <li>2. In the <b>Permission</b> column of the specified table, select <b>Insert</b>.</li> </ol>
Setting the permission to import data to a table of another user in the default database	<ol style="list-style-type: none"> <li>1. In the <b>Permission</b> table, choose <b>Hive &gt; Hive Read Write Privileges</b>.</li> <li>2. In the <b>Permission</b> column of the specified table, select <b>Delete</b> and <b>Insert</b>.</li> </ol>

Scenario	Role Authorization
Setting the permission to submit HQL commands to Yarn for execution	<p>The HQL commands used by some services are converted into MapReduce tasks and submitted to Yarn for execution. You need to set the Yarn permissions. For example, the HQL statements to be run use statements, such as <b>insert</b>, <b>count</b>, <b>distinct</b>, <b>group by</b>, <b>order by</b>, <b>sort by</b>, or <b>join</b>.</p> <ol style="list-style-type: none"> <li>1. In the <b>Permission</b> table, choose <b>Yarn &gt; Scheduler Queue &gt; root</b>.</li> <li>2. In the <b>Permission</b> column of the default queue, select <b>Submit</b>.</li> </ol>

**Step 5** Click **OK**, and return to the **Role** page.

**Step 6** Choose **System > Manage User > Create User**.

**Step 7** Enter the username, set **User Type** to **Human-machine**, set the user password, add a user group bound with the Hive administrator role, bind the new Hive role to the user group, and click **OK**.

**Step 8** After the user is created, you can run the SQL statement using the user.

----End

For MRS 3.x or later, perform the following operations to create a Hive role:

**Step 1** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#)

**Step 2** Choose **System > Permission > Role**.

**Step 3** Click **Create Role**, and set **Role Name** and **Description**.

**Step 4** Set **Configure Resource Permission**. For details, see [Table 10-7](#).

- Grant the read and execution permissions for the HDFS directory.
  - Click *Name of the desired cluster* and select **HDFS** for **Service Name**. On the displayed page, click **File System**, choose **hdfs://hacluster/ > user**, locate the row where **hive** is located, and select **Read** and **Execute** in the **Permission** column.
  - Click *Name of the desired cluster* and select **HDFS** for **Service Name**. On the displayed page, click **File System**, choose **hdfs://hacluster/ > user > hive**, locate the row where **warehouse** is located, and select **Read** and **Execute** in the **Permission** column.
  - Click *Name of the desired cluster* and select **HDFS** for **Service Name**. On the displayed page, click **File System**, choose **hdfs://hacluster/ > tmp**, locate the row where **hive-scratch** is located, and select **Read** and **Execute** in the **Permission** column.
- **Hive Admin Privilege:** Hive administrator permission.
- **Hive Read Write Privileges:** Hive data table management permission, which is the operation permission to set and manage the data of created tables.

 **NOTE**

- In MRS 3.1.0, Hive role management supports the Hive administrator permission, and the permissions of accessing tables and views, without granting the database permission.
- The permissions of the Hive administrator do not include the permission to manage HDFS.
- If there are too many tables in the database or too many files in tables, the permission granting may last a while. For example, if a table contains 10,000 files, the permission granting lasts about 2 minutes.

**Table 10-7** Setting a role

Task	Role Authorization
Setting the Hive administrator permission	<p>In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Hive</b> and select <b>Hive Admin Privilege</b>.</p> <p><b>NOTE</b> After being bound to the Hive administrator role, perform the following operations during each maintenance operation:</p> <ol style="list-style-type: none"> <li>1. Log in to the node where the Hive client is installed as the client installation user.</li> <li>2. Run the following command to configure environment variables: For example, if the Hive client installation directory is <b>/opt/hiveclient</b>, run <b>source /opt/hiveclient/bigdata_env</b>.</li> <li>3. Run the following command to authenticate the user: <b>kinit Hive service user</b></li> <li>4. Run the following command to log in to the client tool: <b>beeline</b></li> <li>5. Run the following command to update the administrator permissions of the Hive user: <b>set role admin;</b></li> </ol>
Setting the permission to query a table of another user in the default database	<ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Hive</b> &gt; <b>Hive Read Write Privileges</b>.</li> <li>2. Click the name of the specified database in the database list. Tables in the database are displayed.</li> <li>3. In the <b>Rights</b> column of the specified table, choose <b>Select</b>.</li> </ol>

Task	Role Authorization
Setting the permission to query a table of another user in the default database	<ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Hive</b> &gt; <b>Hive Read Write Privileges</b>.</li> <li>2. Click the name of the specified database in the database list. Tables in the database are displayed.</li> <li>3. In the <b>Permission</b> column of the specified table, select <b>INSERT</b>.</li> </ol>
Setting the permission to import data to a table of another user in the default database	<ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Hive</b> &gt; <b>Hive Read Write Privileges</b>.</li> <li>2. Click the name of the specified database in the database list. Tables in the database are displayed.</li> <li>3. In the <b>Permission</b> column of the specified indexes, select <b>DELETE</b> and <b>INSERT</b>.</li> </ol>
Setting the permission to submit HQL commands to Yarn for execution	<p>The HQL commands used by some services are converted into MapReduce tasks and submitted to Yarn for execution. You need to set the Yarn permissions. For example, the HQL statements to be run use statements, such as <b>insert</b>, <b>count</b>, <b>distinct</b>, <b>group by</b>, <b>order by</b>, <b>sort by</b>, or <b>join</b>.</p> <ol style="list-style-type: none"> <li>1. In the <b>Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Yarn</b> &gt; <b>Scheduling Queue</b> &gt; <b>root</b>.</li> <li>2. In the <b>Permission</b> column of the <b>default</b> queue, select <b>Submit</b>.</li> </ol>

**Step 5** Click **OK**, and return to the **Role** page.

----End

### 10.4.3 Configuring Permissions for Hive Tables, Columns, or Databases

#### Scenario

You can configure related permissions if you need to access tables or databases created by other users. Hive supports column-based permission control. If a user

needs to access some columns in tables created by other users, the user must be granted the permission for columns. The following describes how to grant table, column, and database permissions to users by using the role management function of MRS Manager.

 **NOTE**

- You can configure permissions for Hive tables, columns, or databases only in security mode.
- MRS 3.x or later supports Ranger. If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for Hive](#).

## Prerequisites

- You have obtained a user account with the administrator permissions, such as **admin**.
- You have created a role, for example, **hrole**, on Manager by referring to instructions in [Creating a Hive Role](#). You do not need to set the Hive permission but need to set the permission to submit the HQL command to Yarn for execution.
- You have created two Hive human-machine users, such as **huser1** and **huser2**, on Manager and added them to the **hive** group. **huser2** has been bound to **hrole**. The **hdb** database has created by user **huser1** and the **htable** table has been created in the database.

## Procedure

- Granting Table Permissions  
Users have complete permission on the tables created by themselves in Hive and the HDFS. To access the tables created by others, they need to be granted the permission. After the Hive metadata permission is granted, the HDFS permission is automatically granted. The procedure for granting a role the permission of querying, inserting, and deleting **htable** data is as follows:  
For versions earlier than MRS 3.x, perform the following operations to grant table permissions:
  - a. On MRS Manager, choose **System > Permission > Manage Role**.
  - b. Locate the row that contains **hrole**, and click **Modify**.
  - c. Choose **Hive > Hive Read Write Privileges**.
  - d. Click the name of the specified database **hdb** in the database list. Table **htable** in the database is displayed.
  - e. In the **Permission** column of the **htable** table, select **Select, Insert, and Delete**.
  - f. Click **OK**.For MRS 3.x or later, perform the following operations to grant table permissions:
  - a. On FusionInsight Manager, choose **System > Permission > Role**.
  - b. Locate the row that contains **hrole**, and click **Modify**.
  - c. Choose *Name of the desired cluster* > **Hive > Hive Read Write Privileges**.

- d. Click the name of the specified database **hdb** in the database list. Table **htable** in the database is displayed.
- e. In the **Permission** column of the **htable** table, select **SELECT**, **INSERT**, and **DELETE**.
- f. Click **OK**.

 **NOTE**

In role management, the procedure for granting a role the permission of querying, inserting, and deleting Hive external table data is the same. After the metadata permission is granted, the HDFS permission is automatically granted.

- **Granting Column Permissions**

Users have all permissions for the tables created by themselves in Hive and HDFS. Users do not have the permission to access the tables created by others. If a user needs to access some columns in tables created by other users, the user must be granted the permission for columns. After the Hive metadata permission is granted, the HDFS permission is automatically granted. The procedure for granting a role the permission of querying and inserting data in **hcol** of **htable** is as follows:

For versions earlier than MRS 3.x, perform the following operations to grant column permissions:

- a. On MRS Manager, choose **System > Permission > Manage Role**.
- b. Locate the row that contains **hrole**, and click **Modify**.
- c. Choose **Hive > Hive Read Write Privileges**.
- d. In the database list, click the specified database **hdb** to display the **htable** table in the database. Click the **htable** table to display the **hcol** column in the table.
- e. In the **Permission** column of the **hcol** column, select **Select** and **Insert**.
- f. Click **OK**.

For MRS 3.x or later, perform the following operations:

- a. On FusionInsight Manager, choose **System > Permission > Role**.
- b. Locate the row that contains **hrole**, and click **Modify**.
- c. Choose *Name of the desired cluster* > **Hive > Hive Read Write Privileges**.
- d. In the database list, click the specified database **hdb** to display the **htable** table in the database. Click the **htable** table to display the **hcol** column in the table.
- e. In the **Permission** column of the **hcol** column, select **SELECT** and **INSERT**.
- f. Click **OK**.

 **NOTE**

In role management, after the metadata permission is granted, the HDFS permission is automatically granted. Therefore, after the column permission is granted, the HDFS ACL permission for all files of the table is automatically granted.

- **Granting Database Permissions**

Users have complete permission on the databases created by themselves in Hive and the HDFS. To access the databases created by others, they need to

be granted the permission. After the Hive metadata permission is granted, the HDFS permission is automatically granted. The procedure for granting a role the permission of querying data and creating tables in database **hdb** is as follows. Other types of database operation permission are not supported.

For versions earlier than MRS 3.x, perform the following database authorization operations:

- a. On MRS Manager, choose **System > Permission > Manage Role**.
- b. Locate the row that contains **hrole**, and click **Modify**.
- c. Choose **Hive > Hive Read Write Privileges**.
- d. In the **Permission** column of the **hdb** database, select **Select** and **Create**.
- e. Click **OK**.

For MRS 3.x or later, perform the following operations to grant database permissions:

- a. On FusionInsight Manager, choose **System > Permission > Role**.
- b. Locate the row that contains **hrole**, and click **Modify**.
- c. Choose *Name of the desired cluster* > **Hive > Hive Read Write Privileges**.
- d. In the **Permission** column of the **hdb** database, select **SELECT** and **CREATE**.
- e. Click **OK**.

 **NOTE**

- Any permission for a table in the database is automatically associated with the HDFS permission for the database directory to facilitate permission management. When any permission for a table is canceled, the system does not automatically cancel the HDFS permission for the database directory to ensure performance. In this case, users can only log in to the database and view table names.
- When the query permission on a database is added to or deleted from a role, the query permission on tables in the database is automatically added to or deleted from the role.

## Concepts

**Table 10-8** Scenarios of using Hive tables, columns, or databases

Scenario	Required Permission
DESCRIBE TABLE	SELECT
SHOW PARTITIONS	SELECT
ANALYZE TABLE	SELECT and INSERT
SHOW COLUMNS	SELECT
SHOW TABLE STATUS	SELECT
SHOW TABLE PROPERTIES	SELECT
SELECT	SELECT
EXPLAIN	SELECT

Scenario	Required Permission
CREATE VIEW	SELECT, Grant Of Select, and CREATE
SHOW CREATE TABLE	SELECT and Grant Of Select
CREATE TABLE	CREATE
ALTER TABLE ADD PARTITION	INSERT
INSERT	INSERT
INSERT OVERWRITE	INSERT and DELETE
LOAD	INSERT and DELETE
ALTER TABLE DROP PARTITION	DELETE
CREATE FUNCTION	Hive Admin Privilege
DROP FUNCTION	Hive Admin Privilege
ALTER DATABASE	Hive Admin Privilege

## 10.4.4 Configuring Permissions to Use Other Components for Hive

### Scenario

Hive may need to be associated with other components. For example, Yarn permissions are required in the scenario of using HQL statements to trigger MapReduce jobs, and HBase permissions are required in the Hive over HBase scenario. The following describes the operations in the two scenarios.

#### NOTE

- In security mode, Yarn and HBase permission management is enabled by default. Therefore, Yarn and HBase permissions need to be configured by default.
- In common mode, Yarn and HBase permission management is disabled by default. That is, any user has permissions. Therefore, YARN and HBase permissions does not need to be configured by default. If a user enables the permission management by modifying the Yarn or HBase configurations, the Yarn and HBase permissions then need to be configured.
- MRS 3.x or later supports Ranger. If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for Hive](#).

### Prerequisites

- The Hive client has been installed. For example, the installation directory is `/opt/client`.
- You have obtained a user account with the administrator permissions, such as `admin`.



## Procedure

### Association with Yarn in MRS Earlier than 3.x

Yarn permissions are required when HQL statements, such as **insert**, **count**, **distinct**, **group by**, **order by**, **sort by**, and **join**, are used to trigger MapReduce jobs. The following uses the procedure for assigning a role the permissions to run the **count** statements in the **thc** table as an example.

- Step 1** Create a role on MRS Manager.
- Step 2** In the **Permission** table, choose **Yarn > Scheduler Queue > root**.
- Step 3** In the **Permission** column of the default queue, select **Submit** and click **OK**.
- Step 4** In the **Permission** table, choose **Hive > Hive Read Write Privileges > default**, select **Select** for **thc**, and click **OK**.

----End

### Association with Yarn in MRS 3.x or Later

Yarn permissions are required when HQL statements, such as **insert**, **count**, **distinct**, **group by**, **order by**, **sort by**, and **join**, are used to trigger MapReduce jobs. The following uses the procedure for assigning a role the permissions to run the **count** statements in the **thc** table as an example.

- Step 1** Create a role on FusionInsight Manager.
- Step 2** In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **Yarn > Scheduler Queue > root**.
- Step 3** In the **Permission** column of the **default** queue, select **Submit** and click **OK**.
- Step 4** In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **Hive > Hive Read Write Privileges > default**. Select **SELECT** for table **thc**, and click **OK**.

----End

### Hive over HBase Authorization in MRS Earlier than 3.x

After the permissions are assigned, you can use HQL statements that are similar to SQL statements to access HBase tables from Hive. The following uses the procedure for assigning a user the rights to query HBase tables as an example.

- Step 1** On the role management page of MRS Manager, create an HBase role, for example, **hive\_hbase\_create**, and grant the permission to create HBase tables.  
  
In the **Permission** table, choose **HBase > HBase Scope > global**, select **create** of the namespace **default**, and click **OK**.
- Step 2** On MRS Manager, create a human-machine user, for example, **hbase\_creates\_user**, add the user to the **hive** group, and bind the **hive\_hbase\_create** role to the user so that the user can create Hive and HBase tables.
- Step 3** Log in to the node where the client is installed.
- Step 4** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

**Step 5** Run the following command to authenticate the user:

```
kinit hbase_creates_user
```

**Step 6** Run the following command to go to the shell environment of the Hive client:

```
beeline
```

**Step 7** Run the following command to create a table in Hive and HBase, for example, the **thh** table.

```
CREATE TABLE thh(id int, name string, country string) STORED BY
'org.apache.hadoop.hive.hbase.HBaseStorageHandler' WITH
SERDEPROPERTIES("hbase.columns.mapping" = "cf1:id,cf1:name,:key")
TBLPROPERTIES ("hbase.table.name" = "thh");
```

The created Hive table and the HBase table are stored in the Hive database **default** and the HBase namespace **default**, respectively.

**Step 8** On the role management page of MRS Manager, create a role, for example, **hive\_hbase\_select**, and assign the role the permission to query the Hive table **thh** and the HBase table **thh**.

1. In the **Permission** table, choose **HBase > HBase Scope > global > default**, select **Read** for the **thh** table, and click **OK** to grant the HBase role the permission to query the table.
2. Edit a role. In the **Permission** table, choose **HBase > HBase Scope > global > hbase**. Select **Execute** for **hbase:meta**, and click **OK**.
3. Edit a role. In the **Permission** table, choose **Hive > Hive Read Write Privileges > default**, select **Select** for **thh**, and click **OK**.

**Step 9** On MRS Manager, create a human-machine user, for example, **hbase\_select\_user**, add the user to the **hive** group, and bind the **hive\_hbase\_select** role to the user so that the user can query Hive and HBase tables.

**Step 10** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

**Step 11** Run the following command to authenticate users:

```
kinit hbase_select_user
```

**Step 12** Run the following command to go to the shell environment of the Hive client:

```
beeline
```

**Step 13** Run the following command to use an HQL statement to query HBase table data:

```
select * from thh;
```

```
----End
```

#### **Hive over HBase Authorization in MRS 3.x or Later**

After the permissions are assigned, you can use HQL statements that are similar to SQL statements to access HBase tables from Hive. The following uses the procedure for assigning a user the rights to query HBase tables as an example.

**Step 1** On the role management page of FusionInsight Manager, create an HBase role, for example, **hive\_hbase\_create**, and grant the permission to create HBase tables.

In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **HBase** > **HBase Scope** > **global**. Select **Create** of the namespace **default**, and click **OK**.

**Step 2** On FusionInsight Manager, create a human-machine user, for example, **hbase\_creates\_user**, add the user to the **hive** group, and bind the **hive\_hbase\_create** role to the user so that the user can create Hive and HBase tables.

**Step 3** If the current component uses Ranger for permission control, grant the create permission for **hive\_hbase\_create** or **hbase\_creates\_user**. For details, see [Adding a Ranger Access Permission Policy for Hive](#).

**Step 4** Log in to the node where the client is installed as the client installation user.

**Step 5** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

**Step 6** Run the following command to authenticate the user:

```
kinit hbase_creates_user
```

**Step 7** Run the following command to go to the shell environment of the Hive client:

```
beeline
```

**Step 8** Run the following command to create a table in Hive and HBase, for example, the **thh** table.

```
CREATE TABLE thh(id int, name string, country string) STORED BY
'org.apache.hadoop.hive.hbase.HBaseStorageHandler' WITH
SERDEPROPERTIES("hbase.columns.mapping" = "cf1:id,cf1:name,:key")
TBLPROPERTIES ("hbase.table.name" = "thh");
```

The created Hive table and the HBase table are stored in the Hive database **default** and the HBase namespace **default**, respectively.

**Step 9** On the role management page of FusionInsight Manager, create a role, for example, **hive\_hbase\_select**, and assign the role the permission to query the Hive table **thh** and the HBase table **thh**.

1. In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **HBase** > **HBase Scope** > **global** > **default**. Select **read** of the **thh** table, and click **OK** to grant the table query permission to the HBase role.
2. Edit the role. In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **HBase** > **HBase Scope** > **global** > **hbase**, select **Execute** for **hbase:meta**, and click **OK**.
3. Edit the role. In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **Hive** > **Hive Read Write Privileges** > **default**. Select **SELECT** for the **thh** table, and click **OK**.

**Step 10** On FusionInsight Manager, create a human-machine user, for example, **hbase\_select\_user**, add the user to the **hive** group, and bind the **hive\_hbase\_select** role to the user so that the user can query Hive and HBase tables.

**Step 11** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

**Step 12** Run the following command to authenticate users:

```
kinit hbase_select_user
```

**Step 13** Run the following command to go to the shell environment of the Hive client:

```
beeline
```

**Step 14** Run the following command to use an HQL statement to query HBase table data:

```
select * from thh;
```

```
----End
```

## 10.5 Using a Hive Client

### Scenario

This section guides users to use a Hive client in an O&M or service scenario.

### Prerequisites

- The client has been installed. For example, the client is installed in the **/opt/hadoopclient** directory. The client directory in the following operations is only an example. Change it to the actual installation directory.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login.

### Using the Hive Client (Versions Earlier Than MRS 3.x)

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/hadoopclient
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Log in to the Hive client based on the cluster authentication mode.

- In security mode, run the following command to complete user authentication and log in to the Hive client:

```
kinit Component service user
```

```
beeline
```

- In common mode, run the following command to log in to the Hive client. If no component service user is specified, the current OS user is used to log in to the Hive client.

```
beeline -n component service user
```

 NOTE

After a beeline connection is established, you can compile and submit HQL statements to execute related tasks. To run the Catalog client command, you need to run the **!q** command first to exit the beeline environment.

**Step 5** Run the following command to execute the HCatalog client command:

```
hcat -e "cmd"
```

*cmd* must be a Hive DDL statement, for example, **hcat -e "show tables"**.

 NOTE

- To use the HCatalog client, choose **More > Download Client** on the service page to download the clients of all services. This restriction does not apply to the beeline client.
- Due to permission model incompatibility, tables created using the HCatalog client cannot be accessed on the HiveServer client. However, the tables can be accessed on the WebHCat client.
- If you use the HCatalog client in Normal mode, the system performs DDL commands using the current user who has logged in to the operating system.
- Exit the beeline client by running the **!q** command instead of by pressing **Ctrl + c**. Otherwise, the temporary files generated by the connection cannot be deleted and a large number of junk files will be generated as a result.
- If multiple statements need to be entered during the use of beeline clients, separate the statements from each other using semicolons (;) and set the value of **entireLineAsCommand** to **false**.

Setting method: If beeline has not been started, run the **beeline --entireLineAsCommand=false** command. If the beeline has been started, run the **!set entireLineAsCommand false** command.

After the setting, if a statement contains semicolons (;) that do not indicate the end of the statement, escape characters must be added, for example, **select concat\_ws('\;', collect\_set(col1)) from tbl**.

----End

## Using the Hive Client (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 go to the client installation directory:

```
cd /opt/hadoopclient
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Log in to the Hive client based on the cluster authentication mode.

- In security mode, run the following command to complete user authentication and log in to the Hive client:

```
kinit Component service user
```

```
beeline
```

- In common mode, run the following command to log in to the Hive client. If no component service user is specified, the current OS user is used to log in to the Hive client.

**beeline -n** *component service user*

**Step 5** Run the following command to execute the HCatalog client command:

**hcat -e** "*cmd*"

*cmd* must be a Hive DDL statement, for example, **hcat -e "show tables"**.

 **NOTE**

- To use the HCatalog client, choose **More > Download Client** on the service page to download the clients of all services. This restriction does not apply to the beeline client.
- Due to permission model incompatibility, tables created using the HCatalog client cannot be accessed on the HiveServer client. However, the tables can be accessed on the WebHCat client.
- If you use the HCatalog client in Normal mode, the system performs DDL commands using the current user who has logged in to the operating system.
- Exit the beeline client by running the **!q** command instead of by pressing **Ctrl + C**. Otherwise, the temporary files generated by the connection cannot be deleted and a large number of junk files will be generated as a result.
- If multiple statements need to be entered during the use of beeline clients, separate the statements from each other using semicolons (;) and set the value of **entireLineAsCommand** to **false**.

Setting method: If beeline has not been started, run the **beeline --entireLineAsCommand=false** command. If the beeline has been started, run the **!set entireLineAsCommand false** command.

After the setting, if a statement contains semicolons (;) that do not indicate the end of the statement, escape characters must be added, for example, **select concat\_ws('\;', collect\_set(col1)) from tbl**.

----End

## Common Hive Client Commands

The following table lists common Hive Beeline commands.

**Table 10-9** Common Hive Beeline commands

Command	Description
set <key>=<value>	Sets the value of a specific configuration variable (key). <b>NOTE</b> If the variable name is incorrectly spelled, the Beeline does not display an error.
set	Prints the list of configuration variables overwritten by users or Hive.
set -v	Prints all configuration variables of Hadoop and Hive.

Command	Description
add FILE[S] <filepath> <filepath>*add JAR[S] <filepath> <filepath>*add ARCHIVE[S] <filepath> <filepath>*	Adds one or more files, JAR files, or ARCHIVE files to the resource list of the distributed cache.
add FILE[S] <ivyurl> <ivyurl>* add JAR[S] <ivyurl> <ivyurl>* add ARCHIVE[S] <ivyurl> <ivyurl>*	Adds one or more files, JAR files, or ARCHIVE files to the resource list of the distributed cache using the Ivy URL in the <b>ivy://goup:module:version?query_string</b> format.
list FILE[S]list JAR[S]list ARCHIVE[S]	Lists the resources that have been added to the distributed cache.
list FILE[S] <filepath>*list JAR[S] <filepath>*list ARCHIVE[S] <filepath>*	Checks whether given resources have been added to the distributed cache.
delete FILE[S] <filepath>*delete JAR[S] <filepath>*delete ARCHIVE[S] <filepath>*	Deletes resources from the distributed cache.
delete FILE[S] <ivyurl> <ivyurl>* delete JAR[S] <ivyurl> <ivyurl>* delete ARCHIVE[S] <ivyurl> <ivyurl>*	Delete the resource added using <ivyurl> from the distributed cache.
reload	Enable HiveServer2 to discover the change of the JAR file <b>hive.reloadable.aux.jars.path</b> in the specified path. (You do not need to restart HiveServer2.) Change actions include adding, deleting, or updating JAR files.
dfs <dfs command>	Runs the <b>dfs</b> command.
<query string>	Executes the Hive query and prints the result to the standard output.

## 10.6 Using HDFS Colocation to Store Hive Tables

### Scenario

HDFS Colocation is the data location control function provided by HDFS. The HDFS Colocation API stores associated data or data on which associated

operations are performed on the same storage node. Hive supports the HDFS Colocation function. When Hive tables are created, after the locator information is set for table files, data files of related tables are stored on the same storage node when data is inserted into tables using the insert statement (other data import modes are not supported). This ensures convenient and efficient data computing among associated tables. The supported table formats are only TextFile and RCFile.

 **NOTE**

This section applies to MRS 3.x or later.

## Procedure

**Step 1** Log in to the node where the client is installed as a client installation user.

**Step 2** Run the following command to switch to the client installation directory, for example, `/opt/client`:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** If the cluster is in security mode, run the following command to authenticate the user:

```
kinit MRS username
```

**Step 5** Create the `groupid` through the HDFS API.

```
hdfs colocationadmin -createGroup -groupId <groupid> -locatorIds
<locatorid1>,<locatorid2>,<locatorid3>
```

 **NOTE**

In the preceding command, `<groupid>` indicates the name of the created group. The group created in this example contains three locators. You can define the number of locators as required.

For details about group ID creation and HDFS Colocation, see HDFS description.

**Step 6** Run the following command to log in to the Hive client:

```
beeline
```

**Step 7** Enable Hive to use colocation.

Assume that `table_name1` and `table_name2` are associated with each other. Run the following statements to create them:

```
CREATE TABLE <[db_name.]table_name1>[(col_name data_type , ...)] [ROW
FORMAT <row_format>] [STORED AS <file_format>]
TBLPROPERTIES("groupId"=" <group> ","locatorId"=" <locator1>");
```

```
CREATE TABLE <[db_name.]table_name2> [(col_name data_type , ...)] [ROW
FORMAT <row_format>] [STORED AS <file_format>]
TBLPROPERTIES("groupId"=" <group> ","locatorId"=" <locator1>");
```

After data is inserted into `table_name1` and `table_name2` using the insert statement, data files of `table_name1` and `table_name2` are distributed to the



same storage position in the HDFS, facilitating associated operations among the two tables.

----End

## 10.7 Using the Hive Column Encryption Function

### Scenario

Hive supports encryption of one or multiple columns in a table. When creating a Hive table, you can specify the column to be encrypted and encryption algorithm. When data is inserted into the table using the insert statement, the related columns are encrypted. Column encryption can be performed in HDFS tables of only the TextFile and SequenceFile file formats. The Hive column encryption does not support views and the Hive over HBase scenario.

Hive supports two column encryption algorithms, which can be specified during table creation:

- AES (the encryption class is org.apache.hadoop.hive.serde2.AESRewriter)
- SMS4 (the encryption class is org.apache.hadoop.hive.serde2.SMS4Rewriter)

#### NOTE

- In national cryptographic cluster scenarios, Hive column encryption supports only table creation using the SMS4 algorithm.
- When you import data from a common Hive table into a Hive column encryption table, you are advised to delete the original data from the common Hive table as long as doing this does not affect other services. Retaining an unencrypted table poses security risks.

### Procedure

- Step 1** Specify the column to be encrypted and encryption algorithm when creating a table.

```
create table <[db_name.]table_name> (<col_name1>
<data_type>, <col_name2> <data_type>, <col_name3>
<data_type>, <col_name4> <data_type>) ROW FORMAT SERDE
'org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe' WITH
SERDEPROPERTIES ('column.encode.columns'='<col_name2>,<col_name3>',
'column.encode.classname'='org.apache.hadoop.hive.serde2.AESRewriter')STO
RED AS TEXTFILE;
```

Alternatively, use the following statement:

```
create table <[db_name.]table_name> (<col_name1>
<data_type>, <col_name2> <data_type>, <col_name3>
<data_type>, <col_name4> <data_type>) ROW FORMAT SERDE
'org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe' WITH
SERDEPROPERTIES ('column.encode.indices'='1,2',
'column.encode.classname'='org.apache.hadoop.hive.serde2.SMS4Rewriter')
STORED AS TEXTFILE;
```

 NOTE

- The numbers used to specify encryption columns start from 0. 0 indicates column 1, 1 indicates column 2, and so on.
- When creating a table with encrypted columns, ensure that the directory where the table resides is empty.

**Step 2** Insert data into the table using the insert statement.

Assume that the test table exists and contains data.

```
insert into table <table_name> select <col_list> from test;
```

----End

## 10.8 Customizing Row Separators

### Scenario

In most cases, a carriage return character is used as the row delimiter in Hive tables stored in text files, that is, the carriage return character is used as the terminator of a row during queries. However, some data files are delimited by special characters, and not a carriage return character.

MRS Hive allows you to use different characters or character combinations to delimit rows of Hive text data. When creating a table, set **inputformat** to **SpecifiedDelimiterInputFormat**, and set the following parameter before search each time. Then the table data is queried by the specified delimiter.

```
set hive.textinput.record.delimiter="";
```

 NOTE

- The Hue component of the current version does not support the configuration of multiple separators when files are imported to a Hive table.
- This section applies to MRS 3.x or later.

### Procedure

**Step 1** Specify **inputFormat** and **outputFormat** when creating a table.

```
CREATE [TEMPORARY] [EXTERNAL] TABLE [IF NOT EXISTS]
[db_name.]table_name [(col_name data_type [COMMENT col_comment], ...)]
[ROW FORMAT row_format] STORED AS inputformat
'org.apache.hadoop.hive.contrib.fileformat.SpecifiedDelimiterInputFormat'
outputformat 'org.apache.hadoop.hive.ql.io.HiveIgnoreKeyTextOutputFormat'
```

**Step 2** Specify the delimiter before search.

```
set hive.textinput.record.delimiter="!@"
```

Hive will use '!@' as the row delimiter.

----End

## 10.9 Configuring Hive on HBase in Across Clusters with Mutual Trust Enabled

For mutually trusted Hive and HBase clusters with Kerberos authentication enabled, you can access the HBase cluster and synchronize its key configurations to HiveServer of the Hive cluster.

### Prerequisites

The mutual trust relationship has been configured between the two security clusters with Kerberos authentication enabled.

### Procedure for Configuring Hive on HBase Across Clusters

**Step 1** Download the HBase configuration file and decompress it.

1. Log in to FusionInsight Manager of the target HBase cluster, click **Cluster** and choose **Services > HBase**.
2. Choose **More > Download Client**.
3. Download the HBase configuration file and choose **Configuration Files only** for **Select Client Type**.

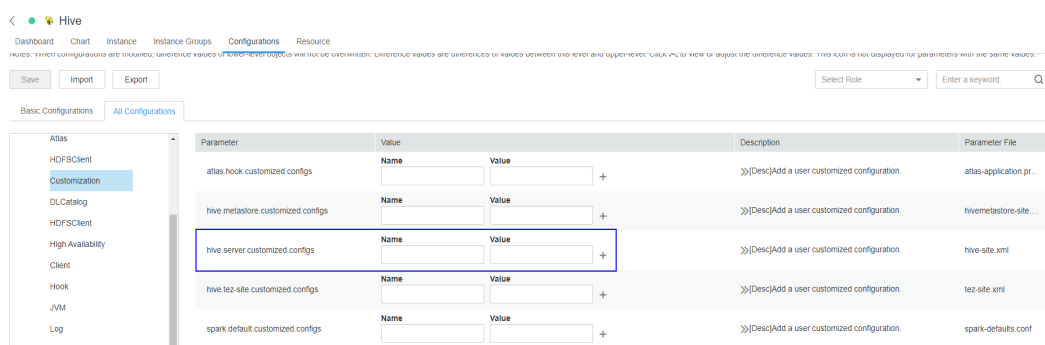
**Step 2** Log in to FusionInsight Manager of the source Hive cluster.

**Step 3** Click **Cluster**, choose **Services > Hive**, click **Configurations** and then **All Configurations**. On the displayed page, add the following parameters to the **hive-site.xml** configuration file of the HiveServer role.

Search for the following parameters in the **hbase-site.xml** configuration file of the downloaded HBase client and add them to HiveServer:

- hbase.security.authentication
- hbase.security.authorization
- hbase.zookeeper.property.clientPort
- hbase.zookeeper.quorum (The domain name needs to be converted into an IP address.)
- hbase.regionserver.kerberos.principal
- hbase.master.kerberos.principal

**Figure 10-1** Custom configurations of the HiveServer role



**Step 4** Save the configurations and restart Hive.

----End

## 10.10 Deleting Single-Row Records from Hive on HBase

### Scenario

Due to the limitations of underlying storage systems, Hive does not support the ability to delete a single piece of table data. In Hive on HBase, MRS Hive supports the ability to delete a single piece of HBase table data. Using a specific syntax, Hive can delete one or more pieces of data from an HBase table.

**Table 10-10** Permissions required for deleting single-row records from the Hive on HBase table

Cluster Authentication Mode	Required Permission
Security mode	SELECT, INSERT, and DELETE
Common mode	None

### Procedure

**Step 1** To delete some data from an HBase table, run the following HQL statement:

```
remove table <table_name> where <expression>;
```

In the preceding information, *<expression>* specifies the filter condition of the data to be deleted. *<table\_name>* indicates the Hive on HBase table from which data is to be deleted.

----End

## 10.11 Configuring HTTPS/HTTP-based REST APIs

### Scenario

WebHCat provides external REST APIs for Hive. By default, the open-source community version uses the HTTP protocol.

MRS Hive supports the HTTPS protocol that is more secure, and enables switchover between the HTTP protocol and the HTTPS protocol.

#### NOTE

The security mode supports HTTPS and HTTP, and the common mode supports only HTTP.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Modify the Hive configuration.

- For versions earlier than MRS 3.x: Enter the parameter name in the search box, search for **templeton.protocol.type**, change the parameter value to **HTTPS** or **HTTP**, and restart the Hive service to use the corresponding protocol.
- For MRS 3.x or earlier: Choose **WebHCat > Security**. On the page that is displayed, select **HTTPS** or **HTTP**. After the modification, restart the Hive service to use the corresponding protocol.

----End

## 10.12 Enabling or Disabling the Transform Function

### Scenario

The Transform function is not allowed by Hive of the open source version.

MRS Hive supports the configuration of the Transform function. The function is disabled by default, which is the same as that of the open-source community version.

Users can modify configurations of the Transform function to enable the function. However, security risks exist when the Transform function is enabled.

 **NOTE**

The Transform function can be disabled only in security mode.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Enter the parameter name in the search box, search for **hive.security.transform.disallow**, change the parameter value to **true** or **false**, and restart all HiveServer instances.

 **NOTE**

- If this parameter is set to **true**, the Transform function is disabled, which is the same as that in the open-source community version.
- If this parameter is set to **false**, the Transform function is enabled, which poses security risks.

----End

## 10.13 Access Control of a Dynamic Table View on Hive

### Scenario

This section describes how to create a view on Hive when MRS is configured in security mode, authorize access permissions to different users, and specify that different users access different data.

In the view, Hive can obtain the built-in function **current\_user()** of the users who submit tasks on the client and filter the users. This way, authorized users can only access specific data in the view.

 **NOTE**

In normal mode, the **current\_user()** function cannot distinguish users who submit tasks on the client. Therefore, the access control function takes effect only for Hive in security mode. If the **current\_user()** function is used in the actual service logic, the possible risks must be fully evaluated during the conversion between the security mode and normal mode.

### Operation Example

- If the **current\_user** function is not used, different views need to be created for different users to access different data.
  - Authorize the view **v1** permission to user **hiveuser1**. The user **hiveuser1** can access data with **type** set to **hiveuser1** in **table1**.  
**create view v1 as select \* from table1 where type='hiveuser1'**
  - Authorize the view **v2** permission to user **hiveuser2**. The user **hiveuser2** can access data with **type** set to **hiveuser2** in **table1**.  
**create view v2 as select \* from table1 where type='hiveuser2'**
- If the **current\_user** function is used, only one view needs to be created. Authorize the view **v** permission to users **hiveuser1** and **hiveuser2**. When user **hiveuser1** queries view **v**, the **current\_user()** function is automatically converted to **hiveuser1**. When user **hiveuser2** queries view **v**, the **current\_user()** function is automatically converted to **hiveuser2**.

```
create view v as select * from table1 where type=current_user()
```

## 10.14 Specifying Whether the ADMIN Permissions Is Required for Creating Temporary Functions

### Scenario

You must have **ADMIN** permission when creating temporary functions on Hive of the open source community version.

MRS Hive supports the configuration of the function for creating temporary functions with **ADMIN** permission. The function is disabled by default, which is the same as that of the open-source community version.

You can modify configurations of this function. After the function is enabled, you can create temporary functions without **ADMIN** permission. If this parameter is set to **false**, security risks exist.

#### NOTE

The security mode supports the configuration of whether the ADMIN permission is required for creating temporary functions, but the common mode does not support this function.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Enter the parameter name in the search box, search for **hive.security.temporary.function.need.admin**, change the parameter value to **true** or **false**, and restart all HiveServer instances.

#### NOTE

- If this parameter is set to **true**, the ADMIN permission is required for creating temporary functions, which is the same as that in the open source community.
- If this parameter is set to **false**, the ADMIN permission is not required for creating temporary functions.

----End

## 10.15 Using Hive to Read Data in a Relational Database

### Scenario

Hive allows users to create external tables to associate with other relational databases. External tables read data from associated relational databases and support Join operations with other tables in Hive.

Currently, the following relational databases can use Hive to read data:

- DB2
- Oracle

#### NOTE

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

### Prerequisites

The Hive client has been installed.

### Procedure

**Step 1** Log in to the node where the Hive client is installed as the Hive client installation user .

**Step 2** Run the following command to go to the client installation directory:

```
cd Client installation directory
```

For example, if the client installation directory is **/opt/client**, run the following command:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Check whether the cluster authentication mode is Security.

- If yes, run the following command to authenticate the user:  

```
kinit Hive service user
```
- If no, go to [Step 5](#).

**Step 5** Run the following command to upload the driver JAR package of the relational database to be associated to an HDFS directory.

```
hdfs dfs -put directory where the JAR package is located HDFS directory to which the JAR is uploaded
```

For example, to upload the Oracle driver JAR package in **/opt** to the **/tmp** directory in HDFS, run the following command:

```
hdfs dfs -put /opt/ojdbc6.jar /tmp
```



**Step 6** Create an external table on the Hive client to associate with the relational database, as shown in the following example.

 **NOTE**

```
If the security mode is used, the user who creates the table must have the ADMIN
permission. The ADD JAR path is subject to the actual path.
-- Example of associating with an Oracle Linux 6 database
-- In security mode, set the admin permission.
set role admin;
-- Upload the driver JAR package of the relational database to be associated. The driver JAR packages
vary according to databases.
ADD JAR hdfs:///tmp/ojdbc6.jar;

CREATE EXTERNAL TABLE ora_test
-- The Hive table must have one more column than the database return result. This column is used for
paging query.
(id STRING,rownum string)
STORED BY 'com.qubitproducts.hive.storage.jdbc.JdbcStorageHandler'
TBLPROPERTIES (
-- Relational database table type
"qubit.sql.database.type" = "ORACLE",
-- Connect to the URL of the relational database through JDBC. (The URL formats vary according to
databases.)
"qubit.sql.jdbc.url" = "jdbc:oracle:thin:@//10.163.0.1:1521/mydb",
-- Relational database driver class type
"qubit.sql.jdbc.driver" = "oracle.jdbc.OracleDriver",
-- SQL statement queried in the relational database. The result is returned to the Hive table.
"qubit.sql.query" = "select name from aaa",
-- (Optional) Match the Hive table columns to the relational database table columns.
"qubit.sql.column.mapping" = "id=name",
-- Relational database user
"qubit.sql.dbc.username" = "test",
-- Relational database password
"qubit.sql.dbc.password" = "xxx");
```

----End

## 10.16 Supporting Traditional Relational Database Syntax in Hive

### Overview

Hive supports the following types of traditional relational database syntax:

- Grouping
- EXCEPT and INTERSECT

### Grouping

Syntax description:

- Grouping takes effect only when the Group by statement contains ROLLUP or CUBE.
- The result set generated by CUBE contains all the combinations of values in the selected columns.
- The result set generated by ROLLUP contains the combinations of a certain layer structure in the selected columns.

- Grouping: If a row is added by using the CUBE or ROLLUP operator, the output value of the added row is 1. If the row is not added by using the CUBE or ROLLUP operator, the output value of the added row is 0.

For example, the **table\_test** table exists in Hive and the table structure is as follows:

```
+-----+-----+
| table_test.id | table_test.value |
+-----+-----+
| 1 | 10 |
| 1 | 15 |
| 2 | 20 |
| 2 | 5 |
| 2 | 13 |
+-----+-----+
```

Run the following statement:

```
select id,grouping(id),sum(value) from table_test group by id with rollup;
```

The result is as follows:

```
+-----+-----+-----+
| id | groupingresult | sum |
+-----+-----+-----+
| 1 | 0 | 25 |
| NULL | 1 | 63 |
| 2 | 0 | 38 |
+-----+-----+-----+
```

## EXCEPT and INTERSECT

Syntax description:

- EXCEPT returns the difference of two result sets (that is, non-duplicated values return only one query).
- INTERSECT returns the intersection of two result sets (that is, non-duplicated values return by both queries).

For example, two tables **test\_table1** and **test\_table2** exist in Hive.

The table structure of **test\_table1** is as follows:

```
+-----+
| test_table1.id |
+-----+
| 1 |
| 2 |
| 3 |
| 4 |
+-----+
```

The table structure of **test\_table2** is as follows:

```
+-----+
| test_table2.id |
+-----+
| 2 |
| 3 |
| 4 |
| 5 |
+-----+
```

- Run the following EXCEPT statement:

```
select id from test_table1 except select id from test_table2;
```

The result is as follows:

```
+-----+--+
|_alias_0.id |
+-----+--+
| 1 |
+-----+--+
```

- Run the following INTERSECT statement:

```
select id from test_table1 intersect select id from test_table2;
```

The result is as follows:

```
+-----+--+
|_alias_0.id |
+-----+--+
| 2 |
| 3 |
| 4 |
+-----+--+
```

## 10.17 Creating User-Defined Hive Functions

When the built-in functions of Hive cannot meet requirements, you can compile user-defined functions (UDFs) and use them in queries.

According to implementation methods, UDFs are classified as follows:

- Common UDFs: used to perform operations on a single data row and export a single data row.
- User-defined aggregating functions (UDAFs): used to input multiple data rows and export a single data row.
- User-defined table-generating functions (UDTFs): used to perform operations on a single data row and export multiple data rows.

According to use methods, UDFs are classified as follows:

- Temporary functions: used only in the current session and must be recreated after a session restarts.
- Permanent functions: used in multiple sessions. You do not need to create them every time a session restarts.

### NOTE

You need to properly control the memory and thread usage of variables in UDFs. Improper control may cause memory overflow or high CPU usage.

The following uses AddDoublesUDF as an example to describe how to compile and use UDFs.

### Function

AddDoublesUDF is used to add two or more floating point numbers. In this example, you can learn how to write and use UDFs.

 NOTE

- A common UDF must be inherited from **org.apache.hadoop.hive.ql.exec.UDF**.
- A common UDF must implement at least one **evaluate()**. The evaluate function supports overloading.
- To develop a UDF, add the **hive-exec-\*.jar** dependency package to the project. You can obtain the package from the Hive service installation directory, for example, **\$ {BIGDATA\_HOME}/components/FusionInsight\_HD\_\*/Hive/disaster/plugin/lib/**.

## Sample Code

The following is a UDF code example:

*xxx* indicates the name of the organization that develops the program.

```
package com.xxx.bigdata.hive.example.udf;
import org.apache.hadoop.hive.ql.exec.UDF;

public class AddDoublesUDF extends UDF {
 public Double evaluate(Double... a) {
 Double total = 0.0;
 // Processing logic
 for (int i = 0; i < a.length; i++)
 if (a[i] != null)
 total += a[i];
 return total;
 }
}
```

## How to Use

- Step 1** Packing programs as **AddDoublesUDF.jar** on the client node, and upload the package to a specified directory in HDFS, for example, **/user/hive\_examples\_jars**.

Both the user who creates the function and the user who uses the function must have the read permission on the file.

The following are example statements:

```
hdfs dfs -put ./hive_examples_jars /user/hive_examples_jars
```

```
hdfs dfs -chmod 777 /user/hive_examples_jars
```

- Step 2** Check the cluster authentication mode.
- In security mode, log in to the beeline client as a user with the Hive management permission and run the following commands:

```
kinit Hive service user
```

```
beeline
```

```
set role admin;
```

- In common mode, run the following command:

```
beeline -n Hive service user
```

- Step 3** Define the function in HiveServer. Run the following SQL statement to create a permanent function:

```
CREATE FUNCTION addDoubles AS
'com.xxx.bigdata.hive.example.udf.AddDoublesUDF' using jar 'hdfs://hacluster/
user/hive_examples_jars/AddDoublesUDF.jar';
```

*addDoubles* is the alias of the function, which is used in **SELECT** queries. *xxx* is typically the name of the organization that develops the program.

Run the following statement to create a temporary function:

```
CREATE TEMPORARY FUNCTION addDoubles AS
'com.xxx.bigdata.hive.example.udf.AddDoublesUDF' using jar 'hdfs://hacluster/
user/hive_examples_jars/AddDoublesUDF.jar';
```

- *addDoubles* indicates the function alias that is used for SELECT query.
- **TEMPORARY** indicates that the function is used only in the current session with the HiveServer.

**Step 4** Run the following SQL statement to use the function on the HiveServer:

```
SELECT addDoubles(1,2,3);
```

 **NOTE**

If an [Error 10011] error is displayed when you log in to the client again, run the **reload function;** command and then use this function.

**Step 5** Run the following SQL statement to delete the function from the HiveServer:

```
DROP FUNCTION addDoubles;
```

----End

## Extended Applications

None

# 10.18 Enhancing beeline Reliability

## Scenario

- When the beeline client is disconnected due to network exceptions during the execution of a batch processing task, tasks submitted before beeline is disconnected can be properly executed in Hive. When you start the batch processing task again, the submitted tasks are not executed and tasks that are not executed are executed in sequence.
- When the HiveServer service breaks down due to some reasons during the execution of a batch processing task, Hive enables that the tasks that have been successfully executed are not executed again when the same batch processing task is started again. The execution starts from the task that has not been executed from the time when HiveServer2 breaks down.

 **NOTE**

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

## Example

1. Beeline is reconnected after being disconnection.  
Example:

```
beeline -e "${SQL}" --hivevar batchid=xxxxx
```

2. Beeline kills the running tasks.

Example:

```
beeline -e "" --hivevar batchid=xxxxx --hivevar kill=true
```

3. Log in to the beeline client and start the mechanism of reconnection after disconnection.

Log in to the beeline client and run the **set hivevar:batchid=xxxxx** command.

#### NOTE

Instructions:

- *xxxx* indicates the batch ID of tasks submitted in the same batch using the beeline client. Batch IDs can be used to identify the task submission batch. If the batch ID is not contained when a task is submitted, this feature is not enabled. The value of *xxxx* is specified during task execution. In the following example, the value of *xxxx* is **012345678901**.

```
beeline -f hdfs://hacluster/user/hive/table.sql --hivevar batchid=012345678901
```

- If the running SQL script depends on the data timeliness, you are advised not to enable the breakpoint reconnection mechanism. You can use a new batch ID to submit tasks. During reexecution of the scripts, some SQL statements have been executed and are not executed again. As a result, expired data is obtained.
- If some built-in time functions are used in the SQL script, it is recommended that you do not enable the breakpoint reconnection mechanism or the use of a new batch ID for each execution. The reason is the same as above.
- A SQL script contains one or more subtasks. If the logic for deleting and creating temporary tables exist in the SQL script, it is recommended that the logic for deleting temporary tables be placed at the end of the script. If the subtasks executed after the temporary table deletion task fail to be executed and the temporary table is used in the subtasks before the temporary table deletion task, when the SQL script is executed using the same batch ID for the next time, the compilation of the subtasks (excluding the task for creating the temporary table because the creation has been completed and is not executed again, and only compilation is allowed) executed before the temporary table deletion task fails because the temporary has been deleted. In this case, you are advised to use a new batch ID to execute the script.

Parameter description:

- **zk.cleanup.finished.job.interval**: indicates the interval for executing the cleanup task. The default interval is 60 seconds.
- **zk.cleanup.finished.job.outdated.threshold**: indicates the threshold of the node validity period. A node is generated for tasks in the same batch. The threshold is calculated from the end time of the execution of the current batch task. If the time exceeds 60 minutes, the node is deleted.
- **batch.job.max.retry.count**: indicates the maximum number of retry times of a batch task. If the number of retry times of a batch task exceeds the value of this parameter, the task execution record is deleted. The task will be executed from the first task when the task is started next time. The default value is **10**.
- **beeline.reconnect.zk.path**: indicates the root node for storing task execution progress. The default value for the Hive service is **/beeline**.

## 10.19 Viewing Table Structures Using the show create Statement as Users with the select Permission

### Scenario

This function is applicable to Hive and Spark2x in MRS 3.x and later.

With this function enabled, if the select permission is granted to a user during Hive table creation, the user can run the **show create table** command to view the table structure.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Choose **HiveServer(Role) > Customization**, add a customized parameter to the **hive-site.xml** parameter file, set **Name** to **hive.allow.show.create.table.in.select.nogrant**, and set **Value** to **true**. Restart all Hive instances after the modification.

**Step 3** Determine whether to enable this function on the Spark/Spark2x client.

- If yes, download and install the Spark/Spark2x client again.
- If no, no further action is required.

----End

## 10.20 Writing a Directory into Hive with the Old Data Removed to the Recycle Bin

### Scenario

This function applies to Hive.

After this function is enabled, run the following command to write a directory into Hive: **insert overwrite directory "/path1" ...** After the operation is successfully performed, the old data is removed to the recycle bin, and the directory cannot be an existing database path in the Hive metastore.

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Choose **HiveServer(Role) > Customization**, add a customized parameter to the **hive-site.xml** parameter file, set **Name** to **hive.override.directory.move.trash**, and set **Value** to **true**. Restart all Hive instances after the modification.

----End

## 10.21 Inserting Data to a Directory That Does Not Exist

### Scenario

This function applies to Hive.

With this function enabled, run the **insert overwrite directory /path1/path2/path3...** command to write a subdirectory. The permission of the **/path1/path2** directory is 700, and the owner is the current user. If the **/path3** directory does not exist, it is automatically created and data is written successfully.

This function is supported when **hive.server2.enable.doAs** is set to **true** in earlier versions. This version supports the function when **hive.server2.enable.doAs** is set to **false**.

 **NOTE**

The parameter adjustment of this function is the same as that of the custom parameters added in [Writing a Directory into Hive with the Old Data Removed to the Recycle Bin](#).

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster >**



*Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**.

**Step 2** Choose **HiveServer(Role) > Customization**, add a customized parameter to the **hive-site.xml** parameter file, set **Name** to **hive.override.directory.move.trash**, and set **Value** to **true**. Restart all Hive instances after the modification.

----End

## 10.22 Creating Databases and Creating Tables in the Default Database Only as the Hive Administrator

### Scenario

This function is applicable to Hive and Spark2x for MRS 3.x or later, or Hive and Spark for versions earlier than MRS 3.x.

After this function is enabled, only the Hive administrator can create databases and tables in the default database. Other users can use the databases only after being authorized by the Hive administrator.

#### NOTE

- After this function is enabled, common users are not allowed to create a database or create a table in the default database. Based on the actual application scenario, determine whether to enable this function.
- Permissions of common users are restricted. In the scenario where common users have been used to perform operations, such as database creation, table script migration, and metadata recreation in an earlier version of database, the users can perform such operations on the database in the condition that this function is disabled temporarily after the database is migrated or after the cluster is upgraded.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components** > **Hive** > **Service Configuration**, and select **All** from the **Basic** drop-down list.

#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**.

**Step 2** Choose **HiveServer(Role) > Customization**, add a customized parameter to the **hive-site.xml** parameter file, set **Name** to **hive.allow.only.admin.create**, and set **Value** to **true**. Restart all Hive instances after the modification.

**Step 3** Determine whether to enable this function on the Spark/Spark2x client.

- If yes, go to [Step 4](#).
- If no, no further action is required.

**Step 4** Choose **SparkResource2x > Customization**, add a customized parameter to the **hive-site.xml** parameter file, set **Name** to **hive.allow.only.admin.create**, and set **Value** to **true**. Then, choose **JDBCServer2x > Customization** and repeat the preceding operations to add the customized parameter. Restart all Spark2x instances after the modification.

**Step 5** Download and install the Spark/Spark2x client again.

----End

## 10.23 Disabling of Specifying the location Keyword When Creating an Internal Hive Table

### Scenario

This function is applicable to Hive and Spark2x for MRS 3.x or later, or Hive and Spark for versions earlier than MRS 3.x.

After this function is enabled, the **location** keyword cannot be specified when a Hive internal table is created. Specifically, after a table is created, the table path following the location keyword is created in the default **/warehouse** directory and cannot be specified to another directory. If the location is specified when the internal table is created, the creation fails.

#### NOTE

After this function is enabled, the location keyword cannot be specified during the creation of a Hive internal table. The table creation statement is restricted. If a table that has been created in the database is not stored in the default directory **/warehouse**, the **location** keyword can still be specified when the database creation, table script migration, or metadata recreation operation is performed by disabling this function temporarily.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

#### NOTE

- If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)
- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Choose **HiveServer(Role) > Customization**, add a customized parameter to the **hive-site.xml** parameter file, set **Name** to **hive.internaltable.notallowlocation**, and set **Value** to **true**. Restart all Hive instances after the modification.

**Step 3** Determine whether to enable this function on the Spark/Spark2x client.

- If yes, download and install the Spark/Spark2x client again.
- If no, no further action is required.

----End

## 10.24 Enabling the Function of Creating a Foreign Table in a Directory That Can Only Be Read

### Scenario

This function is applicable to Hive and Spark2x for MRS 3.x or later, or Hive and Spark for versions earlier than MRS 3.x.

After this function is enabled, the user or user group that has the read and execute permissions on a directory can create foreign tables in the directory without checking whether the current user is the owner of the directory. In addition, the directory of a foreign table cannot be stored in the default directory **\warehouse**. In addition, do not change the permission of the directory during foreign table authorization.

#### NOTE

After this function is enabled, the function of the foreign table changes greatly. Based on the actual application scenario, determine whether to enable this function.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Choose **HiveServer(Role) > Customization**, add a customized parameter to the **hive-site.xml** parameter file, set **Name** to **hive.restrict.create.grant.external.table**, and set **Value** to **true**.

**Step 3** Choose **MetaStore(Role) > Customization**, add a customized parameter to the **hivemetastore-site.xml** parameter file, set **Name** to

**hive.restrict.create.grant.external.table**, and set **Value** to **true**. Restart all Hive instances after the modification.

**Step 4** Determine whether to enable this function on the Spark/Spark2x client.

- If yes, download and install the Spark/Spark2x client again.
- If no, no further action is required.

----End

## 10.25 Authorizing Over 32 Roles in Hive

### Scenario

This function applies to Hive.

The number of OS user groups is limited, and the number of roles that can be created in Hive cannot exceed 32. After this function is enabled, more than 32 roles can be created in Hive.

#### NOTE

- After this function is enabled and the table or database is authorized, roles that have the same permission on the table or database will be combined using vertical bars (|). When the ACL permission is queried, the combined result is displayed, which is different from that before the function is enabled. This operation is irreversible. Determine whether to make adjustment based on the actual application scenario.
- MRS 3.x and later versions support Ranger. If the current component uses Ranger for permission control, you need to configure related policies based on Ranger for permission management. For details, see [Adding a Ranger Access Permission Policy for Hive](#).
- After this function is enabled, a maximum of 512 roles (including **owner**) are supported by default. The number is controlled by the user-defined parameter **hive.supports.roles.max** of MetaStore. You can change the value based on the actual application scenario.

### Procedure

**Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

**Step 2** Choose **MetaStore(Role) > Customization**, add a custom parameter to the **hivemetastore-site.xml** parameter file, and set **Name** to **hive.supports.over.32.roles** and **Value** to **true**.

- Step 3** Choose **HiveServer(Role) > Customization**, add a custom parameter to the **hive-site.xml** parameter file, set **Name** to **hive.supports.over.32.roles**, and set **Value** to **true**. Restart all Hive instances after the modification.

----End

## 10.26 Restricting the Maximum Number of Maps for Hive Tasks

### Scenario

- This function applies to Hive.
- This function is used to limit the maximum number of maps for Hive tasks on the server to avoid performance deterioration caused by overload of the HiveServer service.

### Procedure

- Step 1** The Hive service configuration page is displayed.

- For versions earlier than MRS 3.x, click the cluster name. On the cluster details page that is displayed, choose **Components > Hive > Service Configuration**, and select **All** from the **Basic** drop-down list.

#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). And choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**.

- Step 2** Choose **MetaStore(Role) > Customization**, add a customized parameter to the **hivemetastore-site.xml** parameter file, set **Name** to **hive.mapreduce.per.task.max.splits**, and set the parameter to a large value. Restart all Hive instances after the modification.

----End

## 10.27 HiveServer Lease Isolation

### Scenario

- This function applies to Hive.
- This function can be enabled to specify specific users to access HiveServer services on specific nodes, achieving HiveServer resource isolation.

#### NOTE

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

## Procedure

This section describes how to set lease isolation for user **hiveuser** for existing HiveServer instances.

- Step 1** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **HiveServer**.
- Step 3** In the HiveServer list, select the HiveServer for which lease isolation is configured and choose **HiveServer** > **Instance Configurations** > **All Configurations**.
- Step 4** In the upper right corner of the **All Configurations** page, search for **hive.server2.zookeeper.namespace** and specify its value, for example, **hiveserver2\_zk**.
- Step 5** Click **Save**. In the dialog box that is displayed, click **OK**.
- Step 6** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive**, choose **More** > **Restart Service**, and enter the password to restart the service.
- Step 7** Run the **beeline -u** command to log in to the client and run the following command:

```
beeline -u
"jdbc:hive2://10.5.159.13:2181/?serviceDiscoveryMode=zooKeeper;zooKeeperName
space=hiveserver2_zk;sasl.qop=auth-conf;auth=KERBEROS;principal=hive/
hadoop.<System domain name>@<System domain name>"
```

In the command, **10.5.159.13** is replaced with the IP address of any ZooKeeper instance, which can be viewed through **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Instance**.

**hiveserver2\_zk** following **zooKeeperNamespace=** is set to the value of **hive.server2.zookeeper.namespace** in [Step 4](#).

As a result, only the HiveServer whose lease isolation is configured can be logged in.

### NOTE

- After this function is enabled, you must run the preceding command during login to access the HiveServer for which lease isolation is configured. If you run the **beeline** command to log in to the client, only the HiveServer that is not isolated by the lease is accessed.
- You can log in to FusionInsight Manager, choose **System** > **Permission** > **Domain and Mutual Trust**, and view the value of **Local Domain**, which is the current system domain name. **hive/hadoop.<system domain name>** is the username. All letters in the system domain name contained in the username are lowercase letters.

----End

## 10.28 Hive Supporting Transactions

### Scenario

Hive supports transactions at the table and partition levels. When the transaction mode is enabled, transaction tables can be incrementally updated, deleted, and

read, implementing atomicity, isolation, consistency, and durability of operations on transaction tables.

 **NOTE**

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

## Introduction to Transaction Features

A transaction is a group of unitized operations. These operations are either executed together or not executed together. A transaction is an inseparable unit of work. The four basic elements of a transaction are usually called ACID features, which are as follows:

- **Atomicity:** A transaction is an inseparable unit of work. All operations in a transaction occur or do not occur together.
- **Consistency:** The database integrity constraints are not damaged before and after a transaction starts.
- **Isolation:** When multiple transactions are concurrently accessed, the transactions are isolated from each other. A transaction does not affect the running of other transactions. The impacts between transactions are as follows: dirty read, non-repeatable read, phantom read, and lost update.
- **Durability:** After a transaction is complete, changes made by the transaction lock to the database are permanently stored in the database.

Characteristics of transaction execution:

- A statement can be written to multiple partitions or tables. If the operation fails, the user cannot see partial write or insert. Even if data is frequently changed, operations can still be quickly performed.
- Hive can automatically compress ACID transaction files without affecting concurrent queries. When querying many small partition files, automatic compression can improve query performance and metadata occupation.
- Read semantics include snapshot isolation. When the read operation starts, the Hive data warehouse is logically locked. The read operation is not affected by any changes that occur during the operation.

## Lock Mechanism

Transactions implement the ACID feature through the following two aspects:

- Write-ahead logging ensures atomicity and durability.
- Locking ensures isolation.

Operation	Type of Held Locks
Insert overwrite	If <b>hive.txn.xlock.iow</b> is set to <b>true</b> , the exclusive lock is held. If <b>hive.txn.xlock.iow</b> is set to <b>false</b> , the semi-shared lock is held.
Insert	Shared lock. When performing this operation, you can perform read and write operations on the current table or partition.

Operation	Type of Held Locks
Update/delete	Semi-shared lock. When this operation is performed, an operation of holding a shared lock can be performed, but an operation of holding an exclusive lock or a semi-shared lock cannot be performed.
Drop	Exclusive lock. You cannot perform any other operations on the current table or partition when performing this operation.

 NOTE

If a conflict caused by the lock mechanism exists in the write operation, the operation that preferentially holds the lock succeeds, and other operations fail.

## Procedure

### Starting a Transaction

**Step 1** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations** > **MetaStore(Role)** > **Transaction**.

**Step 2** Set **metastore.compactor.initiator.on** to **true**.

**Step 3** Set **metastore.compactor.worker.threads** to a positive integer.

 NOTE

**metastore.compactor.worker.threads**: Specifies the number of working threads for running the compression program on MetaStore. Set this parameter based on the actual requirements. If the value is too small, the transaction compression task is executed slowly. If the value is too large, the MetaStore execution performance deteriorates.

You are advised to set **metastore.compactor.worker.threads** to **15**. You can run the **show compactions** command to view the compactor tasks that are being executed and have been executed. If compactor tasks cannot be combined in a timely manner due to frequent services, you can choose a larger value for this parameter, but the value should not exceed **50**.

**Step 4** Log in to the Hive client and run the following command to enable the following parameters. For details, see [Using a Hive Client](#).

```
set hive.support.concurrency=true;
```

```
set hive.exec.dynamic.partition.mode=nonstrict;
```

```
set hive.txn.manager=org.apache.hadoop.hive.ql.lockmgr.DbTxnManager;
```

Create a transaction table.

**Step 5** Run the following command to create a transaction table:

```
CREATE TABLE [IF NOT EXISTS] [db_name.]table_name (col_name data_type
[COMMENT col_comment], ...) [ROW FORMAT row_format] STORED AS orc
TBLPROPERTIES ('transactional'='true'[, 'groupId'='group1' ...]);
```



For example:

```
CREATE TABLE acidTbl (a int, b int) STORED AS ORC TBLPROPERTIES ('transactional'='true');
```

 NOTE

- Currently, the transactions support only the ORC format.
- External tables are not supported.
- Sorted tables are not supported.
- To create a transaction table, you must add the table attribute **transactional='true'**.
- The transaction table can be read and written only in transaction mode.

**Use the transaction table.**

**Step 6** Run commands to use the transaction table. The following uses the **acidTbl** table as an example:

- Insert data into an existing transaction table:

```
INSERT INTO acidTbl VALUES(1,1);
```

- Update an existing transaction table:

```
UPDATE acidTbl SET b = 10 where a = 1;
```

The content of **acidTbl** is changed to:

```
+-----+-----+
| acidtbl.a | acidtbl.b |
+-----+-----+
| 1 | 10 |
+-----+-----+
1 row selected (0.775 seconds)
```

- Merge the old and new transaction tables:

The **acidTbl\_update** table contains the following data:

```
+-----+-----+
| acidtbl_update.a | acidtbl_update.b |
+-----+-----+
| 1 | 20 |
| 2 | 10 |
+-----+-----+
2 rows selected (0.537 seconds)
```

```
MERGE INTO acidTbl AS a
USING acidTbl_update AS b ON a.a = b.a
WHEN MATCHED THEN UPDATE SET b = b. b
WHEN NOT MATCHED THEN INSERT VALUES (b.a, b.b);
```

The content of **acidTbl** is changed to:

```
+-----+-----+
| acidtbl.a | acidtbl.b |
+-----+-----+
| 1 | 20 |
| 2 | 10 |
+-----+-----+
2 rows selected (0.666 seconds)
```

 NOTE

If "Error evaluating cardinality\_violation" is displayed when you run the **merge** command, check whether duplicate connection keys exist or run the **set hive.merge.cardinality.check=false** command to avoid this exception.

- Delete records from the transaction table.

```
DELETE FROM acidTbl where a = 2;
```

```

mysql> use mysql;
mysql> show tables;
+-----+
| acidtbl.a | acidtbl.b |
+-----+
| 1 | 20 |
+-----+
1 row selected (1.253 seconds)

```

### Checking the Transaction Execution Status

**Step 7** Run the following command to check the transaction execution status:

- Check the lock:  
**show locks;**
- Check the compression task:  
**show compactions;**
- Check the task execution status:  
**show transactions;**
- Interrupt a transaction:  
**abort transactions *TransactionId*;**  
*TransactionId* is the value in the **Transaction ID** column in the command output of [Check the task execution status](#).

----End

## Configuring the Compression Function

HDFS does not support in-place file changing. For the new content, HDFS does not provide read consistency either. To provide these features on HDFS, we follow the standard approach used in other data warehouse tools: table or partition data is stored in a set of base files, and new, updated, as well as deleted records are stored in incremental files. Each transaction creates a new set of incremental files to change the table or partition. When read, the base files and the incremental files are merged and the changes of the update or deletion are applied.

Writing a transaction table generates some small files in HDFS. Hive provides major and minor compression policies for combining these small files.

## Procedure of Automatic Compression

**Step 1** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations** > **MetaStore(Role)** > **Transaction**.

**Step 2** Set the following parameters as required:

**Table 10-11** Parameter description

Parameter	Description
hive.compactor.check.interval	Interval of executing compression threads. Unit: second. Default value: <b>300</b>

Parameter	Description
hive.compactor.cleaner.run.interval	Interval of executing cleaning threads. Unit: millisecond. Default value: <b>5,000</b> .
hive.compactor.delta.num.threshold	Threshold of the number of incremental files that trigger minor compression. Default value: <b>10</b>
hive.compactor.delta.pct.threshold	Ratio threshold of the total size of incremental files (delta) that trigger Major compression to the size of base files. The value <b>0.1</b> indicates that Major compression is triggered when the ratio of the total size of delta files to the size of base files is 10%. Default value: <b>0.1</b>
hive.compactor.max.num.delta	Maximum number of incremental files that the compressor will attempt to process in a single job. Default value: <b>500</b>
metastore.compactor.initiator.on	Indicates whether to run the startup program thread and cleanup program thread on the MetaStore instance. The value must be <b>true</b> . Default value: <b>false</b> .
metastore.compactor.worker.threads	Number of compression program work threads running on MetaStore. If this parameter is set to <b>0</b> , no compression is performed. To use a transaction, you must set this parameter to a positive number on one or more instances of the MetaStore service. Unit: second Default value: <b>0</b>

**Step 3** Log in to the Hive client and perform compression. For details, see [Using a Hive Client](#).

```
CREATE TABLE table_name (
 id int, name string
)
CLUSTERED BY (id) INTO 2 BUCKETS STORED AS ORC
TBLPROPERTIES ("transactional"="true",
 "compactor.mapreduce.map.memory.mb"="2048", -- Specify the properties of a compression
 "compactorthreshold.hive.compactor.delta.num.threshold"="4", -- If there are more than four incremental
 "compactorthreshold.hive.compactor.delta.pct.threshold"="0.5" -- If the ratio of the incremental file size to
 the basic file size is greater than 50%, deep compression is triggered.
);
```

or

```
ALTER TABLE table_name COMPACT 'minor' WITH OVERWRITE TBLPROPERTIES
("compactor.mapreduce.map.memory.mb"="3072"); -- Specify the properties of a compression map job.
ALTER TABLE table_name COMPACT 'major' WITH OVERWRITE TBLPROPERTIES
("tblprops.orc.compress.size"="8192"); -- Modify any other Hive table attributes.
```

 **NOTE**

After compression, small files are not deleted immediately. After the cleaner thread performs cleaning, the files are deleted in batches.

----End

## 10.29 Switching the Hive Execution Engine to Tez

### Scenario

Hive can use the Tez engine to process data computing tasks. Before executing a task, you can manually switch the execution engine to Tez.

### Prerequisites

The TimelineServer role of the Yarn service has been installed in the cluster and is running properly.

### Switching the Execution Engine on the Client to Tez

**Step 1** Install and log in to the Hive client. For details, see [Using a Hive Client](#).

**Step 2** Run the following commands to switch the engine and enable the `yarn.timeline-service.enabled` parameter:

```
set hive.execution.engine=tez;
set yarn.timeline-service.enabled=true;
```

#### NOTE

- After `yarn.timeline-service.enabled` is enabled, you can view the details about the tasks executed by the Tez engine on TezUI. After this function is enabled, task information will be reported to TimelineServer. If the TimelineServer instance is faulty, the task will fail.
- Tez uses the ApplicationMaster buffer pool. Therefore, `yarn.timeline-service.enabled` must be enabled before Tez tasks are submitted. Otherwise, this parameter cannot take effect and you need to log in to the client again to configure it.
- When the execution engine needs to be switched to another engine, you need to run the `set yarn.timeline-service.enabled=false` command on the client to disable the `yarn.timeline-service.enabled` parameter.
- To specify a Yarn running queue, run the `set tez.queue.name=default` command on the client.

**Step 3** Submit and execute the Tez tasks.

**Step 4** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services** > **Tez** > **TezUI** (*host name*) to view the task execution status on the TezUI page.

For versions earlier than MRS 3.x, log in to MRS Manager, choose **Services**, and click **Tez**. On the displayed page, click the link next to **Tez WebUI** to view the task execution status on the TezUI page.

----End

### Switching the Default Execution Engine of Hive to Tez

**Step 1** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* >

**Services > Hive > Configurations > All Configurations > HiveServer(Role)**, and search for **hive.execution.engine**.

For versions earlier than MRS 3.x, log in to MRS Manager, choose **Services**, and click **Hive**. On the displayed page, click the **Service Configuration** tab, select **All** from the **Type** drop-down list. On the navigation pane on the left, choose **HiveServer** and search for **hive.execution.engine**.

**Step 2** Set **hive.execution.engine** to **tez**.

**Step 3** Choose **Hive(Service) > Customization** and search for **yarn.site.customized.configs**.

**Step 4** Add custom parameter **yarn.timeline-service.enabled** to **yarn.site.customized.configs** and set it to **true**.

 **NOTE**

- After **yarn.timeline-service.enabled** is enabled, you can view the details about the tasks executed by the Tez engine on TezUI. After this function is enabled, task information will be reported to TimelineServer. If the TimelineServer instance is faulty, the task will fail.
- Tez uses the ApplicationMaster buffer pool. Therefore, **yarn.timeline-service.enabled** must be enabled before Tez tasks are submitted. Otherwise, this parameter cannot take effect and you need to log in to the client again to configure it.
- When the execution engine needs to be switched to another one, you need to set the value of parameter **yarn.timeline-service.enabled** to **false**.

**Step 5** Click **Save**. In the displayed confirmation dialog box, click **OK**.

For versions earlier than MRS 3.x, click **Save Configuration** and click **Yes** in the displayed dialog box.

**Step 6** Choose **Dashboard > More > Restart Service** to restart the Hive service. Enter the password to restart the service.

For versions earlier than MRS 3.x, Click the **Service Status** tab and choose **More > Restart Service** to restart the Hive service.

**Step 7** Install and log in to the Hive client. For details, see [Using a Hive Client](#).

**Step 8** Submit and execute the Tez tasks.

**Step 9** Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > Tez > TezUI (host name)**. On the displayed TezUI page, view the task execution status.

For versions earlier than MRS 3.x, log in to MRS Manager, choose **Services**, and click **Tez**. On the displayed page, click the link next to **Tez WebUI** to view the task execution status on the TezUI page.

----End

## 10.30 Hive Materialized View

### Introduction

A Hive materialized view is a special table obtained based on the query results of Hive internal tables. A materialized view can be considered as an intermediate

table that stores actual data and occupies physical space. The tables on which a materialized view depends are called the base tables of the materialized view.

Materialized views are used to pre-compute and save the results of time-consuming operations such as table joining or aggregation. When executing a query, you can rewrite the query statement based on the base tables to the query statement based on materialized views. In this way, you do not need to perform time-consuming operations such as join and group by, thereby quickly obtaining the query result.

#### NOTE

- A materialized view is a special table that stores actual data and occupies physical space.
- Before deleting a base table, you must delete the materialized view created based on the base table.
- The materialized view creation statement is atomic, which means that other users cannot see the materialized view until all query results are populated.
- A materialized view cannot be created based on the query results of another materialized view.
- A materialized view cannot be created based on the results of a tableless query.
- You cannot insert, update, delete, load, or merge materialized views.
- You can perform complex query operations on materialized views, because they are special tables in nature.
- When the data of a base table is updated, you need to manually update the materialized view. Otherwise, the materialized view will retain the old data. That is, the materialized view expires.
- You can use the describe syntax to check whether the materialized view created based on ACID tables has expired.
- The describe statement cannot be used to check whether a materialized view created based on non-ACID tables has expired.
- A materialized view can store only ORC files. You can use TBLPROPERTIES ('transactional='true') to create a transactional Hive internal table.

## Creating a Materialized View

### Syntax

```
CREATE MATERIALIZED VIEW [IF NOT EXISTS] [db_name.]materialized_view_name
[COMMENT materialized_view_comment]
DISABLE REWRITE
[ROW FORMAT row_format]
[STORED AS file_format]
| STORED BY 'storage.handler.class.name' [WITH SERDEPROPERTIES (...)]
]
[LOCATION hdfs_path]
[TBLPROPERTIES (property_name=property_value, ...)]
AS
<query>;
```

 NOTE

- Currently, the following materialized view file formats are supported: PARQUET, TextFile, SequenceFile, RCfile, and ORC. If **STORED AS** is not specified in the creation statement, the default file format is ORC.
- Names of materialized views must be unique in the same database. Otherwise, you cannot create a new materialized view, and data files of the original materialized view will be overwritten by the data files queried based on the base table in the new one. As a result, data may be tampered with. (After being tampered with, the materialized view can be restored by re-creating the materialized view.)

**Cases**

**Step 1** Log in to the Hive client and run the following command to enable the following parameters. For details, see [Using a Hive Client](#).

```
set hive.support.concurrency=true;
```

```
set hive.exec.dynamic.partition.mode=nonstrict;
```

```
set hive.txn.manager=org.apache.hadoop.hive.ql.lockmgr.DbTxnManager;
```

**Step 2** Create a base table and insert data.

```
create table tb_emp(
empno int,ename string,job string,mgr int,hiredate TIMESTAMP,sal float,comm float,deptno int
)stored as orc
tblproperties('transactional'='true');

insert into tb_emp values(7369, 'SMITH', 'CLERK',7902, '1980-12-17 08:30:09',800.00,NULL,20),
(7499, 'ALLEN', 'SALESMAN',7698, '1981-02-20 17:12:00',1600.00,300.00,30),
(7521, 'WARD', 'SALESMAN',7698, '1981-02-22 09:05:34',1250.00,500.00,30),
(7566, 'JONES', 'MANAGER', 7839, '1981-04-02 10:14:13',2975.00,NULL,20),
(7654, 'MARTIN', 'SALESMAN',7698, '1981-09-28 08:36:17',1250.00,1400.00,30),
(7698, 'BLAKE', 'MANAGER',7839, '1981-05-01 11:12:55',2850.00,NULL,30),
(7782, 'CLARK', 'MANAGER',7839, '1981-06-09 15:45:28',2450.00,NULL,10),
(7788, 'SCOTT', 'ANALYST',7566, '1987-04-19 14:05:34',3000.00,NULL,20),
(7839, 'KING', 'PRESIDENT',NULL, '1981-11-17 10:18:25',5000.00,NULL,10),
(7844, 'TURNER', 'SALESMAN',7698, '1981-09-08 09:05:34',1500.00,0.00,30),
(7876, 'ADAMS', 'CLERK',7788, '1987-05-23 15:07:44',1100.00,NULL,20),
(7900, 'JAMES', 'CLERK',7698, '1981-12-03 16:23:56',950.00,NULL,30),
(7902, 'FORD', 'ANALYST',7566, '1981-12-03 08:48:17',3000.00,NULL,20),
(7934, 'MILLER', 'CLERK',7782, '1982-01-23 11:45:29',1300.00,NULL,10);
```

**Step 3** Create a materialized view based on the results of the **tb\_emp** query.

```
create materialized view group_mv disable rewrite
row format serde 'org.apache.hadoop.hive.serde2.JsonSerDe'
stored as textfile
tblproperties('mv_content'='Total compensation of each department')
as select deptno,sum(sal) sum_sal from tb_emp group by deptno;
```

----End

## Applying a Materialized View

Rewrite the query statement based on base tables to the query statement based on materialized views to improve the query efficiency.

**Cases**

Execute the following query statement:

```
select deptno,sum(sal) from tb_emp group by deptno having sum(sal)>10000;
```

Based on the created materialized view, rewrite the query statement:

```
select deptno, sum_sal from group_mv where sum_sal>10000;
```

## Checking a Materialized View

### Syntax

```
SHOW MATERIALIZED VIEWS [IN database_name]
['identifier_with_wildcards'];
```

```
DESCRIBE [EXTENDED | FORMATTED] [db_name.]materialized_view_name;
```

### Cases

```
show materialized views;
```

```
describe formatted group_mv;
```

## Deleting a Materialized View

### Syntax

```
DROP MATERIALIZED VIEW [db_name.]materialized_view_name;
```

### Cases

```
drop materialized view group_mv;
```

## Rebuilding a Materialized View

When a materialized view is created, the base table data is filled in the materialized view. However, the data that is added, deleted, or modified in the base table is not automatically synchronized to the materialized view. Therefore, you need to manually rebuild the view after updating the data.

### Syntax

```
ALTER MATERIALIZED VIEW [db_name.]materialized_view_name REBUILD;
```

### Cases

```
alter materialized view group_mv rebuild;
```

### NOTE

When the base table data is updated but the materialized view data is not updated, the materialized view is in the expired state by default.

The describe statement can be used to check whether a materialized view created based on transaction tables has expired. If the value of **Outdated for Rewriting** is **Yes**, the license has expired. If the value of **Outdated for Rewriting** is **No**, the license has not expired.

## 10.31 Hive Log Overview

### Log Description

**Log path:** The default save path of Hive logs is `/var/log/Bigdata/hive/role name`, the default save path of Hive1 logs is `/var/log/Bigdata/hive1/role name`, and the others follow the same rule.



- HiveServer: **/var/log/Bigdata/hive/hiveserver** (run log) and **var/log/Bigdata/audit/hive/hiveserver** (audit log)
- MetaStore: **/var/log/Bigdata/hive/metastore** (run log) and **/var/log/Bigdata/audit/hive/metastore** (audit log)
- WebHCat: **/var/log/Bigdata/hive/webhcat** (run log) and **/var/log/Bigdata/audit/hive/webhcat** (audit log)

**Log archive rule:** The automatic compression and archiving function of Hive is enabled. By default, when the size of a log file exceeds 20 MB (which is adjustable), the log file is automatically compressed. The naming rule of a compressed log file is as follows: *<Original log name>-<yyyy-mm-dd\_hh-mm-ss>.[ID].log.zip* A maximum of 20 latest compressed files are reserved. The number of compressed files and compression threshold can be configured.

**Table 10-12** Hive log list

Log Type	Log File Name	Description
Run log	/hiveserver/hiveserver.out	Log file that records HiveServer running environment information.
	/hiveserver/hive.log	Run log file of the HiveServer process.
	/hiveserver/hive-omm- <Date> <PID>- gc.log.<No.>	GC log file of the HiveServer process.
	/hiveserver/ prestartDetail.log	Work log file before the HiveServer startup.
	/hiveserver/check- serviceDetail.log	Log file that records whether the Hive service starts successfully
	/hiveserver/ cleanupDetail.log	Cleanup log file about the HiveServer uninstallation
	/hiveserver/startDetail.log	Startup log file of the HiveServer process.
	/hiveserver/stopDetail.log	Shutdown log file of the HiveServer process.
	/hiveserver/localtasklog/ omm_<Date>_<Task ID>.log	Run log file of the local Hive task.
	/hiveserver/localtasklog/ omm_<Date>_<Task ID>- gc.log.<No.>	GC log file of the local Hive task.
	/metastore/metastore.log	Run log file of the MetaStore process.

Log Type	Log File Name	Description
	/metastore/hive-omm- <Date>-<PID>- gc.log.<No.>	GC log file of the MetaStore process.
	/metastore/ postinstallDetail.log	Work log file after the MetaStore installation.
	/metastore/ prestartDetail.log	Work log file before the MetaStore startup
	/metastore/ cleanupDetail.log	Cleanup log file of the MetaStore uninstallation
	/metastore/startDetail.log	Startup log file of the MetaStore process.
	/metastore/stopDetail.log	Shutdown log file of the MetaStore process.
	/metastore/metastore.out	Log file that records MetaStore running environment information.
	/webhcat/webhcat- console.out	Log file that records the normal start and stop of the WebHCat process.
	/webhcat/webhcat- console-error.out	Log file that records the start and stop exceptions of the WebHCat process.
	/webhcat/ prestartDetail.log	Work log file before the WebHCat startup.
	/webhcat/ cleanupDetail.log	Cleanup logs generated during WebHCat uninstallation or before WebHCat installation
	/webhcat/hive-omm- <Date>-<PID>- gc.log.<No.>	GC log file of the WebHCat process.
	/webhcat/webhcat.log	Run log file of the WebHCat process
Audit log	hive-audit.log hive-rangeraudit.log	HiveServer audit log file
	metastore-audit.log	MetaStore audit log file.
	webhcat-audit.log	WebHCat audit log file.
	jetty-<Date>.request.log	Request logs of the jetty service.

## Log Levels

**Table 10-13** describes the log levels supported by Hive.

Levels of run logs are ERROR, WARN, INFO, and DEBUG from the highest to the lowest priority. Run logs of equal or higher levels are recorded. The higher the specified log level, the fewer the logs recorded.

**Table 10-13** Log levels

Level	Description
ERROR	Logs of this level record error information about system running.
WARN	Logs of this level record exception information about the current event processing.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of the Yarn service by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level and save the configuration.

 **NOTE**

The Hive log level takes effect immediately after being configured. You do not need to restart the service.

----End

## Log Formats

The following table lists the Hive log formats:

**Table 10-14** Log formats

Log Type	Format	Example
Run log	<yyyy-MM-dd HH:mm:ss,SSS>  <LogLevel> <Thread that generates the log>  <Message in the log>  <Location of the log event>	2014-11-05 09:45:01,242   INFO   main   Starting hive metastore on port 21088   org.apache.hadoop.hive.metas- tore.HiveMetaStore.main(Hive MetaStore.java:5198)
Audit log	<yyyy-MM-dd HH:mm:ss,SSS>  <LogLevel> <Thread that generates the log> <User Name><User IP><Time><Operation><Re- source><Result><Detail > < Location of the log event >	2018-12-24 12:16:25,319   INFO   HiveServer2-Handler- Pool: Thread-185   UserName=hive UserIP=10.153.2.204 Time=2018/12/24 12:16:25 Operation=CloseSession Result=SUCCESS Detail=   org.apache.hive.service.cli.thrif- t.ThriftCLIService.logAuditEven- t(ThriftCLIService.java:434)

## 10.32 Hive Performance Tuning

### 10.32.1 Creating Table Partitions

#### Scenario

During the Select query, Hive generally scans the entire table, which is time-consuming. To improve query efficiency, create table partitions based on service requirements and query dimensions.

#### Procedure

**Step 1** For versions earlier than MRS 3.x:

Log in to the MRS console. In the left navigation pane, choose **Clusters > Active Clusters**, and click a cluster name. Choose **Nodes > Node**. The ECS page is displayed. Click **Remote Login** to log in to the Hive node.

For MRS 3.x or later:

Log in to the node where the Hive client has been installed as user **root**.

**Step 2** Run the following command to go to the client installation directory, for example, **/opt/client**:

```
cd /opt/client
```

**Step 3** Run the **source bigdata\_env** command to configure environment variables for the client.

**Step 4** Run the following command on the client for login:

```
kinit Username
```

**Step 5** Run the following command to log in to the client tool:

```
beeline
```

**Step 6** Select the static or dynamic partition.

- Static partition:

Manually enter a partition name, and use the keyword **PARTITIONED BY** to specify partition column name and data type when creating a table. During application development, use the **ALTER TABLE ADD PARTITION** statement to add a partition and use the **LOAD DATA INTO PARTITION** statement to load data to the partition, which supports only static partitions.

- Dynamic partition: Use a query command to insert results to a partition of a table. The partition can be a dynamic partition.

The dynamic partition can be enabled on the client tool by running the following command:

```
set hive.exec.dynamic.partition=true;
```

The default mode of the dynamic partition is strict. That is, at least a column must be specified as a static partition, under which dynamic sub-partitions can be created. You can run the following command to enable a completely dynamic partition:

```
set hive.exec.dynamic.partition.mode=nonstrict;
```

 **NOTE**

- The dynamic partition may cause a DML statement to create a large number of partitions and new mapping folders, which deteriorates system performance.
- If there are a large number of files, it takes a long time to run a SQL statement. You can run the **set mapreduce.input.fileinputformat.list-status.num-threads = 100;** statement before running a SQL statement to shorten the time. The parameter **mapreduce.input.fileinputformat.list-status.num-threads** can be set only after being added to the Hive whitelist.

----End

## 10.32.2 Optimizing Join

### Scenario

When the Join statement is used, the command execution speed and query speed may be slow in case of large data volume. To resolve this problem, you can optimize Join.

Join optimization can be classified into the following modes:

- Map Join
- Sort Merge Bucket Map Join
- Optimizing Join Sequences

## Map Join

Hive Map Join applies to small tables (the table size is less than 25 MB) that can be stored in the memory. The table size can be defined using **hive.mapjoin.smalltable.filesize**, and the default table size is 25 MB.

Map Join has two methods:

- Use `/*+ MAPJOIN(join_table) */`.
- Set the following parameter before running the statement. The default value is true in the current version.

**set hive.auto.convert.join=true;**

There is no Reduce task when Map Join is used. Instead, a MapReduce Local Task is created before the Map job. The task uses TableScan to read small table data to the local computer, saves and writes the data in HashTable mode to a hard disk on the local computer, upload the data to DFS, and saves the data in distributed cache. The small table data that the map task reads from the local disk or distributed cache is the output together with the large table join result.

When using Map Join, make sure that the size of small tables cannot be too large. If small tables use up memory, the system performance will deteriorate and even memory leakage occurs.

## Sort Merge Bucket Map Join

The following conditions must be met before using Sort Merge Bucket Map Join:

- The two Join tables are large and cannot be stored in the memory.
- The two tables are bucketed (clustered by (column)) and sorted (sorted by(column)) according to the join key, and the buckets counts of the two tables are in integral multiple relationship.

Set the following parameters to enable Sort Merge Bucket Map Join:

**set hive.optimize.bucketmapjoin=true;**

**set hive.optimize.bucketmapjoin.sortedmerge=true;**

This type of Map Join does not have Reduce tasks too. A MapReduce Local Task is started before the Map job to read small table data by bucket to the local computer. The local computer saves the HashTable backup of multiple buckets and writes the backup into HDFS. The backup is also saved in the distributed cache. The small table data that the map task reads from the local disk or distributed cache by bucket is the output after mapping with the large table.

## Optimizing Join Sequences

If the Join operation is to be performed on three or more tables and different Join sequences are used, the execution time will be greatly different. Using an appropriate Join sequence can shorten the time for task execution.

Rules of a Join sequence:

- A table with small data volume or a combination with fewer results generated after a Join operation is executed first.

- A table with large data volume or a combination with more results generated after a Join operation is executed later.

For example, the **customer** table has the largest data volume, and fewer results will be generated if a Join operation is performed on the **orders** and **lineitem** tables first.

The original Join statement is as follows.

```
select
 l_orderkey,
 sum(l_extendedprice * (1 - l_discount)) as revenue,
 o_orderdate,
 o_shippriority
from
 customer,
 orders,
 lineitem
where
 c_mktsegment = 'BUILDING'
 and c_custkey = o_custkey
 and l_orderkey = o_orderkey
 and o_orderdate < '1995-03-22'
 and l_shipdate > '1995-03-22'
limit 10;
```

After the sequence is optimized, the Join statements are as follows:

```
select
 l_orderkey,
 sum(l_extendedprice * (1 - l_discount)) as revenue,
 o_orderdate,
 o_shippriority
from
 orders,
 lineitem,
 customer
where
 c_mktsegment = 'BUILDING'
 and c_custkey = o_custkey
 and l_orderkey = o_orderkey
 and o_orderdate < '1995-03-22'
 and l_shipdate > '1995-03-22'
limit 10;
```

## Precautions

### Join Data Skew Problem

Data skew refers to the symptom that the task progress is 99% for a long time.

Data skew often exists because the data volume of a few Reduce tasks is much larger than that of others. Most Reduce tasks are complete while a few Reduce tasks are not complete.

To resolve the data skew problem, set **hive.optimize.skewjoin=true** and adjust the value of **hive.skewjoin.key**. **hive.skewjoin.key** specifies the maximum number of keys received by a Reduce task. If the number reaches the maximum, the keys are atomically distributed to other Reduce tasks.

## 10.32.3 Optimizing Group By

### Scenario

Optimize the Group by statement to accelerate the command execution and query speed.

During the Group by operation, Map performs grouping and distributes the groups to Reduce; Reduce then performs grouping again. Group by optimization can be performed by enabling Map aggregation to reduce Map output data volume.

### Procedure

On a Hive client, set the following parameter:

```
set hive.map.aggr=true
```

### Precautions

#### Group By Data Skew

Group by have data skew problems. When `hive.groupby.skewindata` is set to true, the created query plan has two MapReduce jobs. The Map output result of the first job is randomly distributed to Reduce tasks, and each Reduce task performs aggregation operations and generates output result. Such processing may distribute the same Group By Key to different Reduce tasks for load balancing purpose. According to the preprocessing result, the second Job distributes Group By Key to Reduce to complete the final aggregation operation.

#### Count Distinct Aggregation Problem

When the aggregation function `count distinct` is used in deduplication counting, serious Reduce data skew occurs if the processed value is empty. The empty value can be processed independently. If `count distinct` is used, exclude the empty value using the `where` statement and increase the last `count distinct` result by 1. If there are other computing operations, process the empty value independently and then combine the value with other computing results.

## 10.32.4 Optimizing Data Storage

### Scenario

**ORC** is an efficient column storage format and has higher compression ratio and reading efficiency than other file formats.

You are advised to use **ORC** as the default Hive table storage format.

### Prerequisites

You have logged in to the Hive client. For details, see [Using a Hive Client](#).

### Procedure

- Recommended: **SNAPPY** compression, which applies to scenarios with even compression ratio and reading efficiency requirements.



**Create table *xx* (*col\_name data\_type*) stored as orc tblproperties ("orc.compress"="SNAPPY");**

- Available: **ZLIB** compression, which applies to scenarios with high compression ratio requirements.

**Create table *xx* (*col\_name data\_type*) stored as orc tblproperties ("orc.compress"="ZLIB");**

 **NOTE**

*xx* indicates the specific Hive table name.

## 10.32.5 Optimizing SQL Statements

### Scenario

When SQL statements are executed on Hive, if the **(a&b) or (a&c)** logic exists in the statements, you are advised to change the logic to **a & (b or c)**.

### Example

If condition a is **p\_partkey = l\_partkey**, the statements before optimization are as follows:

```
select
 sum(l_extendedprice* (1 - l_discount)) as revenue
from
 lineitem,
 part
where
 (
 p_partkey = l_partkey
 and p_brand = 'Brand#32'
 and p_container in ('SM CASE', 'SM BOX', 'SM PACK', 'SM PKG')
 and l_quantity >= 7 and l_quantity <= 7 + 10
 and p_size between 1 and 5
 and l_shipmode in ('AIR', 'AIR REG')
 and l_shipinstruct = 'DELIVER IN PERSON'
)
 or
 (
 p_partkey = l_partkey
 and p_brand = 'Brand#35'
 and p_container in ('MED BAG', 'MED BOX', 'MED PKG', 'MED PACK')
 and l_quantity >= 15 and l_quantity <= 15 + 10
 and p_size between 1 and 10
 and l_shipmode in ('AIR', 'AIR REG')
 and l_shipinstruct = 'DELIVER IN PERSON'
)
 or
 (
 p_partkey = l_partkey
 and p_brand = 'Brand#24'
 and p_container in ('LG CASE', 'LG BOX', 'LG PACK', 'LG PKG')
 and l_quantity >= 26 and l_quantity <= 26 + 10
 and p_size between 1 and 15
 and l_shipmode in ('AIR', 'AIR REG')
 and l_shipinstruct = 'DELIVER IN PERSON'
)
)
```

The statements after optimization are as follows:

```
select
 sum(l_extendedprice* (1 - l_discount)) as revenue
from
 lineitem,
```

```
part
where p_partkey = l_partkey and
 (
 p_brand = 'Brand#32'
 and p_container in ('SM CASE', 'SM BOX', 'SM PACK', 'SM PKG')
 and l_quantity >= 7 and l_quantity <= 7 + 10
 and p_size between 1 and 5
 and l_shipmode in ('AIR', 'AIR REG')
 and l_shipinstruct = 'DELIVER IN PERSON'
)
or
 (
 p_brand = 'Brand#35'
 and p_container in ('MED BAG', 'MED BOX', 'MED PKG', 'MED PACK')
 and l_quantity >= 15 and l_quantity <= 15 + 10
 and p_size between 1 and 10
 and l_shipmode in ('AIR', 'AIR REG')
 and l_shipinstruct = 'DELIVER IN PERSON'
)
or
 (
 p_brand = 'Brand#24'
 and p_container in ('LG CASE', 'LG BOX', 'LG PACK', 'LG PKG')
 and l_quantity >= 26 and l_quantity <= 26 + 10
 and p_size between 1 and 15
 and l_shipmode in ('AIR', 'AIR REG')
 and l_shipinstruct = 'DELIVER IN PERSON'
)
))
```

## 10.32.6 Optimizing the Query Function Using Hive CBO

### Scenario

When joining multiple tables in Hive, Hive supports Cost-Based Optimization (CBO). The system automatically selects the optimal plan based on the table statistics, such as the data volume and number of files, to improve the efficiency of joining multiple tables. Hive needs to collect table statistics before CBO optimization.

#### NOTE

- The CBO optimizes the joining sequence based on statistics and search criteria. However, the joining sequence may fail to be optimized in some special scenarios, such as data skew occurs and query condition values are not in the table.
- When column statistics collection is enabled, Reduce operations must be performed for aggregation. For insert tasks without the Reduce phase, Reduce operations will be performed to collect statistics.
- This section applies to MRS 3.x or later clusters.

### Prerequisites

You have logged in to the Hive client. For details, see [Using a Hive Client](#).

### Procedure

- Step 1** On the Manager UI, search for the **hive.cbo.enable** parameter in the service configuration of the Hive component, and select **true** to enable the function permanently.
- Step 2** Collect statistics about the existing data in Hive tables manually.

Run the following command to manually collect statistics: Statistics about only one table can be collected. If statistics about multiple tables need to be collected, the command needs to be executed repeatedly.

```
ANALYZE TABLE [db_name.]tablename [PARTITION(partcol1[=val1],
partcol2[=val2], ...)]
```

```
COMPUTE STATISTICS
```

```
[FOR COLUMNS]
```

```
[NOSCAN];
```

 **NOTE**

- When **FOR COLUMNS** is specified, column-level statistics are collected.
- When **NOSCAN** is specified, statistics about the file size and number of files will be collected, but specific files will not be scanned.

For example:

```
analyze table table_name compute statistics;
```

```
analyze table table_name compute statistics for columns;
```

**Step 3** Configure the automatic statistics collection function of Hive. After the function is enabled, new statistics will be collected only when you insert data by running the **insert overwrite/into** command.

- Run the following commands on the Hive client to enable the statistics collection function temporarily:

```
set hive.stats.autogather = true; enables the automatic collection of table/
partition-level statistics.
```

```
set hive.stats.column.autogather = true; enables the automatic collection of
column-level statistics.
```

 **NOTE**

- The column-level statistics collection does not support complex data types, such as Map and Struct.
- The automatic table-level statistics collection does not support Hive on HBase tables.
- On the Manager UI, search for the **hive.stats.autogather** and **hive.stats.column.autogather** parameters in the service configuration of Hive, and select **true** to enable the collection function permanently.

**Step 4** Run the following command to view statistics:

```
DESCRIBE FORMATTED table_name[.column_name] PARTITION
partition_spec;
```

For example:

```
desc formatted table_name;
```

```
desc formatted table_name id;
```

```
desc formatted table_name partition(time='2016-05-27');
```

 NOTE

Partition tables only support partition-level statistics collection, so you must specify partitions to query statistics for partition tables.

----End

## 10.33 Common Issues About Hive

### 10.33.1 How Do I Delete UDFs on Multiple HiveServers at the Same Time?

#### Question

How can I delete permanent user-defined functions (UDFs) on multiple HiveServers at the same time?

#### Answer

Multiple HiveServers share one MetaStore database. Therefore, there is a delay in the data synchronization between the MetaStore database and the HiveServer memory. If a permanent UDF is deleted from one HiveServer, the operation result cannot be synchronized to the other HiveServers promptly.

In this case, you need to log in to the Hive client to connect to each HiveServer and delete permanent UDFs on the HiveServers one by one. The operations are as follows:

**Step 1** Log in to the node where the Hive client is installed as the Hive client installation user.

**Step 2** Run the following command to go to the client installation directory:

```
cd Client installation directory
```

For example, if the client installation directory is **/opt/client**, run the following command:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Run the following command to authenticate the user:

```
kinit Hive service user
```

 NOTE

The login user must have the Hive admin rights.

**Step 5** Run the following command to connect to the specified HiveServer:

```
beeline -u "jdbc:hive2://10.39.151.74:21066/default;sasl.qop=auth-
conf;auth=KERBEROS;principal=hive/hadoop.<system domain name>@<system
domain name>"
```

 NOTE

- *10.39.151.74* is the IP address of the node where the HiveServer is located.
- *21066* is the port number of the HiveServer. The HiveServer port number ranges from 21066 to 21070 by default. Use the actual port number.
- *hive* is the username. For example, if the Hive1 instance is used, the username is **hive1**.
- You can log in to FusionInsight Manager, choose **System > Permission > Domain and Mutual Trust**, and view the value of **Local Domain**, which is the current system domain name.
- **hive/hadoop.<system domain name>** is the username. All letters in the system domain name contained in the username are lowercase letters.

**Step 6** Run the following command to enable the Hive admin rights:

```
set role admin;
```

**Step 7** Run the following command to delete the permanent UDF:

```
drop function function_name;
```

 NOTE

- *function\_name* indicates the name of the permanent function.
- If the permanent UDF is created in Spark, the permanent UDF needs to be deleted from Spark and then from HiveServer by running the preceding command.

**Step 8** Check whether the permanent UDFs are deleted from all HiveServers.

- If yes, no further action is required.
- If no, go to [Step 5](#).

----End

## 10.33.2 Why Cannot the DROP operation Be Performed on a Backed-up Hive Table?

### Question

Why cannot the **DROP** operation be performed for a backed up Hive table?

### Answer

Snapshots have been created for an HDFS directory mapping to the backed up Hive table, so the HDFS directory cannot be deleted. As a result, the Hive table cannot be deleted.

When a Hive table is being backed up, snapshots are created for the HDFS directory mapping to the table. The snapshot mechanism of HDFS has the following limitation: If snapshots have been created for an HDFS directory, the directory cannot be deleted or renamed unless the snapshots are deleted. When the **DROP** operation is performed for a Hive table (except the EXTERNAL table), the system attempts to delete the HDFS directory mapping to the table. If the

directory fails to be deleted, the system displays a message indicating that the table fails to be deleted.

If you need to delete this table, manually delete all backup tasks related to this table.

### 10.33.3 How to Perform Operations on Local Files with Hive User-Defined Functions

#### Question

How to perform operations on local files (such as reading the content of a file) with Hive user-defined functions?

#### Answer

By default, you can perform operations on local files with their relative paths in UDF. The following are sample codes:

```
public String evaluate(String text) {
 // some logic
 File file = new File("foo.txt");
 // some logic
 // do return here
}
```

In Hive, upload the file **foo.txt** used in UDF to HDFS, such as **hdfs://hacluster/tmp/foo.txt**. You can perform operations on the **foo.txt** file by creating UDF with the following sentences:

```
create function testFunc as 'some.class' using jar 'hdfs://hacluster/somejar.jar', file 'hdfs://hacluster/tmp/foo.txt';
```

In abnormal cases, if the value of **hive.fetch.task.conversion** is **more**, you can perform operations on local files in UDF by using absolute path instead of relative path. In addition, you must ensure that the file exists on all HiveServer nodes and NodeManager nodes and **omm** user have corresponding operation rights.

### 10.33.4 How Do I Forcibly Stop MapReduce Jobs Executed by Hive?

#### Question

How do I stop a MapReduce task manually if the task is suspended for a long time?

#### Answer

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

**Step 2** Choose **Cluster > Name of the desired cluster > Services > Yarn**.

**Step 3** On the left pane, click **ResourceManager(Host name, Active)**, and log in to Yarn.

- Step 4** Click the button corresponding to the task ID. On the task page that is displayed, click **Kill Application** in the upper left corner and click **OK** in the displayed dialog box to stop the task.

----End

## 10.33.5 Table Creation Fails Because Hive Complex Fields' Names Contain Special Characters

### Question

Table creation fails because Hive complex fields' names contain special characters.

### Answer

Hive does not support complex fields' names that contain special characters. Special characters refer to characters other than uppercase and lowercase letters, digits, Chinese characters, and Portuguese characters.

## 10.33.6 How Do I Monitor the Hive Table Size?

### Question

How do I monitor the Hive table size?

### Answer

The HDFS refined monitoring function allows you to monitor the size of a specified table directory.

### Prerequisites

- The Hive and HDFS components are running properly.
- The HDFS refined monitoring function is normal.

### Procedure

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Resource**.
- Step 3** Click the first icon in the upper left corner of **Resource Usage (by Directory)**, as shown in the following figure.

Resource Usage (by Directory) 

- Step 4** In the displayed sub page for configuring space monitoring, click **Add**.
- Step 5** In the displayed **Add a Monitoring Directory** dialog box, set **Name** to the name or the user-defined alias of the table to be monitored and **Path** to the path of the monitored table. Click **OK**. In the monitoring result, the horizontal coordinate

indicates the time, and the vertical coordinate indicates the size of the monitored directory.

----End

## 10.33.7 How Do I Prevent Key Directories from Data Loss Caused by Misoperations of the insert overwrite Statement?

### Question

How do I prevent key directories from data loss caused by misoperations of the **insert overwrite** statement?

### Answer

During monitoring of key Hive databases, tables, or directories, to prevent data loss caused by misoperations of the **insert overwrite** statement, configure **hive.local.dir.confblacklist** in Hive to protect directories.

This configuration item has been configured for directories such as **/opt/** and **/user/hive/warehouse** by default.

### Prerequisites

The Hive and HDFS components are running properly.

### Procedure

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster > Name of the desired cluster > Services > Hive > Configurations > All Configurations**, and search for the **hive.local.dir.confblacklist** configuration item.
- Step 3** Add paths of databases, tables, or directories to be protected in the parameter value.
- Step 4** Click **Save** to save the settings.

----End

## 10.33.8 Why Is Hive on Spark Task Freezing When HBase Is Not Installed?

### Scenario

This function applies to Hive.

Perform the following operations to configure parameters. When Hive on Spark tasks are executed in the environment where the HBase is not installed, freezing of tasks can be prevented.



 NOTE

The Spark kernel version of Hive on Spark tasks has been upgraded to Spark2x. Hive on Spark tasks can be executed if Spark2x is not installed. If HBase is not installed, when Spark tasks are executed, the system attempts to connect to the ZooKeeper to access HBase until timeout occurs by default. As a result, task freezing occurs.

If HBase is not installed, perform the following operations to execute Hive on Spark tasks. If HBase is upgraded from an earlier version, you do not need to configure parameters after the upgrade.

## Procedure

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**.
- Step 3** Choose **HiveServer(Role)** > **Customization**. Add a customized parameter to the **spark-defaults.conf** parameter file. Set **Name** to **spark.security.credentials.hbase.enabled**, and set **Value** to **false**.
- Step 4** Click **Save**. In the dialog box that is displayed, click **OK**.
- Step 5** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance**, select all Hive instances, choose **More** > **Restart Instance**, enter the password, and click **OK**.

----End

## 10.33.9 Error Reported When the WHERE Condition Is Used to Query Tables with Excessive Partitions in FusionInsight Hive

### Question

When a table with more than 32,000 partitions is created in Hive, an exception occurs during the query with the WHERE partition. In addition, the exception information printed in **metastore.log** contains the following information:

```
Caused by: java.io.IOException: Tried to send an out-of-range integer as a 2-byte value: 32970
 at org.postgresql.core.PGStream.SendInteger2(PGStream.java:199)
 at org.postgresql.core.v3.QueryExecutorImpl.sendParse(QueryExecutorImpl.java:1330)
 at org.postgresql.core.v3.QueryExecutorImpl.sendOneQuery(QueryExecutorImpl.java:1601)
 at org.postgresql.core.v3.QueryExecutorImpl.sendParse(QueryExecutorImpl.java:1191)
 at org.postgresql.core.v3.QueryExecutorImpl.execute(QueryExecutorImpl.java:346)
```

### Answer

During a query with partition conditions, HiveServer optimizes the partitions to avoid full table scanning. All partitions whose metadata meets the conditions need to be queried. However, the **sendOneQuery** interface provided by GaussDB limits the parameter value to **32767** in the **sendParse** method. If the number of partition conditions exceeds **32767**, an exception occurs.

## 10.33.10 Why Cannot I Connect to HiveServer When I Use IBM JDK to Access the Beeline Client?

### Scenario

When users check the JDK version used by the client, if the JDK version is IBM JDK, the Beeline client needs to be reconstructed. Otherwise, the client will fail to connect to HiveServer.

### Procedure

- Step 1** Log in to FusionInsight Manager and choose **System > Permission > User**. In the **Operation** column of the target user, choose **More > Download Authentication Credential**, select the cluster information, and click **OK** to download the keytab file.
- Step 2** Decompress the keytab file and use WinSCP to upload the decompressed **user.keytab** file to the Hive client installation directory on the node to be operated, for example, **/opt/client**.
- Step 3** Run the following command to open the **Hive/component\_env** configuration file in the Hive client directory:

```
vi Hive client installation directory/Hive/component_env
```

Add the following content to the end of the line where **export CLIENT\_HIVE\_URI** is located:

```
\; user.principal=Username@HADOOP.COM\;user.keytab=user.keytab file path/user.keytab
```

----End

## 10.33.11 Description of Hive Table Location (Either Be an OBS or HDFS Path)

### Question

Can Hive tables be stored in OBS or HDFS?

### Answer

- The location of a common Hive table stored on OBS can be set to an HDFS path.
- In the same Hive service, you can create tables stored in OBS and HDFS, respectively.
- For a Hive partitioned table stored on OBS, the location of the partition cannot be set to an HDFS path. (For a partitioned table stored on HDFS, the location of the partition cannot be changed to OBS.)

## 10.33.12 Why Cannot Data Be Queried After the MapReduce Engine Is Switched After the Tez Engine Is Used to Execute Union-related Statements?

### Question

Hive uses the Tez engine to execute union-related statements to write data. After Hive is switched to the MapReduce engine for query, no data is found.

### Answer

When Hive uses the Tez engine to execute the union-related statement, the generated output file is stored in the **HIVE\_UNION\_SUBDIR** directory. After Hive is switched back to the MapReduce engine, files in the directory are not read by default. Therefore, data in the **HIVE\_UNION\_SUBDIR** directory is not read.

In this case, you can set **mapreduce.input.fileinputformat.input.dir.recursive** to **true** to enable union optimization and determine whether to read data in the directory.

## 10.33.13 Why Does Hive Not Support Concurrent Data Writing to the Same Table or Partition?

### Question

Why Does Data Inconsistency Occur When Data Is Concurrently Written to a Hive Table Through an API?

### Answer

Hive does not support concurrent data insertion for the same table or partition. As a result, multiple tasks perform operations on the same temporary data directory, and one task moves the data of another task, causing task data exception. The service logic is modified so that data is inserted to the same table or partition in single thread mode.

## 10.33.14 Why Does Hive Not Support Vectorized Query?

### Question

When the vectorized parameter **hive.vectorized.execution.enabled** is set to **true**, why do some null pointers or type conversion exceptions occur occasionally when Hive on Tez/MapReduce/Spark is executed?

### Answer

Currently, Hive does not support vectorized execution. Many community issues are introduced during vectorized execution and are not resolved stably. The default value of **hive.vectorized.execution.enabled** is **false**. You are advised not to set this parameter to **true**.

## 10.33.15 Why Does Metadata Still Exist When the HDFS Data Directory of the Hive Table Is Deleted by Mistake?

### Question

The HDFS data directory of the Hive table is deleted by mistake, but the metadata still exists. As a result, an error is reported during task execution.

### Answer

This is an exception caused by misoperation. You need to manually delete the metadata of the corresponding table and try again.

Example:

Run the following command to go to the console:

```
source ${BIGDATA_HOME}/FusionInsight_BASE_8.1.0.1/install/FusionInsight-
dbservice-2.7.0/.dbservice_profile
```

```
gsql -p 20051 -U hive -d hivemeta -W HiveUser@
```

Run the **delete from tbls where tbl\_id='xxx';** command.

## 10.33.16 How Do I Disable the Logging Function of Hive?

### Question

How do I disable the logging function of Hive?

### Answer

**Step 1** Log in to the node where the client is installed as user **root**.

**Step 2** Run the following command to switch to the client installation directory, for example, **/opt/client**:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Log in to the Hive client based on the cluster authentication mode.

- In security mode, run the following command to complete user authentication and log in to the Hive client:

```
kinit Component service user
```

```
beeline
```

- In normal mode, run the following command to log in to the Hive client:

- Run the following command to log in to the Hive client as the component service user:

```
beeline -n component service user
```

- If no component service user is specified, the current OS user is used to log in to the Hive client.

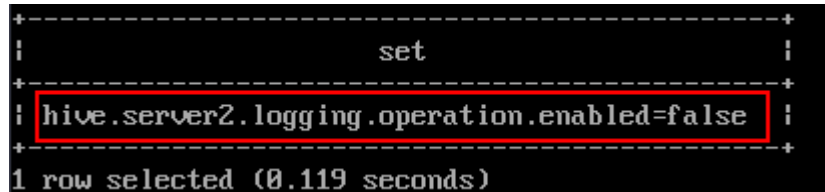
### beeline

**Step 5** Run the following command to disable the logging function:

```
set hive.server2.logging.operation.enabled=false;
```

**Step 6** Run the following command to check whether the logging function is disabled. If the following information is displayed, the logging function is disabled successfully.

```
set hive.server2.logging.operation.enabled;
```



```
+-----+
| set |
+-----+
| hive.server2.logging.operation.enabled=false |
+-----+
1 row selected (0.119 seconds)
```

----End

## 10.33.17 Why Hive Tables in the OBS Directory Fail to Be Deleted?

### Question

In the scenario where the fine-grained permission is configured for multiple MRS users to access OBS, after the permission for deleting Hive tables in the OBS directory is added to the custom configuration of Hive, tables are deleted on the Hive client but still exist in the OBS directory.

### Answer

You do not have the permission to delete directories on OBS. As a result, Hive tables cannot be deleted. In this case, modify the custom IAM policy of the agency and configure Hive with the permission for deleting tables in the OBS directory.

## 10.33.18 Hive Configuration Problems

- The error message "java.lang.OutOfMemoryError: Java heap space." is displayed during Hive SQL execution.  
Solution:
  - For MapReduce tasks, increase the values of the following parameters:  

```
set mapreduce.map.memory.mb=8192;
set mapreduce.map.java.opts=-Xmx6554M;
set mapreduce.reduce.memory.mb=8192;
set mapreduce.reduce.java.opts=-Xmx6554M;
```
  - For Tez tasks, increase the value of the following parameter:  

```
set hive.tez.container.size=8192;
```
- After a column name is changed to a new one using the Hive SQL **as** statement, the error message "Invalid table alias or column reference 'xxx'." is displayed when the original column name is used for compilation.

Solution: Run the **set hive.cbo.enable=true;** statement.

- The error message "Unsupported SubQuery Expression 'xxx': Only SubQuery expressions that are top level conjuncts are allowed." is displayed during Hive SQL subquery compilation.

Solution: Run the **set hive.cbo.enable=true;** statement.

- The error message "CalciteSubquerySemanticException [Error 10249]: Unsupported SubQuery Expression Currently SubQuery expressions are only allowed as Where and Having Clause predicates." is displayed during Hive SQL subquery compilation.

Solution: Run the **set hive.cbo.enable=true;** statement.

- The error message "Error running query: java.lang.AssertionError: Cannot add expression of different type to set." is displayed during Hive SQL compilation.

Solution: Run the **set hive.cbo.enable=false;** statement.

- The error message "java.lang.NullPointerException at org.apache.hadoop.hive.ql.udf.generic.GenericUDAFComputeStats \$GenericUDAFNumericStatsEvaluator.init." is displayed during Hive SQL execution.

Solution: Run the **set hive.map.aggr=false;** statement.

- When **hive.auto.convert.join** is set to **true** (enabled by default) and **hive.optimize.skewjoin** is set to **true**, the error message "ClassCastException org.apache.hadoop.hive.ql.plan.ConditionalWork cannot be cast to org.apache.hadoop.hive.ql.plan.MapredWork" is displayed.

Solution: Run the **set hive.optimize.skewjoin=false;** statement.

- When **hive.auto.convert.join** is set to **true** (enabled by default), **hive.optimize.skewjoin** is set to **true**, and **hive.exec.parallel** is set to **true**, the error message "java.io.FileNotFoundException: File does not exist:xxx/reduce.xml" is displayed.

Solution:

- Method 1: Switch the execution engine to Tez. For details, see [Switching the Hive Execution Engine to Tez](#).
- Method 2: Run the **set hive.exec.parallel=false;** statement.
- Method 3: Run the **set hive.auto.convert.join=false;** statement.

- Error message "NullPointerException at org.apache.hadoop.hive.ql.exec.CommonMergeJoinOperator.mergeJoinComputeKeys" is displayed when Hive on Tez executes bucket map join.

Solution: Run the **set tez.am.container.reuse.enabled=false;** statement.

# 11 Using Hudi

---

## 11.1 Getting Started

### Scenario

This section describes capabilities of Hudi using spark-shell. Using the Spark data source, this section describes how to insert and update a Hudi dataset of the default storage mode Copy-on Write (COW) tables based on code snippets. After each write operation, you will be introduced how to read snapshot and incremental data.

### Prerequisites

- You have created a user and added the user to user groups **hadoop** (primary group) and **hive** on Manager.

### Procedure

- Step 1** Download and install the Hudi client. For details, see [Installing a Client \(MRS 3.x or Later\)](#).

 **NOTE**

Currently, Hudi is integrated in Spark2x. You only need to download the Spark2x client on Manager. For example, the client installation directory is **/opt/client**.

- Step 2** Log in to the node where the client is installed as user **root** and run the following command:

```
cd /opt/client
```

- Step 3** Run the following commands to load environment variables:

```
source bigdata_env
```

```
source Hudi/component_env
```

```
kinit Created user
```

 NOTE

- You need to change the password of the created user, and then run the **kinit** command to log in to the system again.
- In normal mode (Kerberos authentication disabled), you do not need to run the **kinit** command.

**Step 4** Use **spark-shell --master yarn-client** to import Hudi packages to generate test data:

- Import required packages.  

```
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._
```
- Define the table name and storage path to generate test data.  

```
val tableName = "hudi_cow_table"
val basePath = "hdfs://hacluster/tmp/hudi_cow_table"
val dataGen = new DataGenerator
val inserts = convertToStringList(dataGen.generateInserts(10))
val df = spark.read.json(spark.sparkContext.parallelize(inserts, 2))
```

**Step 5** Write data to the Hudi table in overwrite mode.

```
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)
```

**Step 6** Query the Hudi table.

Register a temporary table and query the table.

```
val roViewDF = spark.read.format("org.apache.hudi").load(basePath +
"/**/**/*")
roViewDF.createOrReplaceTempView("hudi_ro_table")
spark.sql("select fare, begin_lon, begin_lat, ts from hudi_ro_table where fare
> 20.0").show()
```

**Step 7** Generate new data and update the Hudi table in append mode.

```
val updates = convertToStringList(dataGen.generateUpdates(10))
```



```
val df = spark.read.json(spark.sparkContext.parallelize(updates, 1))
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(Append).
save(basePath)
```

**Step 8** Query incremental data in the Hudi table.

- Reloading data  

```
spark.read.format("org.apache.hudi").load(basePath + "/*/*/*/*").createOrReplaceTempView("hudi_ro_table")
```
- Perform an incremental query.  

```
val commits = spark.sql("select distinct(_hoodie_commit_time) as commitTime from hudi_ro_table order by commitTime").map(k => k.getString(0)).take(50)
val beginTime = commits(commits.length - 2)
val incViewDF = spark.
read.
format("org.apache.hudi").
option(VIEW_TYPE_OPT_KEY, VIEW_TYPE_INCREMENTAL_OPT_VAL).
option(BEGIN_INSTANTTIME_OPT_KEY, beginTime).
load(basePath);
incViewDF.registerTempTable("hudi_incr_table")
spark.sql("select `_hoodie_commit_time`, fare, begin_lon, begin_lat, ts
from hudi_incr_table where fare > 20.0").show()
```

**Step 9** Perform the point-in-time query.

```
val beginTime = "000"
val endTime = commits(commits.length - 2)
val incViewDF = spark.read.format("org.apache.hudi").
option(VIEW_TYPE_OPT_KEY, VIEW_TYPE_INCREMENTAL_OPT_VAL).
option(BEGIN_INSTANTTIME_OPT_KEY, beginTime).
option(END_INSTANTTIME_OPT_KEY, endTime).
load(basePath);
incViewDF.registerTempTable("hudi_incr_table")
spark.sql("select `_hoodie_commit_time`, fare, begin_lon, begin_lat, ts from
hudi_incr_table where fare > 20.0").show()
```

**Step 10** Delete data.

- Prepare the data to delete.
 

```
val df = spark.sql("select uuid, partitionpath from hudi_ro_table limit 2")
val deletes = dataGen.generateDeletes(df.collectAsList())
```
- Delete the data.
 

```
val df = spark.read.json(spark.sparkContext.parallelize(deletes, 2));
df.write.format("org.apache.hudi").
options(getQuickstartWriteConfigs).
option(OPERATION_OPT_KEY,"delete").
option(PRECOMBINE_FIELD_OPT_KEY, "ts").
option(RECORDKEY_FIELD_OPT_KEY, "uuid").
option(PARTITIONPATH_FIELD_OPT_KEY, "partitionpath").
option(TABLE_NAME, tableName).
mode(Append).
save(basePath);
```
- Query data again.
 

```
val roViewDFAfterDelete = spark.
read.
format("org.apache.hudi").
load(basePath + "/*/*/*/*")
roViewDFAfterDelete.createOrReplaceTempView("hudi_ro_table")
spark.sql("select uuid, partitionPath from hudi_ro_table").show()
```

----End

## 11.2 Basic Operations

### 11.2.1 Hudi Table Schema

When writing data, Hudi generates a Hudi table based on attributes such as the storage path, table name, and partition structure.

Hudi table data files can be stored in the OS file system or distributed file system such as HDFS. To ensure analysis performance and data reliability, HDFS is generally used for storage. Using HDFS as an example, Hudi table storage files are classified into two types.

Permission	Owner	Group	Size	Last Modified	Replication	Block Size	Name
drwxr-xr-x	testcz	hadoop	0 B	Apr 25 15:32	0	0 B	.hoodie
drwxr-xr-x	testcz	hadoop	0 B	Apr 25 15:30	0	0 B	americas
drwxr-xr-x	testcz	hadoop	0 B	Apr 25 15:30	0	0 B	asia

- The **.hoodie** folder stores the log files related to file merging.

<code>drwxr-xr-x</code>	<code>admintest</code>	<code>hadoop</code>	0 B	Mar 30 09:44	0	0 B	<code>.aux</code>
<code>drwxr-xr-x</code>	<code>admintest</code>	<code>hadoop</code>	0 B	Mar 30 11:45	0	0 B	<code>.temp</code>
<code>-rw-r--r--</code>	<code>admintest</code>	<code>hadoop</code>	4.58 KB	Mar 30 09:44	3	128 MB	<code>20210330094435.deltacommith</code>
<code>-rw-r--r--</code>	<code>admintest</code>	<code>hadoop</code>	0 B	Mar 30 09:44	3	128 MB	<code>20210330094435.deltacommith.inflight</code>
<code>-rw-r--r--</code>	<code>admintest</code>	<code>hadoop</code>	0 B	Mar 30 09:44	3	128 MB	<code>20210330094435.deltacommith.requested</code>

- The path containing `_partition_key` stores actual data files and metadata by partition.

Hudi data files of are stored in Parquet base files and Avro log files.

<code>-rw-r--r--</code>	<code>admintest</code>	<code>hadoop</code>	93 B	Mar 30 09:44	3	128 MB	<code>.hoodie_partition_metadata</code>
<code>-rw-r--r--</code>	<code>admintest</code>	<code>hadoop</code>	441.77 KB	Mar 30 09:46	3	128 MB	<code>2b4d098e-4dc8-4633-a22a-dc22f87c57d9-1_0-13-22_20210330094613.parquet</code>
<code>-rw-r--r--</code>	<code>admintest</code>	<code>hadoop</code>	445.28 KB	Mar 30 09:44	3	128 MB	<code>4010e8a8-1b20-4be7-8442-4e30af401e84-0_1-4-8_20210330094435.parquet</code>

## 11.2.2 Write

### 11.2.2.1 Before You Start

Currently, Spark and Flink can be used as write engines for Hudi. The capability of Flink of the current version is weak and not recommended. It will be enhanced in later versions.

### 11.2.2.2 Batch Write

#### Scenario

Hudi provides multiple write modes. For details, see the configuration item `hoodie.datasource.write.operation`. This section describes `upsert`, `insert`, and `bulk_insert`.

- **insert**: The operation process is similar to **upsert**. The query on updated file partitions is not based on indexes. Therefore, **insert** is faster than **upsert**. This operation is recommended for data sources that do not contain updated data. If the data source contains updated data, duplicate data will exist in the data lake.
- **bulk\_insert** (insert in batches): It is used for initial dataset loading. This operation sorts primary keys and then inserts data into a Hudi table by writing data to a common Parquet table. It has the best performance but cannot control small files. The **upsert** and **insert** operations can control small files by using heuristics.
- **upsert** (insert and update): It is the default operation type. Hudi determines whether historical data exists based on the primary key. Historical data is updated, and other data is inserted. This operation is recommended for data sources, such as change data capture (CDC), that include updated data.

#### NOTE

- Primary keys are not sorted during **insert**. Therefore, you are not advised to use **insert** during dataset initialization.
- You are advised to use **insert** if data is new, use **upsert** if data needs to be updated, and use **bulk\_insert** if datasets need to be initialized.

## Writing Data to Hudi Tables In Batches

1. Import the Hudi package to generate test data. For details, see [Step 2](#) to [Step 4](#) in [Getting Started](#).
2. Add the `option("hoodie.datasource.write.operation", "bulk_insert")` parameter to the command for writing data to a Hudi table to set the write mode to `bulk_insert`. For example:

```
df.write.format("org.apache.hudi").
options(getQuickstartWriteConfigs).
option("hoodie.datasource.write.precombine.field", "ts").
option("hoodie.datasource.write.recordkey.field", "uuid").
option("hoodie.datasource.write.partitionpath.field", "").
option("hoodie.datasource.write.operation", "bulk_insert").
option("hoodie.table.name", tableName).
option("hoodie.datasource.write.keygenerator.class",
"org.apache.hudi.keygen.NonpartitionedKeyGenerator").
option("hoodie.datasource.hive_sync.enable", "true").
option("hoodie.datasource.hive_sync.partition_fields", "").
option("hoodie.datasource.hive_sync.partition_extractor_class",
"org.apache.hudi.hive.NonPartitionedExtractor").
option("hoodie.datasource.hive_sync.table", tableName).
option("hoodie.datasource.hive_sync.use_jdbc", "false").
option("hoodie.bulkinsert.shuffle.parallelism", 4).
mode(Overwrite).
save(basePath)
```

### NOTE

- For details about the parameters in the example, see [Table 11-4](#).
- If the Spark DataSource API is used to update the MOR table, small files of the updated data may be merged when a small volume of data is inserted. As a result, some updated data can be found in the read-optimized view of the MOR table.
- If the base file of the data to be updated is a small file, the data to be inserted and new data for update are merged with the base file to generate a new base file instead of being written to logs.

## Configuring Partitions

Hudi supports multiple partitioning modes, such as multi-level partitioning, non-partitioning, single-level partitioning, and partitioning by date. You can select a proper partitioning mode as required. The following describes how to configure different partitioning modes for Hudi.

- Multi-level partitioning

Multi-level partitioning indicates that multiple fields are specified as partition keys. Pay attention to the following configuration items:

Configuration Item	Description
hoodie.datasource.write.partitionpath.field	Configure multiple partition fields, for example, <b>p1</b> , <b>p2</b> , and <b>p3</b> .
hoodie.datasource.hive_sync.partition_fields	Set this parameter to <b>p1</b> , <b>p2</b> , and <b>p3</b> . The values must be the same as the partition fields of <b>hoodie.datasource.write.partitionpath.field</b> .

Configuration Item	Description
hoodie.datasource.write.keygenerator.class	Set this parameter to <b>org.apache.hudi.keygen.ComplexKeyGenerator</b> .
hoodie.datasource.hive_sync.partition_extractor_class	Set this parameter to <b>org.apache.hudi.hive.MultiPartKeysValueExtractor</b> .

- Non-partitioning

Hudi supports non-partitioned tables. Pay attention to the following configuration items:

Configuration Item	Description
hoodie.datasource.write.partitionpath.field	Leave this parameter blank.
hoodie.datasource.hive_sync.partition_fields	Leave this parameter blank.
hoodie.datasource.write.keygenerator.class	Set this parameter to <b>org.apache.hudi.keygen.NonpartitionedKeyGenerator</b> .
hoodie.datasource.hive_sync.partition_extractor_class	Set this parameter to <b>org.apache.hudi.hive.NonPartitionedExtractor</b> .

- Single-level partitioning

It is similar to multi-level partitioning. Pay attention to the following configuration items:

Configuration Item	Description
hoodie.datasource.write.partitionpath.field	Set this parameter to one field, for example, <b>p</b> .
hoodie.datasource.hive_sync.partition_fields	Set this parameter to <b>p</b> . The value must be the same as the partition field of <b>hoodie.datasource.write.partitionpath.field</b>
hoodie.datasource.write.keygenerator.class	(Optional) The default value is <b>org.apache.hudi.keygen.SimpleKeyGenerator</b> .
hoodie.datasource.hive_sync.partition_extractor_class	Set this parameter to <b>org.apache.hudi.hive.MultiPartKeysValueExtractor</b> .

- Partitioning by date

The **date** field is specified as the partition field. Pay attention to the following configuration items:

Configuration Item	Description
hoodie.datasource.write.partitionpath.field	Set this parameter to the <b>date</b> field, for example, <b>operationTime</b> .
hoodie.datasource.hive_sync.partition_fields	Set this parameter to <b>operationTime</b> . The value must be the same as the preceding partition field.
hoodie.datasource.write.keygenerator.class	(Optional) The default value is <b>org.apache.hudi.keygen.SimpleKeyGenerator</b> .
hoodie.datasource.hive_sync.partition_extractor_class	Set this parameter to <b>org.apache.hudi.hive.SlashEncodedDayPartitionValueExtractor</b> .

 **NOTE**

Date format for **SlashEncodedDayPartitionValueExtractor** must be *yyyy/mm/dd*.

- Partition sorting

Configuration Item	Description
hoodie.bulkinsert.user.defined.partition_sorter.class	Specifies the partition sorting class. You can customize a sorting method. For details, see the sample code.

 **NOTE**

By default, **bulk\_insert** sorts data by character and applies only to primary keys of StringType.

### 11.2.2.3 Stream Write

#### Stream Write Using HoodieDeltaStreamer

The HoodieDeltaStreamer tool provided by Hudi supports stream write. You can also use SparkStreaming to write data in microbatch mode. HoodieDeltaStreamer provides the following functions:

- Supports multiple data sources, such as Kafka and DFS.
- Manages checkpoints, rollback, and recovery to ensure exactly-once semantics.

- Supports user-defined transformations.

Example:

Prepare the configuration file **kafka-source.properties**.

```
#Hudi configuration
hoodie.datasource.write.recordkey.field=id
hoodie.datasource.write.partitionpath.field=age
hoodie.upsert.shuffle.parallelism=100
#hive config
hoodie.datasource.hive_sync.table=hudimor_deltastreamer_partition
hoodie.datasource.hive_sync.partition_fields=age
hoodie.datasource.hive_sync.partition_extractor_class=org.apache.hudi.hive.MultiPartKeysValueExtractor
hoodie.datasource.hive_sync.use_jdbc=false
hoodie.datasource.hive_sync.support_timestamp=true
Kafka Source topic
hoodie.deltastreamer.source.kafka.topic=hudimor_deltastreamer_partition
#checkpoint
hoodie.deltastreamer.checkpoint.provider.path=hdfs://hacluster/tmp/huditest/
hudimor_deltastreamer_partition
Kafka props
The kafka cluster we want to ingest from
bootstrap.servers= xx.xx.xx.xx:xx
auto.offset.reset=earliest
#auto.offset.reset=latest
group.id=hoodie-delta-streamer
offset.rang.limit=10000
```

Run the following commands to specify the HoodieDeltaStreamer execution parameters:

**spark-submit --master yarn**

**--jars /opt/hudi-java-examples-1.0.jar** // Specify the Hudi **jars** directory required for Spark running.

**--driver-memory 1g**

**--executor-memory 1g --executor-cores 1 --num-executors 2 --conf spark.kryoserializer.buffer.max=128m**

**--driver-class-path /opt/client/Hudi/hudi/conf:/opt/client/Hudi/hudi/lib/\*:/opt/client/Spark2x/spark/jars\*/:/opt/hudi-examples-0.6.1-SNAPSHOT.jar:/opt/hudi-examples-0.6.1-SNAPSHOT-tests.jar** // Specify the Hudi **jars** directory required by the Spark driver.

**--class org.apache.hudi.utilities.deltastreamer.HoodieDeltaStreamer spark-internal**

**--props file:///opt/kafka-source.properties** // Specify the configuration file. You need to set the configuration file path to the HDFS path when submitting tasks in yarn-cluster mode.

**--target-base-path /tmp/huditest/hudimor1\_deltastreamer\_partition** // Specify the path of the Hudi table.

**--table-type MERGE\_ON\_READ** // Specify the type of the Hudi table to be written.

**--target-table hudimor\_deltastreamer\_partition** // Specify the Hudi table name.

**--source-ordering-field name** // Specify the columns to be pre-combined in the Hudi table.

**--source-class org.apache.hudi.utilities.sources.JsonKafkaSource** // Set the consumed data source to **JsonKafkaSource**. Different source classes are specified based on different data sources.

**--schemaprovider-class com..bigdata.hudi.examples.DataSchemaProviderExample** // Specify the schema required by the Hudi table.

**--transformer-class com..bigdata.hudi.examples.TransformerExample** // Specify how to process the data obtained from the data source. Set this parameter based on service requirements.

**--enable-hive-sync** // Enable Hive synchronization to synchronize the Hudi table to Hive.

**--continuous** // Set the stream processing mode to **continuous**.

### 11.2.2.4 Synchronizing Hudi Table Data to Hive

You can run **run\_hive\_sync\_tool.sh** to synchronize data in the Hudi table to Hive.

For example, run the following command to synchronize the Hudi table in the **hdfs://hacluster/tmp/huditest/hudimor1\_deltastreamer\_partition** directory on HDFS to the Hive table **table hive\_sync\_test3** with **unite**, **country**, and **state** as partition keys:

```
run_hive_sync_tool.sh --partitioned-by unite,country,state --base-path hdfs://hacluster/tmp/huditest/hudimor1_deltastreamer_partition --table hive_sync_test3 --partition-value-extractor org.apache.hudi.hive.MultiPartKeyValueExtractor --support-timestamp
```

**Table 11-1** Parameter description

Command	Description	Mandatory or Not (Yes or No)	Default Value
<b>--database</b>	Specifies the Hive database name.	No	default
<b>--table</b>	Specifies the Hive table name.	Yes	-
<b>--base-file-format</b>	Specifies the file format ( <b>PARQUET</b> or <b>HFILE</b> ).	No	PARQUET
<b>--user</b>	Specifies the Hive username.	No	-
<b>--pass</b>	Specifies the Hive password.	No	-
<b>--jdbc-url</b>	Specifies the Hive JDBC connection URL.	No	-



Command	Description	Mandatory or Not (Yes or No)	Default Value
--base-path	Specifies the storage path of the Hudi table to be synchronized.	Yes	-
--partitioned-by	Specifies the partition key.	No	-
--partition-value-extractor	Specifies the partition class. PartitionValueExtractor needs to be implemented. The partition value can be extracted from the HDFS path.	No	SlashEncodedDayPartitionValueExtractor
--assume-date-partitioning	Creates partitions in yyyy/mm/dd format to support backward compatibility.	No	false
--use-pre-apache-input-format	Use InputFormat in the <b>com.uber.hoodie</b> package to replace the one in the <b>org.apache.hudi</b> package. Do not use this command except for migrating projects from <b>com.uber.hoodie</b> to <b>org.apache.hudi</b> .	No	false
--use-jdbc	Uses Hive JDBC connection.	No	true
--auto-create-database	Specifies whether to automatically create a Hive database.	No	true
--skip-ro-suffix	Specifies whether to skip the read-optimized view with the <b>_ro</b> suffix during registration.	No	false
--use-file-listing-from-metadata	Specifies whether to obtain the file list from the Hudi metadata.	No	false
--verify-metadata-file-listing	Specifies whether to verify the file list in the Hudi metadata based on the file system.	No	false

Command	Description	Mandatory or Not (Yes or No)	Default Value
--help/-h	Specifies whether to display help information.	No	false
--support-timestamp	Specifies whether to convert <b>TIMESTAMP_MICROS</b> of INT64 to Hive timestamp.	No	false
--decode-partition	Specifies whether to decode the partition value if the partition is encoded during the write process.	No	false
--batch-sync-num	Specifies the number of Hive partitions to be synchronized in each batch.	No	1000

 **NOTE**

During Hive synchronization, if the table does not exist, an external table is created and partitions are added. If the table exists, check whether table schemas are different. If they are different, replace the table. Check whether new partitions exist. If new partitions exist, partitions are added accordingly.

Therefore, there are the following restrictions when Hive synchronization is used:

- Fields can only be added to the schema and cannot be modified or deleted.
- Partition directories can only be added but cannot be deleted.
- **Overwrite** can only overwrite the Hudi table. The Hive table cannot be overwritten synchronously.
- Do not use the timestamp type as the partition column when synchronizing a Hudi table to Hive.

## 11.2.3 Read

### 11.2.3.1 Overview

Read operations on Hudi tables are based on three types of views. You can select a proper view for query as required.

Hudi supports multiple query engines, including Spark and Hive. For details, see [Table 11-2](#) and [Table 11-3](#).

**Table 11-2** COW tables

Query Engine	Real-time View/Read-optimized View	Incremental View
Hive	Y	Y
Spark (SparkSQL)	Y	Y
Spark (SparkDataSource API)	Y	Y

**Table 11-3** MOR tables

Query Engine	Real-time View	Incremental View	Read-optimized View
Hive	Y	Y	Y
Spark (SparkSQL)	Y	Y	Y
Spark (SparkDataSource API)	Y	Y	Y

 **CAUTION**

- Currently, the partition deduction capability is not supported when Hudi uses the Spark DataSource API to read data. For example, when the DataSource API is used to query a bootstrap table, the partition field may not be displayed or may be displayed as null.
- For an incremental view, set **hoodie.hudicow.consume.mode** to **INCREMENTAL**. This parameter applies only to queries on the incremental view and cannot be used for queries on other types of Hudi tables or queries on other tables. You can set **hoodie.hudicow.consume.mode** to **SNAPSHOT** or any value to restore the configuration.

### 11.2.3.2 Reading COW Table Views

- Reading the real-time view (using Hive and SparkSQL as an example):  
Directly read the Hudi table stored in Hive.

```
select count(*) from test;
```

- Reading the real-time view (using the Spark DataSource API as an example):  
This is similar to reading a common DataSource table.

**QUERY\_TYPE\_OPT\_KEY** must be set to **QUERY\_TYPE\_SNAPSHOT\_OPT\_VAL**.

```
spark.read.format("hudi")
.option(QUERY_TYPE_OPT_KEY, QUERY_TYPE_SNAPSHOT_OPT_VAL) // Set the query type to the real-time view.
.load("/tmp/default/cow_bugx/") // Specify the path of the Hudi table to read.
.createTempView("mycall")
spark.sql("select * from mycall").show(100)
```

- Reading the incremental view (using Hive as an example):  

```
set hoodie.test.consume.mode=INCREMENTAL; // Specify the incremental reading mode.
set hoodie.test.consume.max.commits=3; // Specify the maximum number of commits to be consumed.
set hoodie.test.consume.start.timestamp=20201227153030; // Specify the initial incremental pull commit.
select count(*) from default.test where `_hoodie_commit_time`>'20201227153030'; // This filtering condition must be added, and the value is the initial incremental pull commit.
```
- Reading the incremental view (using Spark SQL as an example):  

```
set hoodie.test.consume.mode=INCREMENTAL; // Specify the incremental reading mode.
set hoodie.test.consume.start.timestamp=20201227153030; // Specify the initial incremental pull commit.
set hoodie.test.consume.end.timestamp=20210308212318; // Specify the end commit of the incremental pull. If this parameter is not specified, the latest commit is used.
select count(*) from default.test where `_hoodie_commit_time`>'20201227153030'; // This filtering condition must be added, and the value is the initial incremental pull commit.
```
- Reading the incremental view (using the Spark DataSource API as an example):  

**QUERY\_TYPE\_OPT\_KEY** must be set to **QUERY\_TYPE\_INCREMENTAL\_OPT\_VAL**.

```
spark.read.format("hudi")
.option(QUERY_TYPE_OPT_KEY, QUERY_TYPE_INCREMENTAL_OPT_VAL) // Set the query type to the incremental mode.
.option(BEGIN_INSTANTTIME_OPT_KEY, "20210308212004") // Specify the initial incremental pull commit.
.option(END_INSTANTTIME_OPT_KEY, "20210308212318") // Specify the end commit of the incremental pull.
.load("/tmp/default/cow_bugx/") // Specify the path of the Hudi table to read.
.createTempView("mycall") // Register as a Spark temporary table.
spark.sql("select * from mycall where `_hoodie_commit_time`>'20210308211131'")// Start the query. The statement is the same as the Hive incremental query statement.
.show(100, false)
```
- Reading the read-optimized view: The read-optimized view of COW tables is equivalent to the real-time view.

### 11.2.3.3 Reading MOR Table Views

After the MOR table is synchronized to Hive, the following two tables are synchronized to Hive: *Table name<sub>rt</sub>* and *Table name<sub>ro</sub>*. The table suffixed with **rt** indicates the real-time view, and the table suffixed with **ro** indicates the read-optimized view. For example, the name of the Hudi table to be synchronized to Hive is **test**. After the table is synchronized to Hive, two more tables **test<sub>rt</sub>** and **test<sub>ro</sub>** are generated in the Hive table.

- Reading the real-time view (using Hive and SparkSQL as an example):  
 Directly read the Hudi table with suffix **\_rt** stored in Hive.  

```
select count(*) from test_rt;
```
- Reading the real-time view (using the Spark DataSource API as an example):  
 The operations are the same as those for the COW table. For details, see the operations for the COW table.
- Reading the incremental view (using Hive as an example):  

```
set hive.input.format=org.apache.hudi.hadoop.hive.HoodieCombineHiveInputFormat; // This parameter does not need to be specified for SparkSQL.
set hoodie.test.consume.mode=INCREMENTAL;
set hoodie.test.consume.max.commits=3;
set hoodie.test.consume.start.timestamp=20201227153030;
select count(*) from default.test_rt where `_hoodie_commit_time`>'20201227153030';
```
- Reading the incremental view (using Spark SQL as an example):  

```
set hoodie.test.consume.mode=INCREMENTAL;
set hoodie.test.consume.start.timestamp=20201227153030; // Specify the initial incremental pull
```

```
commit.
set hoodie.test.consume.end.timestamp=20210308212318; // Specify the end commit of the
incremental pull. If this parameter is not specified, the latest commit is used.
select count(*) from default.test_rt where `_hoodie_commit_time`>'20201227153030';
```

- Incremental view (using the Spark DataSource API as an example): The operations are the same as those for the COW table. For details, see the operations for the COW table.
- Reading the read-optimized view (using Hive and SparkSQL as an example): Directly read the Hudi table with suffix **\_ro** stored in Hive.  

```
select count(*) from test_ro;
```
- Reading the read-optimized view (using the Spark DataSource API as an example): This is similar to reading a common DataSource table.

**QUERY\_TYPE\_OPT\_KEY** must be set to  
**QUERY\_TYPE\_READ\_OPTIMIZED\_OPT\_VAL**.

```
spark.read.format("hudi")
.option(QUERY_TYPE_OPT_KEY, QUERY_TYPE_READ_OPTIMIZED_OPT_VAL) // Set the query type to
the read-optimized view.
.load("/tmp/default/mor_bugx/") // Specify the path of the Hudi table to read.
.createTempView("mycall")
spark.sql("select * from mycall").show(100)
```

## 11.2.4 Data Management and Maintenance

### 11.2.4.1 Clustering

#### Introduction

Clustering reorganizes data layout to improve query performance without affecting the ingestion speed.

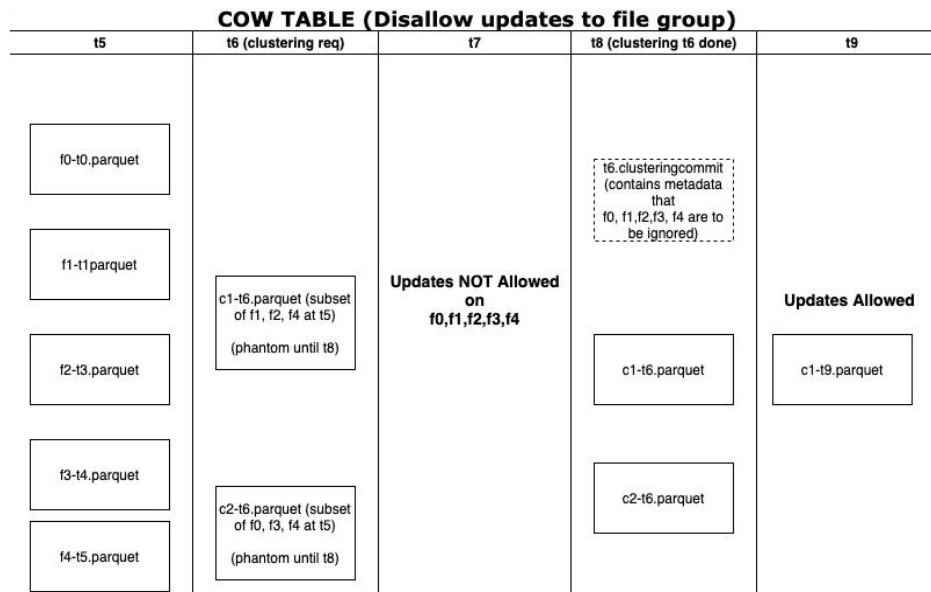
#### Architecture

Hudi provides different operations, such as **insert**, **upsert**, and **bulk\_insert**, through its write client API to write data to a Hudi table. To weight between file size and speed of importing data into the data lake, Hudi provides **hoodie.parquet.small.file.limit** to configure the minimum file size. You can set it to **0** to force new data to be written to new file groups, or to a higher value to ensure that new data is "padded" to existing small file groups until it reaches the specified size, but this increases ingestion latency.

To support fast ingestion without affecting query performance, the clustering service is introduced to rewrite data to optimize the layout of Hudi data lake files.

The clustering service can run asynchronously or synchronously. It adds a new operation type called **REPLACE**, which will mark the clustering operation in the Hudi metadata timeline.

Clustering service is based on the MVCC design of Hudi to allow new data to be inserted. Clustering operations run in the background to reformat data layout, ensuring snapshot isolation between concurrent readers and writers.



Clustering is divided into two parts:

- Scheduling clustering: Create a clustering plan using a pluggable clustering strategy.
  - a. Identify files that are eligible for clustering: Depending on the selected clustering strategy, the scheduling logic will identify the files eligible for clustering.
  - b. Group files that are eligible for clustering based on specific criteria. The data size of each group must be a multiple of **targetFileSize**. Grouping is a part of the strategy defined in the plan. Additionally, there is an option to control group size to improve parallelism and avoid shuffling large volumes of data.
  - c. Save the clustering plan to the timeline in Avro metadata format.
- Execute clustering: Process the plan using an execution strategy to create new files and replace old files.
  - a. Read the clustering plan and obtain **clusteringGroups** that marks the file groups to be clustered.
  - b. Instantiate appropriate strategy class for each group using **strategyParams** (for example, **sortColumns**) and apply the strategy to rewrite data.
  - c. Create a **REPLACE** commit and update the metadata in HoodieReplaceCommitMetadata.

## How to Execute Clustering

1. Executing clustering synchronously

Add the following configuration parameters when the data write operation is performed:

**option("hoodie.clustering.inline", "true").**

**option("hoodie.clustering.inline.max.commits", "4").**

**option("hoodie.clustering.plan.strategy.target.file.max.bytes", "1073741824").**

```
option("hoodie.clustering.plan.strategy.small.file.limit", "629145600").
option("hoodie.clustering.plan.strategy.sort.columns",
"column1,column2").
```

2. Executing clustering asynchronously

```
spark-submit --master yarn --class
org.apache.hudi.utilities.HoodieClusteringJob /opt/client/Hudi/hudi/lib/
hudi-utilities*.jar --schedule --base-path <table_path> --table-name
<table_name> --props /tmp/clusteringjob.properties --spark-memory 1g
spark-submit --master yarn --driver-memory 16G --executor-memory 12G
--executor-cores 4 --num-executors 4 --class
org.apache.hudi.utilities.HoodieClusteringJob /opt/client/Hudi/hudi/lib/
hudi-utilities*.jar --base-path <table_path> --instant-time
20210605112954 --table-name <table_name> --props /tmp/
clusteringjob.properties --spark-memory 12g
```

3. Specifying the ordering mode and sequence of clustering

Currently, clustering supports three sorting modes: Linear, Z-Order, and Hilbert, which can be configured in option or set mode.

- Linear ordering: a common but default ordering mode, which applies to ordering one field or multiple low-level fields.
- Z-order or Hilbert: multi-dimensional ordering, which is available when you set **hoodie.layout.optimize.strategy** to **z-order** or **hilbert**.

These two ordering modes apply to sorting 2 to 4 fields, for example, multiple fields involved in a query condition.

Hilbert has a better multi-dimensional ordering effect than Z-order but lower ordering efficiency.

For details, see [Configuration Reference](#).

---

 CAUTION

1. The sorting column of clustering cannot be null. This is restricted by Spark RDD.
2. If the value of **target.file.max.bytes** is large, increase the value of **--spark-memory** to execute clustering. Otherwise, the executor memory overflow occurs.
3. Currently, the clean mechanism cannot be used to delete junk files generated after the clustering fails.
4. After the clustering, sizes of new files may be different, causing data skew.
5. Clustering and upsert operations cannot be performed at the same time.
6. If the clustering is in the **inflight** state, the files in the file group do not support the **update** operation.

---

### 11.2.4.2 Cleaning

Cleaning is used to delete data of versions that are no longer required.

Hudi uses the cleaner working in the background to continuously delete unnecessary data of old versions. You can configure **hoodie.cleaner.policy** and

**hoodie.cleaner.commits.retained** to use different cleaning policies and determine the number of saved commits.

You can use either of the following methods to perform cleaning:

- Synchronous cleaning is controlled by the **hoodie.clean.automatic** parameter, which is automatically enabled by default.

Disable synchronous cleaning:

When a data source is written, you can use **.option("hoodie.clean.automatic", "false")** to disable automatic cleaning.

When spark-sql is written, you can use **set hoodie.clean.automatic=false;** to disable automatic cleaning.

- You can use spark-sql to perform asynchronous cleaning.

For more cleaning parameters, see [Compaction and Cleaning Configurations](#).

### 11.2.4.3 Compaction

A compaction merges base and log files of MOR tables.

For MOR tables, data is stored in columnar Parquet files and row-based Avro files, updates are recorded in incremental files, and then a synchronous or asynchronous compaction is performed to generate new versions of columnar files. MOR tables can reduce data ingestion latency, so an asynchronous compaction that does not block ingestion is useful.

- An asynchronous compaction is performed in the following two steps:
  - a. Scheduling a compaction: A compaction is completed by the job of importing data into the data lake. In this step, Hudi scans partitions and selects the file slices to be compacted. A compaction plan is finally written to the Hudi timeline.
  - b. Executing a compaction: A separate process or thread reads the compaction plan and performs the compaction of file slices.
- Compaction can be synchronous or asynchronous.
  - The synchronization mode is controlled by the **hoodie.compact.inline** parameter. The default value is **true**, indicating that the compaction scheduling plan is automatically generated and compaction is executed.

- Disable synchronous compaction.

When a data source is written, run the **.option("hoodie.compact.inline", "false")** command to disable automatic compaction.

When spark-sql is written, run the **set hoodie.compact.inline=false;** command to disable automatic compaction.

- Only compaction scheduling is generated synchronously, but compaction is not executed.

- · A data source can be written by configuring the following option parameters:

**option("hoodie.compact.inline", "true").**

**option("hoodie.schedule.compact.only.inline", "true").**



```
option("hoodie.run.compact.only.inline", "false").
```

- spark-sql can be written by configuring the following set parameters:

```
set hoodie.compact.inline=true;
```

```
set hoodie.schedule.compact.only.inline=true;
```

```
set hoodie.run.compact.only.inline=false;
```

- The asynchronous mode is implemented by spark-sql.

To execute only the compaction scheduling plan that has been generated during asynchronous compaction without creating a new scheduling plan, run the following commands to configure set parameters:

```
set hoodie.compact.inline=true;
```

```
set hoodie.schedule.compact.only.inline=false;
```

```
set hoodie.run.compact.only.inline=true;
```

For more compaction parameters, see [Compaction and Cleaning Configurations](#).

#### NOTE

To ensure the maximum efficiency of data import into the lake, you are advised to generate compaction scheduling plans synchronously and execute compaction scheduling plans asynchronously.

### 11.2.4.4 Savepoint

Savepoints are used to save and restore custom version data.

Savepoints provided by Hudi can save different commits so that the cleaner program does not delete them. You can use rollback to restore them later.

Use spark-sql to manage savepoints.

Example:

- Creating a savepoint  
call `create_savepoints('hudi_test1', '20220908155421949');`
- Viewing all existing savepoints  
call `show_savepoints(table =>'hudi_test1');`
- Rolling back a savepoint  
call `rollback_savepoints('hudi_test1', '20220908155421949');`

#### NOTE

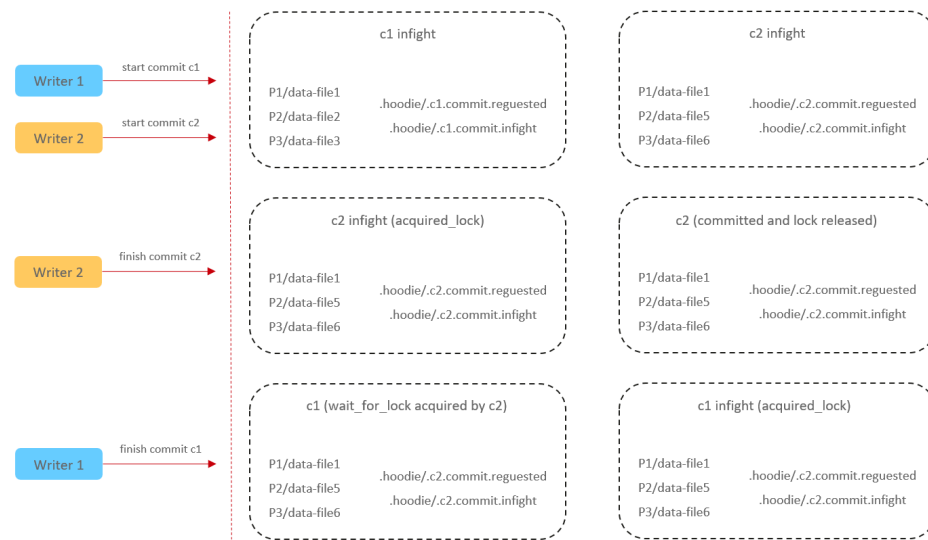
MOR tables do not support savepoints.

### 11.2.4.5 Single-Table Concurrency Control

By default, Hudi does not support concurrent write and compaction operations on a single table. When Flink or Spark is used to write data or Spark is used to perform compaction operations, Hudi attempts to obtain the lock corresponding to the table. (ZooKeeper in the cluster provides the distributed lock service and the configuration takes effect automatically.) If the lock fails to obtain, the task exits directly to prevent table damage caused by concurrent operations of the lock task. If the concurrent write function is enabled for a single Hudi table, these functions automatically become invalid.

## Hudi Single-Table Concurrent Write Solution

1. Uses an external service (ZooKeeper or Hive MetaStore) as the distributed mutex lock service.
2. Files can be concurrently written, but commits cannot be concurrent. The commit operation is encapsulated in a transaction.
3. When the commit operation is performed, the system performs conflict check. If the modified file list in the current commit operation overlaps with the file list in the commit operation after the instance time, the commit operation fails and the write operation is invalid.



## Precautions for Using the Concurrency Mechanism

1. The current Hudi concurrency mechanism cannot ensure that the primary key of the table is unique after data is written. You need to ensure that the primary key is unique.
2. For incremental queries, data consumption and checkpoints may be out of order. As a result, multiple concurrent write operations are completed at different time points.
3. Concurrent write is supported only after this feature is enabled.

## How to Use the Concurrency Mechanism

1. Enable the concurrent write mechanism.

**hoodie.write.concurrency.mode=optimistic\_concurrency\_control**

**hoodie.cleaner.policy.failed.writes=LAZY**

2. Sets the concurrent lock mode.

Hive MetaStore:

**hoodie.write.lock.provider=org.apache.hudi.hive.HiveMetastoreBasedLockProvider**

**hoodie.write.lock.hivemetastore.database=<database\_name>**

**hoodie.write.lock.hivemetastore.table=<table\_name>**

ZooKeeper:

```
hoodie.write.lock.provider=org.apache.hudi.client.transaction.lock.ZookeeperBasedLockProvider
hoodie.write.lock.zookeeper.url=<zookeeper_url>
hoodie.write.lock.zookeeper.port=<zookeeper_port>
hoodie.write.lock.zookeeper.lock_key=<table_name>
hoodie.write.lock.zookeeper.base_path=<table_path>
```

For details about more parameters, see [Configuration Reference](#).

**CAUTION**

If **cleaner policy** is set to **Lazy**, the system can only check whether the written files expire but cannot check and clear junk files generated by historical writes. That is, junk files cannot be automatically cleared in concurrent scenarios.

## 11.2.5 Using the Hudi Client

### 11.2.5.1 Operating a Hudi Table Using hudi-cli.sh

#### Prerequisites

- For a cluster with Kerberos authentication enabled, a user has been created on FusionInsight Manager of the cluster and associated with user groups **hadoop** and **hive**.
- The Hudi cluster client has been downloaded and installed.

#### Basic Operations

1. Log in to the cluster client as user **root** and run the following commands:  

```
cd {Client installation directory}
source bigdata_env
source Hudi/component_env
kinit Created user
```
2. Run the **hudi-cli.sh** command to access the Hudi client.  

```
cd {Client installation directory}/Hudi/hudi/bin/
./hudi-cli.sh
```

```
[root@kwephispra44948 bin]# hudi-cli.sh
Running : java -cp /opt/prober/client/Hudi/hudi/conf:/opt/prober/client/Hudi/hudi/lib/*:/opt/prober/client/Spark2x/spark/jars/* -Djava.security.krb5.conf=/opt/prober/client/KrbClient/kerberos/var/krb5kdc/krb5.conf -Dzookeeper.server.principal=zookeeper/hadoop.p.hadooptest.com -Djava.security.auth.login.config=/opt/prober/client/Hudi/hudi/conf/jaas.conf -Dzookeeper.kinit=/opt/prober/client/KrbClient/kerberos/bin/kinit -DSPARK_CONF_DIR=/opt/prober/client/Hudi/hudi/conf -DHADOOP_CONF_DIR=/opt/prober/client/Hudi/hudi/conf org.springframework.shell.Bootstrap
2021-09-17 15:24:08.035 | INFO | main | Loading XML bean definitions from URL [jar:file:/opt/prober/client/Hudi/hudi/lib/hudi-cli-0.9.0-hw-e1-312081-SNAPSHOT.jar!/META-INF/spring/spring-shell-plugin.xml] | org.springframework.beans.factory.xml.XmlBeanDefinitionReader.loadBeanDefinitions(XmlBeanDefinitionReader.java:317)
2021-09-17 15:24:08.627 | INFO | main | Refreshing org.springframework.context.support.GenericApplicationContext@59906517: start up date [Fri Sep 17 15:24:08 CST 2021]; root of context hierarchy | org.springframework.context.support.GenericApplicationContext.prepareRefresh(AbstractApplicationContext.java:578)
2021-09-17 15:24:08.827 | INFO | main | JSR-330 'javax.inject.Inject' annotation found and supported for autowiring | org.springframework.beans.factory.annotation.AutowiredAnnotationBeanPostProcessor.<init>(AutowiredAnnotationBeanPostProcessor.java:153)
Table command getting loaded
HoodieSplashScreen loaded

 / ___ /
 / / / /
 / / / /
 / / / /
 / / / /
/ / / /

 Apache Hudi CLI

Welcome to Apache Hudi CLI. Please type help if you are looking for help.
```

3. Run the following example commands as required.
  - Viewing help information  
**help** // View all Hudi CLI commands.  
**help 'command'** // View the help information and parameter list of a certain command.
  - Connecting to a table  
**connect --path '/tmp/huditest/test\_table'**
  - Viewing table information  
**desc**
  - Viewing compaction plans  
**compactions show all**
  - Viewing cleaning plans  
**cleans show**
  - Performing the cleaning operation  
**cleans run**
  - Viewing commit information  
**commits show**
  - Viewing the partition where the commit is written to  
**commit showpartitions --commit 20210127153356**

 NOTE

- 20210127153356* indicates the commit timestamp.
- Viewing the file where the commit is written to  
**commit showfiles --commit 20210127153356**
- Comparing the commit information of two tables  
**commits compare --path /tmp/hudimor/mytest100**
- Rolling back a commit (Only the last commit can be rolled back.)  
**commit rollback --commit 20210127164905**
- Scheduling a compaction  
**compaction schedule --hoodieConfigs**  
*'hoodie.compaction.strategy=org.apache.hudi.table.action.compact.strategy.BoundedIOCompactionStrategy,hoodie.compaction.target.io=1,hoodie.compaction.inline.max.delta.commits=1'*
- Performing a compaction  
**compaction run --parallelism 100 --sparkMemory 1g --retry 1 --compactionInstant 20210602101315 --hoodieConfigs**  
*'hoodie.compaction.strategy=org.apache.hudi.table.action.compact.strategy.BoundedIOCompactionStrategy,hoodie.compaction.target.io=1,hoodie.compaction.inline.max.delta.commits=1' --propsFilePath hdfs://hacluster/tmp/default/tb\_test\_mor/.hoodie/hoodie.properties --schemaFilePath /tmp/default/tb\_test\_mor/.hoodie/compact\_tb\_base.json*
- Creating a savepoint  
**savepoint create --commit 20210318155750**

- Rolling back a specified savepoint  
**savepoint rollback --savepoint 20210318155750**

---

 **CAUTION**

1. If the commit operation causes metadata conflicts, you can run the **commit rollback** and **savepoint rollback** commands to roll back data, but the Hive metadata cannot be rolled back. In this case, you can delete the Hive table and manually synchronize data.
  2. The **commit rollback** command rolls back only the latest commit, and the **savepoint rollback** command rolls back only the latest savepoint. You cannot specify a commit or savepoint to roll back.
- 

## 11.2.6 Configuration Reference

### 11.2.6.1 Write Configuration

Table 11-4 Write configuration

Parameter	Description	Default Value
hoodie.datasource.write.table.name	Specifies the name of the Hudi table to be written.	None

Parameter	Description	Default Value
hoodie.datasource.write.operation	<p>Specifies the operation type of writing the Hudi table. Currently, <b>upsert</b>, <b>delete</b>, <b>insert</b>, and <b>bulk_insert</b> are supported.</p> <ul style="list-style-type: none"> <li>• <b>upsert</b>: updates and inserts data.</li> <li>• <b>delete</b>: deletes data.</li> <li>• <b>insert</b>: inserts data.</li> <li>• <b>bulk_insert</b>: imports data during initial table creation. Do not <b>upsert</b> or <b>insert</b> during initial table creation.</li> <li>• <b>insert_overwrite</b>: performs insert and overwrite operations on static partitions.</li> <li>• <b>insert_overwrite_table</b>: performs insert and overwrite operations on dynamic partitions. It does not immediately delete the entire table or overwrite the table. Instead, it overwrites the metadata of the Hudi table logically, and Hudi deletes useless data through the clean mechanism. Its efficiency is higher than that of the combination of <b>bulk_insert</b> and <b>overwrite</b>.</li> </ul>	upsert
hoodie.datasource.write.table.type	<p>Specifies the Hudi table type. Once the table type is specified, this parameter cannot be modified. The value can be <b>MERGE_ON_READ</b>.</p>	COPY_ON_WRITE
hoodie.datasource.write.precombine.field	<p>Merges and reduplicates rows with the same key before write.</p>	ts

Parameter	Description	Default Value
hoodie.datasource.write.payload.class	Specifies the class used to merge the records to be updated and the updated records during update. This parameter can be customized. You can compile it yourself to implement your merge logic.	org.apache.hudi.common.model.DefaultHoodieRecordPayload
hoodie.datasource.write.recordkey.field	Specifies the primary key of the Hudi table. The Hudi table must have a unique primary key.	uuid
hoodie.datasource.write.partitionpath.field	Specifies the partition key. This parameter is used together with <b>hoodie.datasource.write.keygenerator.class</b> to meet the requirements of different partition scenarios.	None
hoodie.datasource.write.hive_style_partitioning	Specifies whether the partition mode is the same as that of Hive. You are advised to set this parameter to <b>true</b> .	true
hoodie.datasource.write.keygenerator.class	Generates the primary key and partition mode when used together with <b>hoodie.datasource.write.partitionpath.field</b> and <b>hoodie.datasource.write.recordkey.field</b> . <b>NOTE</b> If the value of this parameter is different from that saved in the table, a message is displayed, indicating that the value must be the same.	org.apache.hudi.keygen.ComplexKeyGenerator

### 11.2.6.2 Configuration of Hive Table Synchronization

Parameter	Description	Default Value
hoodie.datasource.hive_sync.enable	Specifies whether to synchronize the Hudi table information to Hive MetaStore. <b>CAUTION</b> You are advised to set this parameter to <b>true</b> to use Hive to manage the Hudi table.	false
hoodie.datasource.hive_sync.database	Specifies the name of the database to be synchronized to Hive.	default
hoodie.datasource.hive_sync.table	Specifies the name of the table to be synchronized to Hive. You are advised to set this parameter to the value of <b>hoodie.datasource.write.table.name</b> .	unknown
hoodie.datasource.hive_sync.username	Specifies the username used for Hive synchronization.	hive
hoodie.datasource.hive_sync.password	Specifies the password used for Hive synchronization.	hive
hoodie.datasource.hive_sync.jdbcurl	Specifies the Hive JDBC URL for connection.	""
hoodie.datasource.hive_sync.use_jdbc	Specifies whether to use Hive JDBC to connect to Hive for the Hudi table information synchronization. You are advised to set this parameter to <b>false</b> , which indicates that the JDBC connection configuration is invalid.	true
hoodie.datasource.hive_sync.partition_fields	Specifies the Hive partition columns.	""
hoodie.datasource.hive_sync.partition_extractor_class	Specifies the class used to extract Hudi partition field values and convert them into Hive partition columns.	org.apache.hudi.hive.SlashEncodedDayPartitionValueExtractor



Parameter	Description	Default Value
hoodie.datasource.hive_sync.support_timestamp	If the Hudi table contains a field of the timestamp type, set this parameter to <b>true</b> to synchronize the timestamp type to the Hive metadata. The default value is <b>false</b> , indicating that the timestamp type is converted to bigint during synchronization by default. In this case, an error may occur when you query a Hudi table that contains a field of the timestamp type using SQL statements.	true

### 11.2.6.3 Index Configuration

Parameter	Description	Default Value
hoodie.index.class	Specifies the full path of user-defined index class. It must be a subclass of HoodieIndex and takes precedence over the <b>hoodie.index.type</b> configuration if specified.	""
hoodie.index.type	Specifies the type of index to be used. The Bloom filter is used by default. The possible option is <b>[BLOOM   HBASE   GLOBAL_BLOOM   SIMPLE   GLOBAL_SIMPLE]</b> . The Bloom filter eliminates the dependency on an external system and is stored in the footer of a Parquet data file.	BLOOM

Parameter	Description	Default Value
hoodie.index.bloom.num_entries	<p>Specifies the number of entries to be stored in the Bloom filter. If the <b>maxParquetFileSize</b> is 128 MB and <b>averageRecordSize</b> is 1,024 bytes, a total of 130 KB records exist in a file. The default value (60000) is about half of this approximate value.</p> <p><b>CAUTION</b> If the value is too low, a large number of false positives will occur, and index lookup will have to scan more files than it needs. If the value is too high, the size every data file will be increased linearly (about 4 KB for every 50000 entries).</p>	60000
hoodie.index.bloom.fpp	<p>Specifies the allowed error rate based on the number of entries. This is used to calculate how many bits should be assigned for the Bloom filter and the number of hash functions. Generally, it is set to a small value (0.000000001 by default). You need to balance the disk space to reduce the false positive rate.</p>	0.000000001
hoodie.bloom.index.parallelism	<p>Specifies the parallelism for index lookup, which involves Spark Shuffle. By default, this is automatically computed based on input workload characteristics.</p>	0
hoodie.bloom.index.pruned.by.ranges	<p>If this parameter is set to <b>true</b>, information is ranged from files to speed up index lookups. It is particularly useful if the key has a monotonously increasing prefix, such as timestamp.</p>	true
hoodie.bloom.index.use.caching	<p>If this parameter is set to <b>true</b>, the input RDD is cached to speed up index lookups by reducing I/O for computing parallelism or affected partitions.</p>	true

Parameter	Description	Default Value
hoodie.bloom.index.use.treebased.filter	If this parameter is set to <b>true</b> , the file pruning optimization based on interval tree is enabled. This mode speeds up file pruning based on key ranges compared with the brute-force mode.	true
hoodie.bloom.index.bucketized.checking	If this parameter is set to <b>true</b> , the bucketized Bloom filtering is enabled. This reduces skew in the sort-based Bloom index lookup.	true
hoodie.bloom.index.keys.per.bucket	This parameter is applicable only when <b>bloomIndexBucketizedChecking</b> is enabled and the index type is <b>BLOOM</b> . This parameter controls the "bucket" size which tracks the number of record-key checks made against a single file and is the unit of work allocated to each partition performing the Bloom filter lookup. A higher value would amortize the fixed cost of reading the Bloom filter to memory.	10000000
hoodie.bloom.index.update.partition.path	This parameter is applicable only when the index type is <b>GLOBAL_BLOOM</b> . If this parameter is set to <b>true</b> , an update including the partition path of a record that already exists will result in the insertion of the incoming record into the new partition and the deletion of the original record in the old partition. If this parameter is set to <b>false</b> , the original record will only be updated in the old partition.	true
hoodie.index.hbase.zk.quorum	Mandatory. This parameter is available only when the index type is <b>HBASE</b> . HBase ZooKeeper quorum URL to be connected.	None

Parameter	Description	Default Value
hoodie.index.hbase.zk port	Mandatory. This parameter is available only when the index type is <b>HBASE</b> . HBase ZooKeeper quorum port to be connected.	None
hoodie.index.hbase.zk node.path	Mandatory. This parameter is available only when the index type is <b>HBASE</b> . It is the root znode that will contain all the znodes created and used by HBase.	None
hoodie.index.hbase.table	Mandatory. This parameter is available only when the index type is <b>HBASE</b> . HBase table name to be used as an index. Hudi stores the <b>row_key</b> and <b>[partition_path, fileID, commitTime]</b> mapping in the table.	None

#### 11.2.6.4 Storage Configuration

Parameter	Description	Default Value
hoodie.parquet.max.file.size	Specifies the target size for Parquet files generated in Hudi write phases. For DFS, this parameter needs to be aligned with the underlying file system block size for optimal performance.	120 x 1024 x 1024 bytes
hoodie.parquet.block.size	Specifies the Parquet page size. Page is the unit of read in a Parquet file. In a block, pages are compressed separately.	120 x 1024 x 1024 bytes
hoodie.parquet.compression.ratio	Specifies the expected compression ratio of Parquet data when Hudi attempts to adjust the size of a new Parquet file. If the size of the file generated by <b>bulk insert</b> is smaller than the expected size, increase the value.	0.1

Parameter	Description	Default Value
hoodie.parquet.compression.codec	Specifies the name of the Parquet compression encoding or decoding mode. The default value is <b>gzip</b> . Possible options are <b>[gzip   snappy   uncompressed   lzo]</b> .	snappy
hoodie.logfile.max.size	Specifies the maximum size of LogFile. It is the maximum size allowed for a log file before it is rolled over to the next version.	1GB
hoodie.logfile.data.block.max.size	Specifies the maximum size of a LogFile data block. It is the maximum size allowed for a single data block to be appended to a log file. It helps to ensure that the data appended to the log file is broken up into sizable blocks to prevent OOM errors. The size should be greater than the JVM memory.	256MB
hoodie.logfile.to.parquet.compression.ratio	Specifies the expected additional compression when records move from log files to Parquet files. It is used for MOR tables to send inserted content into log files and control the size of compacted Parquet files.	0.35

### 11.2.6.5 Compaction and Cleaning Configurations

Parameter	Description	Default Value
hoodie.clean.automati c	Specifies whether to perform automatic cleanup.	true

Parameter	Description	Default Value
hoodie.cleaner.policy	Specifies the cleaning policy to be used. Hudi will delete the Parquet file of an old version to reclaim space. Any query or computation referring to this version of the file will fail. You are advised to ensure that the data retention time exceeds the maximum query execution time.	KEEP_LATEST_COMMITS
hoodie.cleaner.commits.retained	Specifies the number of commits to retain. Data will be retained for <b>num_of_commits * time_between_commits</b> (scheduled). This also directly translates into the number of datasets can be incrementally pulled.	10
hoodie.keep.max.commits	Number of commits that triggers the archiving operation.	30
hoodie.keep.min.commits,	Number of commits reserved by the archiving operation.	20
hoodie.commits.archival.batch	This parameter controls the number of commit instants read in memory as a batch and archived together.	10
hoodie.parquet.small.file.limit	The value must be smaller than that of <b>maxFileSize</b> . If <b>maxFileSize</b> is set to <b>0</b> , this function is disabled. Small files always exist because of the large number of insert records in a partition of batch processing. Hudi provides an option to solve the problem of small files by masking inserts into this partition as updates to existing small files. The size here is the minimum file size that is considered as a "small file size".	104857600 byte

Parameter	Description	Default Value
hoodie.copyonwrite.insert.split.size	Specifies the parallelism for inserting and writing data. It is the number of inserts grouped for a single partition. Writing out 100 MB files with at least 1 KB records means 100 KB records exist in each file. Overprovision to 500 KB by default. To improve insert latency, adjust the value to match the number of records in a single file. If it is set to a smaller value, the file size will shrink (especially when <b>compactionSmallFileSize</b> is set to 0).	500000
hoodie.copyonwrite.insert.auto.split	Specifies whether Hudi dynamically computes <b>insertSplitSize</b> based on the last 24 commit metadata. This function is disabled by default.	true
hoodie.copyonwrite.record.size.estimate	Specifies the average record size. If specified, Hudi will use this parameter and not compute dynamically based on the last 24 commit metadata. There is no default value. This is critical in computing the insert parallelism and packing inserts into small files.	1024
hoodie.compact.inline	If this parameter is set to <b>true</b> , compaction is triggered by the ingestion itself right after a commit or delta commit action as part of <b>insert</b> , <b>upsert</b> , or <b>bulk_insert</b> .	true
hoodie.compact.inline.max.delta.commits	Specifies the maximum number of delta commits to be retained before inline compression is triggered.	5

Parameter	Description	Default Value
hoodie.compaction.lazy.block.read	When <b>CompactedLogScanner</b> merges all log files, this parameter helps to choose whether the logblocks should be read lazily. Set it to <b>true</b> to use I/O-intensive lazy block read (low memory usage) or <b>false</b> to use memory-intensive immediate block read (high memory usage).	true
hoodie.compaction.reverse.log.read	<b>HoodieLogFormatReader</b> reads a log file in the forward direction from <b>pos=0</b> to <b>pos=file_length</b> . If this parameter is set to <b>true</b> , Reader reads a log file in reverse direction from <b>pos=file_length</b> to <b>pos=0</b> .	false
hoodie.cleaner.parallelism	Increase this parameter if cleaning becomes slow.	200
hoodie.compaction.strategy	Determines which file groups are selected for compaction during each compaction run. By default, Hudi selects the log file with most accumulated unmerged data.	org.apache.hudi.table.action.compact.strategy. LogFileSizeBasedCompactionStrategy
hoodie.compaction.target.io	Specifies the number of MBs to spend during compaction run for <b>LogFileSizeBasedCompactionStrategy</b> . This parameter can limit ingestion latency when compaction is run in inline mode.	500 * 1024 MB
hoodie.compaction.daybased.target.partitions	Used by <b>org.apache.hudi.io.compact.strategy.DayBasedCompactionStrategy</b> to denote the number of latest partitions to compact during a compaction run.	10



Parameter	Description	Default Value
hoodie.compaction.payload.class	It needs to be same as class used during insert or upsert. Similar to writing, compaction also uses the record payload class to merge records in the log against each other, merge again with the base file, and produce the final record to be written after compaction.	org.apache.hudi.common.model.Defaulthoodierecordpayload
hoodie.schedule.compact.only.inline	Specifies whether to generate only a compression plan during a write operation. This parameter is valid only when <b>hoodie.compact.inline</b> is set to <b>true</b> .	false
hoodie.run.compact.only.inline	Specifies whether to perform only the compression operation when the <b>run compact</b> command is executed using SQL. If the compression plan does not exist, no action is needed.	false

### 11.2.6.6 Single-Table Concurrency Control Configuration

Parameter	Description	Default Value
hoodie.write.lock.provider	Specifies the lock provider. You are advised to set the parameter to <b>org.apache.hudi.hive.HiveMetastoreBasedLockProvider</b> .	org.apache.hudi.client.transaction.lock.ZookeeperBasedLockProvider
hoodie.write.lock.hive.metastore.database	Specifies the Hive database.	None
hoodie.write.lock.hive.metastore.table	Specifies the Hive table name.	None
hoodie.write.lock.client.num_retries	Specifies the retry times.	10
hoodie.write.lock.client.wait_time_ms_between_retry	Specifies the retry interval.	10000

Parameter	Description	Default Value
hoodie.write.lock.conflict.resolution.strategy	Specifies the lock provider class, which must be a subclass of <b>ConflictResolutionStrategy</b> .	org.apache.hudi.client.transaction.SimpleConcurrentFileWritesConflictResolutionStrategy
hoodie.write.lock.zookeeper.base_path	Path for storing ZNodes. The parameter must be the same for all concurrent write configurations of the same table.	None
hoodie.write.lock.zookeeper.lock_key	ZNode name. It is recommended that the ZNode name be the same as the Hudi table name.	None
hoodie.write.lock.zookeeper.connection_timeout_ms	ZooKeeper connection timeout period.	15000
hoodie.write.lock.zookeeper.port	ZooKeeper port number.	None
hoodie.write.lock.zookeeper.url	URL of the ZooKeeper.	None
hoodie.write.lock.zookeeper.session_timeout_ms	Session expiration time of ZooKeeper.	60000

## 11.3 Hudi Performance Tuning

### Performance Tuning Methods

In the current version, Spark is recommended for Hudi write operations. Therefore, the tuning methods of Hudi are similar to those of Spark. For details, see [Spark2x Performance Tuning](#).

### Recommended Resource Configuration

- For MOR tables:

The essence of MOR tables is to write incremental files, so the tuning is based on the data size (dataSize) of Hudi.

If dataSize is only several GBs, you are advised to run Spark in single-node mode or run Spark in Yarn mode with only one container allocated.

Parallelism (**p**) of programs for importing data to the lake:  $p = \text{dataSize} / 128 \text{ MB}$ . The number of cores allocated to programs must be the same as the value of **p**. It is recommended that the ratio of the memory size to the number of cores be greater than 1.5:1. That is, a core is configured with 1.5

GB memory. For off-heap memory, it is recommended that the ratio of the memory size to the number of cores be greater than 0.5:1.

- For COW tables:

The principle of COW tables is to rewrite the original data. Therefore, dataSize and the number of rewritten files must be considered during tuning. Typically, more cores lead to better performance. The number of cores is directly related to the number of rewritten files. The settings of parallelism (**p**) and memory size are similar to those of MOR tables.

## 11.4 Common Issues About Hudi

### 11.4.1 Data Write

#### 11.4.1.1 Parquet/Avro schema Is Reported When Updated Data Is Written

##### Question

The following error is reported when data is written:

```
org.apache.parquet.io.InvalidRecordException: Parquet/Avro schema mismatch: Avro field 'col1' not found
```

##### Answer

You are advised to evolve schemas in backward compatible mode while using Hudi. This error usually occurs when you delete some columns, such as **col1**, in backward incompatible mode and then update **col1** written with the old schema in the Parquet file. In this case, the Parquet file attempts to search for all the current fields in the input record, if **col1** does not exist, the preceding exception is thrown.

To solve this problem, create an uber schema using all the schema versions evolved and use this uber schema as the target schema. You can obtain a schema from Hive MetaStore and merge it with the current schema.

#### 11.4.1.2 UnsupportedOperationException Is Reported When Updated Data Is Written

##### Question

The following error is reported when data is written:

```
java.lang.UnsupportedOperationException: org.apache.parquet.avro.AvroConverters$FieldIntegerConverter
```

##### Answer

This error will occur again because schema evolutions are in non-backwards compatible mode. Basically, there is some update U for a record R which is already written to the Hudi dataset in the Parquet file. R contains field F which includes certain data type, that is long. U has the same field F with the int data type. Parquet FS does not support incompatible data type conversions.

For such errors, perform valid data type conversions in the data source where you collect data.

### 11.4.1.3 SchemaCompatibilityException Is Reported When Updated Data Is Written

#### Question

The following error is reported when data is written:

```
org.apache.hudi.exception.SchemaCompatibilityException: Unable to validate the rewritten record <record>
against schema <schema>at
org.apache.hudi.common.util.HoodieAvroUtils.rewrite(HoodieAvroUtils.java:215)
```

#### Answer

This error may occur if a schema contains some **non-nullable** field whose value is not present or is null.

You are advised to evolve schemas in backward compatible mode. Essentially, this means either you need to set each newly added field to null or to default values. In Hudi 0.5.1 and later versions, the troubleshooting is invalid if fields rely on default values.

### 11.4.1.4 What Should I Do If Hudi Consumes Much Space in a Temporary Folder During Upsert?

#### Question

Hudi consumes much space in a temporary folder during upsert.

#### Answer

Hudi will spill part of input data to disk if the maximum memory for merge is reached when much input data is upserted.

If the memory is sufficient, increase the memory of the Spark executor and add the **hoodie.memory.merge.fraction** option, for example, **option("hoodie.memory.merge.fraction", "0.8")**.

### 11.4.1.5 Hudi Fails to Write Decimal Data with Lower Precision

#### Question

Decimal data is initially written to a Hudi table using the **BULK\_INSERT** command. Then when data is subsequently written using **UPSERT**, the following error is reported:

```
java.lang.UnsupportedOperationException: org.apache.parquet.avro.AvroConverters$FieldFixedConverter
```

#### Answer

**Cause:**

The Hudi table contains decimal data.

The initial bulk insert of data is implemented using the Spark class for writing Parquet files. However, Spark processes the decimal data with different precisions differently.

When data is written using the **UPSERT** command, Hudi uses the Avro-compliant class for writing Parquet files, which is incompatible with the Spark class.

#### Solutions:

When executing the **BULK\_INSERT** command, set **hoodie.datasource.write.row.writer.enable** to **false** to enable Hoodie to use the Avro-compliant class for writing Parquet files.

## 11.4.2 Data Collection

### 11.4.2.1 IllegalArgumentException Is Reported When Kafka Is Used to Collect Data

#### Question

The error "org.apache.kafka.common.KafkaException: Failed to construct kafka consumer" is reported in the **main** thread, and the following error is reported.

```
java.lang.IllegalArgumentException: Could not find a 'KafkaClient' entry in the JAAS configuration. System property 'java.security.auth.login.config' is not set
```

#### Answer

This error may occur when you try to collect data from the Kafka source with SSL enabled and the installation program cannot read the **jaas.conf** file and its properties.

To solve this problem, pass the required property as part of the command submitted through Spark. Example: **--files jaas.conf,failed\_tables.json --conf 'spark.driver.extraJavaOptions=-Djava.security.auth.login.config=jaas.conf' --conf 'spark.executor.extraJavaOptions=-Djava.security.auth.login.config=jaas.conf'**

### 11.4.2.2 HoodieException Is Reported When Data Is Collected

#### Question

The following error is reported when data is collected:

```
com.uber.hoodie.exception.HoodieException: created_at(Part -created_at) field not found in record. Acceptable fields were :[col1, col2, col3, id, name, dob, created_at, updated_at]
```

#### Answer

This error usually occurs when a field marked as recordKey or partitionKey is not present in the input record. Cross verify the input record.

### 11.4.2.3 HoodieKeyException Is Reported When Data Is Collected

#### Question

Is it possible to use a nullable field that contains null records as a primary key when creating a Hudi table?

#### Answer

No. HoodieKeyException will be thrown.

```
Caused by: org.apache.hudi.exception.HoodieKeyException: recordKey value: "null" for field: "name" cannot be null or empty.
at org.apache.hudi.keygen.SimpleKeyGenerator.getKey(SimpleKeyGenerator.java:58)
at org.apache.hudi.HoodieSparkSqlWriter$$anonfun$1.apply(HoodieSparkSqlWriter.scala:104)
at org.apache.hudi.HoodieSparkSqlWriter$$anonfun$1.apply(HoodieSparkSqlWriter.scala:100)
```

## 11.4.3 Hive Synchronization

### 11.4.3.1 SQLException Is Reported During Hive Data Synchronization

#### Question

The following error is reported during Hive data synchronization:

```
Caused by: java.sql.SQLException: Error while processing statement: FAILED: Execution Error, return code 1 from org.apache.hadoop.hive.ql.exec.DDLTask. Unable to alter table. The following columns have types incompatible with the existing columns in their respective positions :
__col1,__col2
```

#### Answer

This error usually occurs when you try to add a new column to an existing Hive table using the **HiveSyncTool.java** class. Databases usually do not allow the modification of a column data type from a higher order to lower order or cases where the data types may conflict with the data that is already stored or will be stored in the table. To solve this problem,

set **hive.metastore.disallow.incompatible.col.type.changes** to **false**.

### 11.4.3.2 HoodieHiveSyncException Is Reported During Hive Data Synchronization

#### Question

The following error is reported during Hive data synchronization:

```
com.uber.hoodie.hive.HoodieHiveSyncException: Could not convert field Type from <type1> to <type2> for field col1
```

#### Answer

This error occurs because HiveSyncTool currently supports only few compatible data type conversions. The exception is thrown if any other incompatible changes are made.

Check the data type evolution for the related field and verify if it indeed can be considered as a valid data type conversion based on the Hudi code base.

### 11.4.3.3 SemanticException Is Reported During Hive Data Synchronization

#### Question

The following error is reported during Hive data synchronization:

```
org.apache.hadoop.hive.ql.parse.SemanticException: Database does not exist: test_db
```

#### Answer

This error typically occurs when Hive synchronization is performed on the Hudi data set but the configured **hive\_sync** database does not exist.

Create the corresponding database on your Hive cluster and try again.

# 12 Using Hue (Versions Earlier Than MRS 3.x)

---

## 12.1 Using Hue from Scratch

Hue provides the file browser function using a graphical user interface (GUI) so that you can view files and directories on Hive.

### Prerequisites

You have installed Hive and Hue, and the Kerberos authentication cluster in the running state.

### Procedure

**Step 1** Access the Hue web UI.

**Step 2** Open the Hue web UI and choose **Query Editors > Hive**.


**Step 3** In **Databases**, select a Hive database, the default database is **default**.


The system displays all available tables. You can enter a keyword of the table name to search for the desired table.

**Step 4** Click the desired table name. All columns in the table are displayed.

**Step 5** Enter the HiveQL statements in the area for editing.


```
create table hue_table(id int,name string,company string) row format delimited fields terminated by ',' stored as textfile;
```

Click  and select **Explain**. The editor checks the syntax and execution plan of the entered HiveQL statements. If the statements have syntax errors, the editor reports **Error while compiling statement**.

**Step 6** Click , and select the engine for executing the HiveQL statements.

**Step 7** Click  to execute the HiveQL statements.



**Step 8** In the command text box, enter **show tables;** and click . Check whether the **hue-table** table created in [Step 5](#) exists in the result.

----End

## 12.2 Accessing the Hue Web UI

### Scenario

After Hue is installed in an MRS cluster, users can use Hadoop and Hive on the Hue web UI.

This section describes how to open the Hue web UI on the MRS cluster.

#### NOTE

To access the Hue web UI, you are advised to use a browser that is compatible with the Hue WebUI, for example, Google Chrome 50. The Internet Explorer may be incompatible with the Hue web UI.

### Impact on the System

Site trust must be added to the browser when you access Manager and Hue web UI for the first time. Otherwise, the Hue web UI cannot be accessed.

### Prerequisites

When Kerberos authentication is enabled, you have been assigned the permission to access Hive by the MRS cluster administrator. For example, create a human-machine user named **hueuser**, add the user to user groups **hive** (the primary group), **hadoop**, and **supergroup**, and role **System\_administrator**.

This user is used to log in to the Hue WebUI.

### Procedure

**Step 1** Log in to the service page, click the cluster name to go to the cluster details page, and choose **Components**.



#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

**Step 2** Select **Hue**. On the right side of **Hue WebUI**, click the link to log in to the Hue web UI as user **hueuser**.

Hue WebUI provides the following functions:

- If Hive is installed in the MRS cluster, you can use **Query Editors** to execute query statements of Hive. Hive has been installed in the MRS cluster.
- If Hive is installed in the MRS cluster, you can use **Data Browsers** to manage Hive tables.

- If HDFS is installed in the MRS cluster, you can use  to view directories and files in HDFS.
- If Yarn is installed in the MRS cluster, you can use  to view all jobs in the MRS cluster.

 **NOTE**

- When you log in to the Hue web UI as user **hueuser** for the first time, you need to change the password.
- After obtaining the URL for accessing the Hue web UI, you can give the URL to other users who cannot access MRS Manager for accessing the Hue web UI.
- If you perform operations on the Hue WebUI only but not on Manager, you must enter the password of the current login user when accessing Manager again.

----End

## 12.3 Hue Common Parameters

### Navigation Path

For details about how to set parameters, see [Modifying Cluster Service Configuration Parameters](#).

### Parameters

**Table 12-1** Hue common parameters

Parameter	Description	Default Value	Value Range
HANDLER_ACCESSLOG_LEVEL	Hue access log level	DEBUG	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>
HANDLER_AUDITLOG_LEVEL	Hue audit log level	DEBUG	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>
HANDLER_ERRORLOG_LEVEL	Hue error log level	ERROR	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>

Parameter	Description	Default Value	Value Range
HANDLER_LOGFILE_LEVEL	Hue run log level	INFO	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>
HANDLER_LOGFILE_MAXBACKUPINDEX	Maximum number of Hue log files.	20	1 to 999
HANDLER_LOGFILE_SIZE	Maximum size of a Hue log file.	5 MB	-

## 12.4 Using HiveQL Editor on the Hue Web UI

### Scenario


Users can use the Hue web UI to execute HiveQL statements in a cluster.

### Accessing Query Editors

**Step 1** Access the Hue web UI.

**Step 2** Choose **Query Editors** > **Hive**. The **Hive** page is displayed.

**Hive** supports the following functions:

- Executes and manages HiveQL statements.
- View the HiveQL statements saved by the current user in **Saved Queries**.
- Query HiveQL statements executed by the current user in **Query History**.
- Click  to display all databases included in **Databases** of Hive.

----End


### Executing HiveQL Statements

**Step 1** Choose **Query Editors** > **Hive**. The **Hive** page is displayed.


**Step 2** Click  and select a database from **Databases**. The default database is **default**.


The system displays all available tables in the database. You can enter a keyword of the table name to search for the desired table.

**Step 3** Click the desired table name. All columns in the table are displayed.

Move the cursor to the row of the table and click . Column details are displayed.

**Step 4** Enter the query statements in the area for editing HiveQL statements.

Click  and select **Explain**. The editor checks the syntax and execution plan of the entered statements. If the statements have syntax errors, the editor reports **Error while compiling statement**.

**Step 5** Click  and select the engine for executing the HiveQL statements.







- **mr**: MapReduce computing framework
- **spark**: Spark computing framework
- **tez**: Tez computing framework

 **NOTE**

Tez is applicable to MRS 1.9.x and later versions.

**Step 6** Click  to execute the HiveQL statements.

 **NOTE**

- If you want to use the entered HiveQL statements again, click  to save them.
- To format HiveQL statements, click  and select **Format**.
- To delete an entered HiveQL statement, click  and select **Clear**.
- Clear the entered statement and execute a new statement. Click  and select **New query**.
- Viewing history:  
Click **Query History** to view the HiveQL running status. You can view the history of all the statements or only the saved statements. If many historical records exist, you can enter keywords in the text box to search for desired records.
- Advanced query configuration:  
Click  in the upper right corner to configure information such as files, functions, and settings.
- Viewing the information of shortcut keys:  
Click  in the upper right corner to view all shortcut keys.

----End

## Viewing Execution Results

**Step 1** In the **Hive** execution area, **Query History** is displayed by default.

**Step 2** Click **Results** to view the execution result of the executed statement.

----End

## Managing Query Statements

**Step 1** Choose **Query Editors > Hive**. The **Hive** page is displayed.

**Step 2** Click **Saved Queries**.


Click a saved statement. The system automatically adds the statement to the editing area.


----End


## Modifying Query Editors Settings


**Step 1** On the **Hive** tab page, click .

**Step 2** Click  next to **Files** and click  to specify the directory for storing the file.

You can click  to add a file resource.

**Step 3** Click  next to **Functions** and enter the names of user-defined function and function class.

You can click  to add a customized function.

**Step 4** Click  next to **Settings**, enter the Hive parameter name in **Key**, and the value in **Value**. The current Hive session connects to Hive based on the customized configuration.

You can click  to add a parameter.

----End

## 12.5 Using the Metadata Browser on the Hue Web UI

### Scenario


Users can use the Hue web UI to manage Hive metadata in an MRS cluster.

### Using Metastore Manager




Access the Hue web UI.

Choose **Data Browsers > Metastore Tables**, and access **Metastore Manager**.

- Viewing metadata of Hive tables

In the left navigation pane, move the cursor to a table and click  on the right. The metadata of the Hive table is displayed.



- Managing metadata of Hive tables

On the metadata page of a Hive table, you can click  in the upper right corner to import data, click  to browse data, and click  to view the location of the table file.

 **CAUTION**

The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

- Managing Hive metadata tables

Click  in the upper right corner to create a table in the database based on the uploaded files. Or click  in the upper right corner to manually create a table.

## Accessing Metastore Manager

**Step 1** Access the Hue web UI.

**Step 2** Choose **Data Browsers > Metastore Tables**, and access **Metastore Manager**.

**Metastore Manager** supports the following functions:

- Creating a Hive table from a file
- Manually creating a Hive table
- Viewing Hive table metadata

----End


## Creating a Hive table from a File

**Step 1** Access **Metastore Manager** and select a database in **Databases**.

The default database is **default**.

**Step 2** Click . The **Create a new table from a file** page is displayed.

**Step 3** Select a file.

1. In **Table Name**, enter a Hive table name.  
A Hive table name contains no more than 128 characters, including letters, numbers, or underscores (\_), and must start with a letter or number.
2. In **Description**, enter description about the Hive table as required.
3. In **Input File or Location**, click  and select a Hive table file from HDFS. The file is used to store new data of the Hive table.  
If the file is not stored in HDFS, click **Upload a file** to upload the file from the local directory to HDFS. Multiple files can be simultaneously uploaded. The files cannot be empty.
4. If you need to import the data in the file to the Hive table, select **Import data as Load method**. By default, **Import data** is selected.

If you select **Create External Table**, a Hive external table is created.

 **NOTE**

If you select **Create External Table**, set **Input File or Location** to a path.

If you select **Leave Empty**, an empty Hive table is created.

5. Click **Next**.

**Step 4** Set a delimiter.


1. In **Delimiter**, select one.

If your desired delimiter is not in the list, select **Other..** and enter a delimiter.

2. Click **Preview** to preview data processing.

3. Click **Next**.


**Step 5** Define a column.

1. If you click  on the right side of **Use first row as column names**, the first row of data in the file is used as a column name. If you do not click it, the first row of data is not used as the column name.

2. In **Column name**, set a name for each column.

A Hive table name contains no more than 128 characters, including letters, numbers, or underscores (\_), and must start with a letter or number.

 **NOTE**

You can rename columns in batches by clicking  on the right side of **Bulk edit column names**. Enter all column names and separate them by commas (,).

3. In **Column Type**, select a type for each column.

**Step 6** Click **Create Table** to create the table. Wait for Hue to display information about the Hive table.

----End

## Manually Creating a Hive Table

**Step 1** Access **Metastore Manager** and select a database in **Databases**.

The default database is **default**.

**Step 2** Click . The **Create a new table manually** page is displayed.

**Step 3** Set a table name.

1. In **Table Name**, enter a Hive table name.

A Hive table name contains no more than 128 characters, including letters, numbers, or underscores (\_), and must start with a letter or number.

2. In **Description**, enter description about the Hive table as required.

3. Click **Next**.

**Step 4** Select a data storage format.

- If data needs to be separated by delimiters, select **Delimited** and perform [Step 5](#).

- If data needs to be stored in serialization format, select **SerDe** and perform [Step 6](#).

**Step 5** Set a delimiter.

1. In **Field terminator**, set a column delimiter.  
If your desired delimiter is not in the list, select **Other..** and enter a delimiter.
2. In **Collection terminator**, set a delimiter to separate the data set of columns of the **array** type in Hive. For example, the type of a column is array. A value needs to store **employee** and **manager**. The user specifies a colon (:) as the delimiter. Therefore, the final value is **employee:manager**.
3. In **Map key terminator**, set a delimiter to separate the data set of columns of the **map** type in Hive. For example, the type of a column is map. A value needs to store **home** of **aaa** and **company** of **bbb**. The user defines | as the delimiter. Therefore, the final value is **home|aaa:company|bbb**.
4. Click **Next** and perform [Step 7](#).

**Step 6** Set serialization properties.

1. In **SerDe Name**, enter the class name of the serialization format:  
**org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe**  
Users can expand Hive to support more customized serialization classes.
2. In **Serde properties**, enter the value of the serialization format:  
**"field.delim"="," "collection.delim"=":" "mapkey.delim"="|"**
3. Click **Next** and perform [Step 7](#).

**Step 7** Select a data table format and click **Next**.

- **TextFile**: indicates that data is stored in text files.
- **SequenceFile**: indicates that data is stored in binary files.
- **InputFormat**: indicates that data in files is used in the customized input and output formats.

Users can expand Hive to support more customized formatting classes.

- a. In **InputFormat Class**, enter the class used by input data:  
**org.apache.hadoop.hive.ql.io.RCFileInputFormat**
- b. In **OutputFormat Class**, enter the class used by output data:  
**org.apache.hadoop.hive.ql.io.RCFileOutputFormat**

**Step 8** Select a file storage location and click **Next**.

**Use default location** is selected by default. If you want to customize a storage location, deselect the default value and specify a file storage location in **External location** by clicking **\*\***.

**Step 9** Set columns of the Hive table.

1. In **Column name**, set a column name.  
A Hive table name contains no more than 128 characters, including letters, numbers, or underscores (\_), and must start with a letter or number.
2. In **Column type**, select a type for each column.  
Click **Add a column** to add a new column.



3. Click **Add a partition** to add a new partition for the Hive table to improve the query efficiency.

**Step 10** Click **Create Table** to create a new table. Wait for Hue to display information about the Hive table.

----End

## Managing the Hive Table

**Step 1** Access **Metastore Manager** and select a database in **Databases**. All tables in the database are displayed on the page.

The default database is **default**.

**Step 2** Click a table name in the database to view table details.

The following operations are supported: importing data, browsing data,, or viewing file storage location. When viewing all tables in the database, you can select tables and perform the following operations such as viewing tables and browsing data.

---

 **CAUTION**

The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

---

----End

## 12.6 Using File Browser on the Hue Web UI

### Scenario

Users can use the Hue web UI to manage files in HDFS in a cluster.

---

 **CAUTION**

The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

---

## File Browser (File Browser)

**Step 1** Access the Hue web UI.

**Step 2** Click . The **File Browser** page is displayed.

You can view the home directory of the current login user.

On the **File Browser** page, the following information about subdirectories for files in the directory is displayed.

**Table 12-2** HDFS file attributes


Attribute	Description
<b>Name</b>	Name of a directory or file
<b>Size</b>	File size
<b>User</b>	Owner of a directory or file
<b>Group</b>	Group of a directory or file
<b>Permissions</b>	Permission of a directory or file
<b>Date</b>	Time when a directory or file is created

**Step 3** In the search box, enter a keyword. The system automatically searches directories or files in the current directory.

**Step 4** Clear the search criteria. The system displays all directories or files.

----End

## Performing Actions

**Step 1** Click  and select one or more directories or files.

**Step 2** Click **Actions**. On the menu that is displayed, select an operation.

- **Rename**: renames a directory or file.
- **Move**: moves a file. In **Move to**, select a new directory and click **Move**.
- **Copy**: copies the selected files or directories.
- **Change permissions**: changes permission to access the selected directory or file.
  - You can grant the owner, the group, or other users with the **Read**, **Write**, and **Execute** permissions.
  - **Sticky**: indicates that only HDFS administrators, directory owners, and file owners can move files in the directory.
  - **Recursive**: indicates that permission is granted to subdirectories recursively.
- **Storage policies**: indicates the policies for storing files or directories in HDFS.

- **Summary:** indicates that you can view HDFS storage information about the selected file or directory.

----End

## Accessing Other Directories

**Step 1** Click the directory name, type a full path you want to access, for example, **/mr-history/tmp**, and press **Enter**.

The current user must have permission to access other directories.

**Step 2** Click **Home** to go to the home directory.

**Step 3** Click **History**. The history records of directory access are displayed and the directories can be accessed again.

**Step 4** Click **Trash** to access the recycle bin of the current directory.

Click **Empty Trash** to clean up the recycle bin.

----End

## Uploading User Files

**Step 1** Click  and click **Upload**.

**Step 2** Select an operation.

- **Files:** uploads user files to the current user.
- **Zip/Tgz/Bz2 file:** uploads a compressed file. In the dialog box that is displayed, click **Select ZIP, TGZ or BZ2 files** to select the compressed file to be uploaded. The system automatically decompresses the file in HDFS. Compressed files in **ZIP, TGZ, and BZ2** formats are supported.

----End

## Creating a New File or Directory

**Step 1** Click  and click **New**.

**Step 2** Select an operation.

- **File:** creates a file. Enter a file name and click **Create**.
- **Directory:** creates a directory. Enter a directory name and click **Create**.

----End


## Storage Policy Definition and Usage

### NOTE

If the value of Hue parameter **fs\_defaultFS** is set to **viewfs://ClusterX**, the big data storage policy cannot be enabled.

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

**Step 2** On MRS Manager, choose **System > Permission > Manage Role > Create Role**.

1. Set **Role Name**.
  2. Choose **Configure Resource Permission > Hue**, select **Storage Policy Admin**, and click **OK** to grant the storage policy administrator permission to the role.
- Step 3** Choose **System > Permission > Manage User Group > Create User Group**, set **Group Name**, and click **Select and Add Role** next to **Role**. On the displayed page, select the created role and click **OK** to add the role to the group.
- Step 4** Choose **System > Permission > Manage User > Create User**.
1. Specify the **Username** of a user who can log in to the Hue web UI and has the **Storage Policy Admin** permission.
  2. Set **User Type** to **Human-machine**.
  3. Set **Password** and **Confirm Password** for logging in to the Hue web UI.
  4. Click **Select and Join User Group** next to **User Group**. On the page that is displayed, select the created user group, **supergroup**, **hadoop**, and **hive**, and click **OK**.
  5. Set **Primary Group** to **hive**.
  6. Click **Select and Add Role** on the right of **Assign Rights by Role**. On the Select snf page that is displayed, select the newly created role and the **System\_administrator** role, and click **OK**.
  7. Click **OK**. The user is added successfully.
- Step 5** Access the Hue web UI.
- Step 6** Click  in the upper right corner.
- Step 7** Select the check box of the directory and click **Action** on the upper part of the page. Then select **Storage policies**.
- Step 8** In the dialog box that is displayed, set a new storage policy and click **OK**.
- End

## 12.7 Using Job Browser on the Hue Web UI

### Scenario

You can use the Hue web UI to query all jobs in the cluster.

### Accessing Job Browser

**Step 1** Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).

**Step 2** Click **Job Browser**.


View the jobs in the cluster.

#### NOTE

The number on **Job Browser** indicates the total number of jobs in the cluster.

**Job Browser** displays the following job information.

**Table 12-3** MRS job attributes

Attribute	Description
Logs	Log information. If a job has logs,  is displayed.
ID	Job ID, which is generated by the system automatically.
Name	Job name
Application Type	Job type
Status	Job status. Possible values are <b>RUNNING</b> , <b>SUCCEEDED</b> , <b>FAILED</b> , and <b>KILLED</b> .
User	User who starts the job
Maps	Map progress
Reduces	Reduce progress
Queue	Yarn queue used for job running
Priority	Job running priority
Duration	Job running duration
Submitted	Time when the job is submitted to the MRS cluster

 **NOTE**

If the MRS cluster has Spark, the **Spark-JDBCServer** job is started by default to execute tasks.

----End

## Searching for Jobs

**Step 1** Enter keywords in **Username** or **Text** on the **Job Browser** page to search for the desired jobs.

**Step 2** Clear the search criteria. The system displays all jobs.


----End

## Querying Job Details

**Step 1** In the job list on the **Job Browser** page, click the row that contains the desired job to view details.

**Step 2** On the **Metadata** tab page, you can view the metadata of the job.

 **NOTE**

You can click  to open job running logs.

----End

# 13 Using Hue (MRS 3.x or Later)

---

## 13.1 Using Hue from Scratch


Hue aggregates interfaces which interact with most Apache Hadoop components and enables you to use Hadoop components with ease on a web UI. You can operate components such as HDFS, Hive, HBase, Yarn, MapReduce, Oozie, and Spark SQL on the Hue web UI.

### Prerequisites

You have installed Hue, and the Kerberos authentication cluster is in the running state.

### Procedure

**Step 1** Access the Hue web UI.

**Step 2** In the navigation tree on the left, click the editor icon  and choose **Hive**.

**Step 3** Select a Hive database from the **Database** drop-down list box. The default database is **default**.


The system displays all available tables. You can enter a keyword of the table name to search for the desired table.

**Step 4** Click the desired table name. All columns in the table are displayed.

**Step 5** Enter the HiveQL statements in the area for editing.

```
create table hue_table(id int,name string,company string) row format
delimited fields terminated by ',' stored as textfile;
```

**Step 6** Click  to execute the HiveQL statements.

**Step 7** In the command text box, enter **show tables;** and click . Check whether the **hue\_table** table created in [Step 5](#) exists in the **Result**.

----End

## 13.2 Accessing the Hue Web UI

### Scenario

After Hue is installed in an MRS cluster, users can use Hadoop-related components on the Hue web UI.

This section describes how to open the Hue web UI on the MRS cluster.

#### NOTE

To access the Hue web UI, you are advised to use a browser that is compatible with the Hue WebUI, for example, Google Chrome 50. The Internet Explorer may be incompatible with the Hue web UI.

### Impact on the System

Site trust must be added to the browser when you access Manager and Hue web UI for the first time. Otherwise, the Hue web UI cannot be accessed.

### Prerequisites

When Kerberos authentication is enabled, you have been assigned the permission to access Hive by the MRS cluster administrator. For example, create a human-machine user **hueuser**, add the user to groups **hive** (primary group), **hadoop**, and **supergroup**, and assign the **System\_administrator** role.

This user is used to log in to Manager.

### Procedure





**Step 1** Log in to the service page.





For versions earlier than MRS 3.x, click the cluster name on the MRS console and choose **Components > Hue**.

For MRS 3.x or later, log in to FusionInsight Manager and choose **Cluster > Services > Hue**.

**Step 2** On the right of **Hue WebUI**, click the link to open the Hue web UI.

Hue WebUI provides the following functions:

- Click  to execute query statements of Hive and SparkSQL as well as Notebook code. Make sure that Hive and Spark2x have been installed in the MRS cluster before this operation.
- Click  to submit workflow tasks, scheduled tasks, and bundle tasks.
- Click  to view, import, and export tasks on the Hue web UI, such as workflow tasks, scheduled tasks, and bundle tasks.
- Click  to manage metadata in Hive and SparkSQL. Make sure that Hive and Spark2x have been installed in the MRS cluster before this operation.

- Click  to view the directories and files in HDFS. Make sure that HDFS has been installed in the MRS cluster before this operation.
- Click  to view all jobs in the MRS cluster. Make sure that Yarn has been installed in the MRS cluster before this operation.
- Use  to create or query HBase tables. Make sure that the HBase component has been installed in the MRS cluster and the Thrift1Server instance has been added before this operation.
- Use  to import data that is in the CSV or TXT format.

 **NOTE**

- When you log in to the Hue web UI as user **hueuser** for the first time, you need to change the password.
- After obtaining the URL for accessing the Hue web UI, you can give the URL to other users who cannot access MRS Manager for accessing the Hue web UI.
- If you perform operations on the Hue WebUI only but not on Manager, you must enter the password of the current login user when accessing Manager again.

----End

## 13.3 Hue Common Parameters

### Page Access

Go to the **All Configurations** page of the Hue service by referring to [Modifying Cluster Service Configuration Parameters](#).

### Parameter Description

For details about Hue common parameters, see [Table 13-1](#).

**Table 13-1** Hue common parameters

Configuration	Description	Default Value	Value Range
HANDLER_ACCESSLOG_LEVEL	Hue access log level.	DEBUG	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>
HANDLER_AUDITLOG_LEVEL	Hue audit log level.	DEBUG	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>



Configuration	Description	Default Value	Value Range
HANDLER_ERRORLOG_LEVEL	Hue error log level.	ERROR	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>
HANDLER_LOGFILE_LEVEL	Hue run log level.	INFO	<ul style="list-style-type: none"> <li>• ERROR</li> <li>• WARN</li> <li>• INFO</li> <li>• DEBUG</li> </ul>
HANDLER_LOGFILE_MAXBACKUPINDEX	Maximum number of Hue log files.	20	1 to 999
HANDLER_LOGFILE_SIZE	Maximum size of a Hue log file.	5 MB	-


## 13.4 Using HiveQL Editor on the Hue Web UI

### Scenario

Users can use the Hue web UI to execute HiveQL statements in an MRS cluster.

### Access Editor

**Step 1** Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).

**Step 2** In the navigation tree on the left, click  and choose **Hive**. The **Hive** page is displayed.

**Hive** supports the following functions:

- Executes and manages HiveQL statements.
- Views the HiveQL statements saved by the current user in **Saved Queries**.
- Queries HiveQL statements executed by the current user in **Query History**.


----End

### Executing HiveQL Statements

**Step 1** Select a Hive database from the **Database** drop-down list box. The default database is **default**.

The system displays all available tables. You can enter a keyword of the table name to search for the desired table.





**Step 2** Click the desired table name. All columns in the table are displayed.

Move the cursor to the row where the table or column is located and click . Column details are displayed.

**Step 3** Enter the query statements in the area for editing HiveQL statements.

**Step 4** Click  to execute the HiveQL statements.

 **NOTE**

- If you want to use the entered HiveQL statements again, click  to save them.
- Advanced query configuration:  
Click  in the upper right corner to configure information such as files, functions, and settings.
- Viewing the information of shortcut keys:  
Click  in the upper right corner to view the syntax and keyboard shortcut information.
- To delete an entered HiveQL statement, click the triangle next to  and select **Clear**.
- Viewing history:  
Click **Query History** to view the HiveQL running status. You can view the history of all the statements or only the saved statements. If many historical records exist, you can enter keywords in the text box to search for desired records.

----End

## Viewing Execution Results

**Step 1** View the execution results below the execution area on **Hive**. The **Query History** tab page is displayed by default.

**Step 2** Click a result to view the execution result of the executed statement.

----End

## Managing Query Statements

**Step 1** Click **Saved Queries**.

**Step 2** Click a saved statement. The system automatically adds the statement to the editing area.


----End

## Modifying the Session Configuration of the Hue Editor


**Step 1** On the editor page, click .


**Step 2** Click  on the right of **Files**, and then click  to select files.

You can click  next to **Files** to add a file resource.

**Step 3** In the **Functions**  area, enter a user-defined name and the class name of the function.

You can click  next to **Functions** to add a customized function.

**Step 4** In the **Settings**  area, enter the Hive parameter name in the **Key**, and value in **Value**. The current Hive session connects to Hive based on the customized configuration.

You can click  to add a parameter.

----End

## 13.5 Using the SparkSql Editor on the Hue Web UI

### Scenario

You can use Hue to execute SparkSql statements in a cluster on a graphical user interface (GUI).

### Configuring Spark2x

Before using the SparkSql editor, you need to modify the Spark2x configuration.

**Step 1** Go to the Spark2x configuration page. For details, see [Modifying Cluster Service Configuration Parameters](#).

**Step 2** Set the Spark2x multi-instance mode. Search for and modify the following parameters of the Spark2x service:

Parameter	Value
spark.thriftserver.proxy.enabled	false
spark.scheduler.allocation.file	#{conf_dir}/fairscheduler.xml

**Step 3** Go to the JDBCServer2x customization page and add the following customized items to the **spark.core-site.customized.configs** parameter:

**Table 13-2** Custom parameters


Parameter	Value
hadoop.proxyuser.hue.groups	*
hadoop.proxyuser.hue.hosts	*

**Step 4** Save the configuration and restart the meta and Spark2x services.

----End

## Accessing the Editor

**Step 1** Access the Hue web UI.

**Step 2** In the navigation tree on the left, click  and choose **SparkSql**. The **SparkSql** page is displayed.

**SparkSql** supports the following functions:

- Executes and manages SparkSql statements.
- Views the SparkSql statements saved by the current user in **Saved Queries**.
- Queries SparkSql statements executed by the current user in **Query History**.


----End

## Executing SparkSql Statements


**Step 1** Select a SparkSql database from the **Database** drop-down list box. The default database is **default**.

The system displays all available tables. You can enter a keyword of the table name to search for the desired table.

**Step 2** Click the desired table name. All columns in the table are displayed.


Move the cursor to the row of the table and click . Column details are displayed.


**Step 3** In the SparkSql statement editing area, enter the query statement.

Click the triangle next to  and select **Explain**. The editor checks the syntax and execution plan of the entered statements. If the statements have syntax errors, the editor reports **Error while compiling statement**.


**Step 4** Click  to execute the SparkSql statement.



### NOTE

- If you want to use the entered SparkSql statements again, click  to save them.
- Advanced query configuration:

Click  in the upper right corner to configure information such as files, functions, and settings.

- Viewing the information of shortcut keys:

Click  in the upper right corner to view the syntax and keyboard shortcut information.

- To format the SparkSql statement, click the triangle next to  and select **Format**.
- To delete an entered SparkSql statement, click the triangle next to  and select **Clear**.
- Viewing historical records:

Click **Query History** to view the SparkSql running status. You can view the history of all the statements or only the saved statements. If many historical records exist, you can enter keywords in the text box to search for desired records.

----End

## Viewing Execution Results

**Step 1** View the execution results below the execution area on **SparkSql**. The **Query History** tab page is displayed by default.

**Step 2** Click a result to view the execution result of the executed statement.

----End

## Managing Query Statements

**Step 1** Click **Saved Queries**.

**Step 2** Click a saved statement. The system automatically adds the statement to the editing area.

----End

# 13.6 Using the Metadata Browser on the Hue Web UI

## Scenario


Users can use the Hue web UI to manage Hive metadata in an MRS cluster.

## Using Metadata Manager

Access the Hue web UI.

- Viewing metadata of Hive tables



Click  in the navigation tree on the left and click a table name. The metadata of the Hive table is displayed.

- Managing metadata of Hive tables


On the metadata information page of a Hive table:

- Click **Import** in the upper right corner to import data.
- Click **Overview** to view the location of the table file in the **PROPERTIES** field.

View the field information of each column in a Hive table and manually add description information. Note that the added description information is not the field comments in the Hive table.

- Click **Sample** to browse data.

- Managing Hive metadata tables

Click  in the left list to create a table based on the uploaded file in the database. You can also manually create a table.

 CAUTION

The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

## 13.7 Using File Browser on the Hue Web UI

### Scenario

Users can use the Hue web UI to manage files in HDFS in a cluster.

 CAUTION

The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

### Accessing File Browser

**Step 1** Access the Hue web UI.

**Step 2** In the left navigation pane, click . The **File Browser** page is displayed.

By default, the homepage of **File Browser** is the home directory of the current login user. On the displayed page, the following information about subdirectories for files in the directory is displayed:

**Table 13-3** HDFS file attributes

Attribute	Description
Name	Name of a directory or file
Size	File size
User	Owner of a directory or file
Group	Group of a directory or file
Permission	Permission of a directory or file
Date	Time when a directory or file is created

**Step 3** In the search box, enter a keyword. The system automatically searches directories or files in the current directory.

**Step 4** Clear the search criteria. The system displays all directories or files.

----End

## Performing Actions

**Step 1** On the **File Browser** page, select one or more directories or files.

**Step 2** Click **Actions**. On the menu that is displayed, select an operation.

- **Rename**: renames a directory or file.
- **Move**: moves a file. In **Move to**, select a new directory and click **Move**.
- **Copy**: copies the selected files or directories.
- **Change permissions**: changes permission to access the selected directory or file.
  - You can grant the owner, the group, or other users with the **Read**, **Write**, and **Execute** permissions.
  - **Sticky**: indicates that only HDFS administrators, directory owners, and file owners can move files in the directory.
  - **Recursive**: indicates that permission is granted to subdirectories recursively.
- **Storage policies**: indicates the policies for storing files or directories in HDFS.
- **Summary**: indicates that the HDFS storage information about the selected file or directory can be viewed.

----End

## Uploading User Files

**Step 1** On the **File Browser** page, click **Upload**.

**Step 2** In the displayed dialog box for uploading files, click **Select files** or drag the file to the dialog box.

----End

## Creating a New File or Directory

**Step 1** On the **File Browser** page, click **New**.

**Step 2** Select an operation.


- **File**: creates a file. Enter a file name and click **Create**.
- **Directory**: creates a directory. Enter a directory name and click **Create**.

----End

## Storage Policy Definition and Usage

### NOTE

If the value of Hue parameter **fs\_defaultFS** is set to **viewfs://ClusterX**, the big data storage policy cannot be enabled.

- Step 1** Log in to FusionInsight Manager.
- Step 2** On FusionInsight Manager, choose **System > Permission > Manage Role > Create Role**.
1. Set **Role Name**.
  2. In the **Configure Resource Permission** area, choose *Name of the desired cluster* > **Hue**, select **Storage Policy Admin**, and click **OK**. Then, grant the permission to the role.
- Step 3** Choose **System > Permission > User Group > Create User Group**. Set **Group Name** and click **Select and Add Role** next to **Role**. On the displayed page, select the role created in **Step 2** and click **OK** to add the role to the group.
- Step 4** Choose **System > Permission > User > Create**.
1. **Username**: Enter the name of the user to be added.
  2. Set **User Type** to **Human-machine**.
  3. Set **Password** and **Confirm Password** for logging in to the Hue web UI.
  4. Click **Add** next to **User Group**. On the page that is displayed, select the user group created in **Step 3**, **supergroup**, **hadoop**, and **hive**, and click **OK**.
  5. Set **Primary Group** to **hive**.
  6. Click **Add** on the right of **Role**. On the page that is displayed, select the role created in **Step 2** and **System\_administrator** role, and click **OK**.
  7. Click **OK**. The user is added successfully.
- Step 5** Access the Hue web UI as the created user.
- Step 6** In the left navigation tree, click . The **File Browser** page is displayed.
- Step 7** Select the check box of the directory and click **Actions** on the top of the page. Choose **Storage policies**.
- Step 8** In the dialog box that is displayed, set a new storage policy and click **OK**.
- End

## 13.8 Using Job Browser on the Hue Web UI

### Scenario

Users can use the Hue web UI to query all jobs in an MRS cluster.



## Accessing Job Browser

**Step 1** Access the Hue web UI.

**Step 2** Click .

View the jobs in the current cluster.

 **NOTE**

The number on **Job Browser** indicates the total number of jobs in the cluster.

**Job Browser** displays the following job information:

**Table 13-4** MRS job attributes

Attribute	Description
Name	Job name
User	User who starts a job
Type	Job type
Status	Job status, including <b>Succeeded</b> , <b>Running</b> , and <b>Failed</b> .
Progress	Job running progress
Group	Group to which a job belongs
Start	Start time of a job
Duration	Job running duration
Id	Job ID, which is generated by the system automatically.

 **NOTE**

If the MRS cluster has Spark, the **Spark-JDBCServer** job is started by default to execute tasks.

----End

## Searching for Jobs

**Step 1** In the search box of **Job Browser**, enter the specified character. The system automatically searches for all jobs that contain the keyword by ID, name, or user.

**Step 2** Clear the search criteria. The system displays all jobs.

----End

## Querying Job Details

**Step 1** In the job list on the **Job Browser** page, click the row that contains the desired job to view details.

**Step 2** On the **Metadata** tab page, you can view the metadata of the job.

 **NOTE**

You can click **Log** to open the job running log.

----End

## 13.9 Using HBase on the Hue Web UI

### Scenario

You can use Hue to create or query HBase tables in a cluster and run tasks on the Hue web UI.

Make sure that the HBase component has been installed in the MRS cluster and the Thrift1Server instance has been added before this operation.

### Accessing Job Browser

**Step 1** Access the Hue web UI.

**Step 2** Click HBase . The **HBase Browser** page is displayed.

----End

### Creating an HBase Table

**Step 1** Access the Hue web UI.

**Step 2** Click HBase . The **HBase Browser** page is displayed.

**Step 3** Click **New Table** on the right, enter the table name and column family parameters, and click **Submit**.

----End

### Querying Data in an HBase Table

**Step 1** Access the Hue web UI.

**Step 2** Click HBase . The **HBase Browser** page is displayed.

**Step 3** Click the HBase table to be queried. Then, click the key value next to search box in the upper part, and query the HBase table.

----End

## 13.10 Typical Scenarios

## 13.10.1 HDFS on Hue

Hue provides the file browser function for users to use HDFS in GUI mode.

---


### CAUTION

The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

---

### How to Use File Browser

Access the Hue web UI.

Click . The **File Browser** page is displayed. You can perform the following operations:

- Viewing files or directories  
By default, the directory and files in the directory of the login user are displayed. You can view **Name**, **Size**, **User**, **Group**, **Permission**, and **Date**.  
Click a file name to view the text information or binary data in the text file. The file content can be edited.  
If there are a large number of files and directories, you can enter keywords in the search text box to search for specific files or directories.
- Creating files or directories  
Click **New** in the upper right corner. Choose **File** to create the file. Choose **Directory** to create a directory.
- Managing files or directories  
Select the check box of a file or director, and click **Actions**. In the displayed menu, choose **Rename**, **Move**, **Copy**, and **Change permissions** to rename, move, copy, or change the file or directory permissions.
- Uploading files  
Click **Upload** in the upper right corner and click **Select files** or drag the file to the window.

### How to Use Storage Policies

#### NOTE

If the value of Hue parameter **fs\_defaultFS** is set to **viewfs://ClusterX**, the big data storage policy cannot be enabled.

**Storage policies on the Hue web UI are classified into the following two types:**

- Static Storage Policies

Current storage policy

According to the access frequency and importance of documents in HDFS, specify a storage policy for an HDFS directory, such as ONE\_SSD or ALL\_SSD. The files in this directory can be migrated to the storage media.

- Dynamic Storage Policies

Set rules for an HDFS directory. The system can automatically change the storage policy, change the number of file copies, or move the file directory.

Before configuring a dynamic storage policy on the Hue WebUI, you must set the CRON expressions for cold and hot data migration and start automatic cold and hot data migration on Manager.

Operations:

Modify the following NameNode parameters of HDFS. For details, see [Modifying Cluster Service Configuration Parameters](#).

Parameter	Description	Example Value
dfs.auto.data.mover.enable	Whether to enable automatic hot and cold data migration. The default value is <b>false</b> .	true
dfs.auto.data.mover.cron.expression	CRON expression for hot and cold data migration in HDFS, which is used to control the start time of data migration. This parameter is available only when <b>dfs.auto.data.mover.enable</b> is set to <b>true</b> . The default value is <b>0 * * * *</b> , indicating that the task is executed on the hour.	0 * * * *

[Table 13-5](#) describes the expression for modifying the **dfs.auto.data.mover.cron.expression** parameter. \* indicates consecutive time segments.

**Table 13-5** Parameters in the execution expression

Column	Description
1	Minute. The value ranges from 0 to 59.
2	Hour. The value ranges from 0 to 23.
3	Date. The value ranges from 1 to 31.
4	Month. The value ranges from 1 to 12.
5	Week. The value ranges from 0 to 6. <b>0</b> indicates Sunday.

**To set storage policies on the web UI, perform the following operations:**

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

**Step 2** On FusionInsight Manager, choose **System > Permission > Role > Create Role**.

1. Set **Role Name**.
2. In the **Configure Resource Permission** area, choose *Name of the desired cluster* > **Hue**, select **Storage Policy Admin**, and click **OK**. Then, grant the permission to the role.

**Step 3** Choose **System > Permission > User Group > Create User Group**. Set **Group Name**, and click **Add** next to **Role**. On the displayed page, select the created role, click **OK** to add the role to the group, and click **OK**.

**Step 4** Choose **System > Permission > User > Create**.

1. **Username**: Enter the name of the user to be added.
2. Set **User Type** to **Human-machine**.
3. Set **Password** and **Confirm Password** for logging in to the Hue web UI.
4. Click **Add** next to **User Group**. On the page that is displayed, select the created user group in [Step 3](#), **supergroup**, **hadoop**, and **hive**, and click **OK**.
5. Set **Primary Group** to **hive**.
6. Click **Add** next to **Role**. On the page that is displayed, select the created role in [Step 2](#) and the **System\_administrator** role, and click **OK**.
7. Click **OK**. The user is added successfully.

**Step 5** Access the Hue web UI as the created user.

**Step 6** In the left navigation pane, click . The **File Browser** page is displayed.

**Step 7** Select the check box of a directory and choose **Action** on the top of the page. Choose **Storage policies**.

**Step 8** In the dialog box that is displayed, set a new storage policy and click **OK**.

- On the **Static Storage Policy** page, you can set a static storage policy and click **Save**.
- On the **Dynamic Storage Policy** page, you can create, delete, or modify a dynamic storage policy. [Table 13-6](#) describes the parameters.

**Table 13-6** Parameters of the dynamic storage policy

Category	Parameter	Description
Rule	Last Access to File	Indicates the time when the file is last accessed.
	Last File Modification	Indicates the time when the file is last modified.
Operation	Change Number of Copies	Indicates the number of file copies.

Category	Parameter	Description
	Modify Storage Policy	Indicates that you can modify storage policies to the following: HOT, WARM, COLD, ONE_SSD, and ALL_SSD.
	Move to Directory	Indicates that you can move the file to another directory.

 **NOTE**

- You need to consider whether the rules conflict with each other and whether the rules damage the system when setting rules.
- When a directory is configured with multiple rules and operations, the rule that is triggered first is located at the bottom of the rule/operation list, and the rules that are triggered later are placed from bottom to top to prevent repeated operations.
- The system checks whether the files under the directory specified by the dynamic storage policy meet the rules on an hourly basis. If the files meet the rules, the execution is triggered. Execution logs are recorded in the `/var/log/Bigdata/hdfs/nn/hadoop.log` directory of the active NameNode.

----End

## Typical Scenarios

On the Hue page, view and edit HDFS files in text or binary mode as follows:

### Viewing a File

**Step 1** Access the Hue web UI.

**Step 2** In the left navigation pane, click . The **File Browser** page is displayed.

**Step 3** Click the name of the file to be viewed.

**Step 4** Click **View as binary** to switch from the text mode to the binary mode. Click **View as file** to switch from the binary mode to the text mode.

### Editing a file

**Step 5** Click **Edit File**. The file content can be edited.

**Step 6** Click **Save** or **Save As** to save the file.

----End

## 13.10.2 Hive on Hue


Hue provides the Hive GUI management function so that users can query Hive data in GUI mode.


### How to Use Query Editor

Access the Hue web UI.


In the navigation tree on the left, click  and choose **Hive**. The **Hive** page is displayed.

- Running Hive HQL statements


Select the target database on the left. You can also click **default**  in the upper right corner and enter the target database name to search for the target database.

Enter a Hive HQL statement in the text box and click  or press **Ctrl+Enter** to run the HQL statement. The execution result is displayed on the **Result** tab page.

- Analyzing Hive HQL statements

Select the target database on the left, enter the Hive HQL statement in the text box, and click  to compile the HQL statement and check whether the statement is correct. The execution result is displayed under the text editing box.

- Saving HQL statements

Enter the Hive HQL statement in the text box, click  in the upper right corner, and enter the name and description. You can view the saved statements on the **Saved Queries** tab page.

- Viewing historical records

Click **Query History** to view the HQL running status. You can view the history of all the statements or only the saved statements. If many historical records exist, you can enter keywords in the text box to search for desired records.

- Configuring advanced query

Click  in the upper right corner to configure the file, function, and settings.


- Viewing the information of shortcut keys

Click  in the upper right corner to view information about all shortcut keys.

## How to Use Metadata Browser

Access the Hue web UI.

- Viewing metadata of Hive tables

Click  in the navigation tree on the left and click a table name. The metadata of the Hive table is displayed.


- Managing metadata of Hive tables

On the metadata information page of a Hive table:

- Click **Import** in the upper right corner to import data.
- Click **Overview** to view the location of the table file in the **PROPERTIES** field.

View the field information of each column in a Hive table and manually add description information. Note that the added description information is not the field comments in the Hive table.

- Click **Sample** to browse data.

- Managing Hive metadata tables  
Click  in the left list to create a table based on the uploaded file in the database. You can also manually create a table.

---


 **CAUTION**

The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

---


## Typical Scenarios

On the Hue page, create a Hive table as follows:

**Step 1** Click  at the upper left corner of Hue web UI and select the Hive instance to be operated to enter the Hive command execution page.


**Step 2** Enter an HQL statement in the command input box, for example:

```
create table hue_table(id int,name string,company string) row format
delimited fields terminated by ',' stored as textfile;
```

Click  to execute the HQL statements.

**Step 3** Enter the following command in the command input box:

```
show tables;
```

Click  to view the created table **hue\_table** in **Result**.

----End

### 13.10.3 Oozie on Hue

Hue provides the Oozie job manager function, in this case, you can use Oozie in GUI mode.

---

 **CAUTION**


The Hue page is used to view and analyze data such as files and tables. Do not perform high-risk management operations such as deleting objects on the page. If an operation is required, you are advised to perform the operation on each component after confirming that the operation has no impact on services. For example, you can use the HDFS client to perform operations on HDFS files and use the Hive client to perform operations on Hive tables.

---



## How to Use Oozie Job Designer

Access the Hue web UI.

In the navigation tree on the left, click  and choose **Workflow**.


The job designer allows users to create MapReduce, Java, Streaming, Fs, SSH, Shell and DistCp jobs.

## How to Use Dashboard

Access the Hue web UI.


Click **Jobs** in the upper right corner. The **Job Browser** page is displayed.

View the running status of the Workflow, Coordinator, and Bundles jobs.



## How to Use Editor

Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).

In the navigation tree on the left, click  and choose **Workflow**.

Workflows, Schedule, and Bundle tasks can be created. Existing applications can be submitted for running, shared, copied, and exported.

- Each Workflow can contain one or more jobs to form a complete workflow for a specified service.  
When creating a Workflow, you can design jobs in the Hue editor and add the jobs to the Workflow.
- Each Schedule can define a time trigger to periodically execute a specified Workflow. One time trigger cannot execute multiple Workflows.
- Each Bundles can define a set to execute multiple Schedules so that different Workflows can be executed in batches.

## 13.11 Hue Log Overview

### Log Description

**Log paths:** The default paths of Hue logs are `/var/log/Bigdata/hue` (for storing run logs) and `/var/log/Bigdata/audit/hue` (for storing audit logs).

**Log archive rules:** The automatic compression and archiving function of the Hue logs is enabled. By default, when the size of a log file (**access.log**, **error.log**, **runcpserver.log**, or **hue-audits.log**) exceeds 5 MB, logs are automatically compressed. A maximum of 20 latest compressed files are reserved. The number of compressed files and compression threshold can be configured.

**Table 13-7** Hue log list

Type	Log File Name	Description
Run log	access.log	Access log file
	error.log	Error log file
	gsdb_check.log	Log file of the GaussDB check information
	kt_renewer.log	Log file of Kerberos authentication
	kt_renewer.out.log	Log file of the abnormal Kerberos authentication logs
	runcpserver.log	Log file of operation records
	runcpserver.out.log	Log file of process running exceptions
	supervisor.log	Log file of process startup
	supervisor.out.log	Log file of process startup exceptions
	dbDetail.log	Log file of database initialization
	initSecurityDetail.log	Download initialization log file of the Keytab file
	postinstallDetail.log	Work log file generated after the Hue service is installed
	prestartDetail.log	Prestart log file
	statusDetail.log	Log file of the Hue health status
	startDetail.log	Startup log
	get-hue-ha.log	Log file of the Hue HA status
	hue-ha-status.log	Log file of the Hue HA status monitoring
	get-hue-health.log	Log file of the Hue health status
	hue-health-check.log	Log file of the Hue health check
	hue-refresh-config.log	Log file of the Hue configuration update
	hue-script-log.log	Log file of the Hue operations on the Manager console
hue-service-check.log	Log file of the Hue service status monitoring	

Type	Log File Name	Description
	db_pwd.log	Log that records the changes of the password for Hue to connect to the DBService database
	modifyDBPwd_Date.log	-
	watch_config_update.log	Parameter update log file
Audit log	hue-audits.log	Audit log file

## Log Level

**Table 13-8** describes the log levels supported by Hue.

Levels of logs are ERROR, WARN, INFO, and DEBUG from the highest to the lowest priority. Run logs of equal or higher levels are recorded. The higher the specified log level, the fewer the logs recorded.

**Table 13-8** Log levels

Level	Description
ERROR	Logs of this level record error information about system running.
WARN	Logs of this level record exception information about the current event processing.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of the Hue service by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** In the navigation tree on the left, select **Log** corresponding to the role to be modified.
- Step 3** Select the log level to be changed on the right.
- Step 4** Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.
- Step 5** Restart the service or instance whose configuration has expired for the configuration to take effect.

----End

## Log Format

The following table lists the Hue log formats:

**Table 13-9** Log formats

Type	Format	Example
Run log	<i>&lt;dd-MM-yy HH:mm:ss,SSS&gt;&lt;Location where the log event occurs&gt;&lt;Log level&gt;&lt;Message in the log&gt;</i>	[03/Nov/2014 11:57:19 ] middleware   INFO   Unloading MimeTypeJSFileFixStrea- mingMiddleware.
	<i>&lt;Log level&gt;&lt;Time format&gt;&lt;yyyy-MM-dd HH:mm:ss,SSS&gt;&lt;Location where the log event occurs&gt;&lt;Message in the log&gt;</i>	INFO : CST 2014-11-06 11:22:52 hue-ha- status.sh : update 4 <= 15:myHostName=10.0.0. 250 ACTIVE=10.0.0.250
Audit log	<i>&lt;UserName&gt;&lt;yyyy-MM- dd HH:mm:ss,SSS&gt;&lt; Audit operation description&gt; &lt;Resource parameter&gt; &lt;URL&gt; &lt;Whether to allow&gt; &lt;Audit operation&gt; &lt;IP address&gt;</i>	{"username": "admin", "eventTime": "2014-11-06 10:28:34", "operationText": "Successful login for user: admin", "service": "accounts", "url": "/" accounts/login/", "allowed": true, "operation": "USER_LOGIN", "ipAddress": "10.0.0.250"}

## 13.12 Common Issues About Hue

### 13.12.1 Why Do HQL Statements Fail to Execute in Hue Using Internet Explorer?

#### Question

What do I do if all HQL statements fail to be executed when I use Internet Explorer to access Hive Editor in Hue and the message "There was an error with your query" is displayed?

## Answer

Internet Explorer does not support processing of AJAX POST requests containing form data in 307 redirection. You are advised to use a compatible browser, for example, Google Chrome.

## 13.12.2 Why Does the use database Statement Become Invalid in Hive?

### Question

When Hive is used, the **use database** statement is entered in the text box to switch the database, and other statements are also entered, why does the database fail to be switched?


### Answer

Using Hive on Hue is different from using Hive on the Hive client. There is an option to select a database on the Hue interface, and the database where the current SQL is executed is the one that is displayed on the interface. You are advised to use functions on the Hue interface instead of using statements to perform session-level and one-off operations, for example, setting parameters. If you must enter specific statements to perform an operation, ensure that all statements you enter are in one text box.

## 13.12.3 Why Do HDFS Files Fail to Access Through the Hue Web UI?

### Question

What can I do if an error message shown in the following figure is displayed, indicating that the HDFS file cannot be accessed when I use Hue web UI to access the HDFS file?



### Answer

1. Check whether the user who logs in to the Hue web UI has the permissions of the **hadoop** user group.
2. Check whether the HttpFS instance has been installed for the HDFS service and is running properly. If the HttpFS instance is not installed, manually install and restart the Hue service.

## 13.12.4 Why Do Large Files Fail to Upload on the Hue Page

### Question

What can I do when a large file fails to be uploaded on the Hue page?

### Answer

1. You are advised to run commands on the client to upload large files instead of using the Hue file browser.
2. If you must use Hue to upload the file, perform the following steps to modify Httpd parameters:

- a. Log in to the active management node as user **omm**.
- b. Run the following command to edit the **httpd.conf** file:  
**vi \$BIGDATA\_HOME/om-server/Apache-httpd-\*/conf/httpd.conf**
- c. Search for **21201** and add **RequestReadTimeout handshake=0 header=0 body=0** to the **</VirtualHost>** configuration, as shown in the following:

```
...
<VirtualHost *:21201>
 ServerName https://10.112.16.93:21201
 AllowEncodedSlashes On
 SSLProxyEngine On
 ProxyRequests Off
 TraceEnable off
 ProxyTimeout 1200
 RewriteEngine on
 RewriteMap proxylist dbm:${BIGDATA_ROOT_HOME}/om-server_*/Apache-httpd-*/conf/
 proxylist.dbm

 RewriteRule ^(\/*)$ ${proxylist:/Hue/Hue/21201}$1 [E=TARGET_PATH:$1,L,P]

 Header edit Location ^(!https://10.112.16.93:20009|https://
10.112.16.93:21201)http[s]?://[^\/*]*$ https://10.112.16.93:21201$1

 ProxyPassReverseCookiePath / / interpolate

 SSLEngine On
 SSLProxyProtocol All +TLSv1.2 -SSLv2 -SSLv3 -TLSv1 -TLSv1.1
 SSLProtocol ALL +TLSv1.2 -SSLv2 -SSLv3 -TLSv1 -TLSv1.1
 SSLCipherSuite ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES256-GCM-
SHA384:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES128-GCM-SHA256:DHE-DSS-
AES256-GCM-SHA384:DHE-RSA-AES256-GCM-SHA384:DHE-DSS-AES128-GCM-SHA256:DHE-
RSA-AES128-GCM-SHA256
 SSLProxyCheckPeerName off
 SSLProxyCheckPeerCN off
 SSLCertificateFile "${BIGDATA_ROOT_HOME}/om-server_*/Apache-httpd-*/conf/security/
proxy_ssl.cert"
 SSLCertificateKeyFile "${BIGDATA_ROOT_HOME}/om-server_*/Apache-httpd-*/conf/security/
server.key"
 SSLProxyCACertificateFile ${BIGDATA_ROOT_HOME}/om-server_*/apache-tomcat-*/conf/
security/tomcat.crt
 SSLCertificateChainFile "${BIGDATA_ROOT_HOME}/om-server_*/Apache-httpd-2.4.39/conf/
security/proxy_chain.cert"
 RequestReadTimeout handshake=0 header=0 body=0
</VirtualHost>
...
```

- d. Run the **kill -9 httpd** command to stop the httpd process and wait for it to automatically restart.

## 13.12.5 Why Is the Hue Native Page Cannot Be Properly Displayed If the Hive Service Is Not Installed in a Cluster?

### Question

Why is the native Hue page blank if the Hive service is not installed in a cluster?

### Answer

In MRS 3.x, Hue depends on Hive. If this problem occurs, check whether the Hive component is installed in the current cluster. If not, install it.

# 14 Using Impala

## 14.1 Using Impala from Scratch

Impala is a massively parallel processing (MPP) SQL query engine for processing vast amounts of data stored in Hadoop clusters. It is an open source software written in C++ and Java. It provides high performance and low latency compared with other SQL engines for Hadoop.

### Background

Suppose a user develops an application to manage users who use service A in an enterprise. The procedure of operating service A on the Impala client is as follows:

#### Operations on common tables:

- Create the **user\_info** table.
- Add users' educational backgrounds and titles to the table.
- Query user names and addresses by user ID.
- Delete the user information table after service A ends.

**Table 14-1** User information

No.	Name	Gender	Age	Address
12005000201	A	Male	19	City A
12005000202	B	Female	23	City B
12005000203	C	Male	26	City C
12005000204	D	Male	18	City D
12005000205	E	Female	21	City E
12005000206	F	Male	32	City F
12005000207	G	Female	29	City G



No.	Name	Gender	Age	Address
12005000208	H	Female	30	City H
12005000209	I	Male	26	City I
12005000210	J	Female	25	City J

## Prerequisites

The client has been installed. For example, the client is installed in the `/opt/hadoopclient` directory. The client directory in the following operations is only an example. Change it to the actual installation directory.

## Procedure

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/hadoopclient
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Run the Impala client command to implement service A.

- **Operations on internal tables:**

Run the Impala client command **impala-shell**.

 **NOTE**

By default, **impala-shell** attempts to connect to the Impala daemon on port 21000 of **localhost**. To connect to another host, use the **-i <host:port>** option, for example, **impala-shell -i xxx.xxx.xxx.xxx:21000**. To automatically connect to a specific Impala database, use the **-d <database>** option. For example, if all your Kudu tables are in the **impala\_kudu** database, **-d impala\_kudu** can use this database. To exit Impala Shell, run the **quit** command.

a. Create the **user\_info** user information table according to [Table 14-1](#) and add data to it.

```
create table user_info(id string,name string,gender string,age int,addr string);
insert into table user_info(id,name,gender,age,addr) values("12005000201", "A", "Male", 19,
"City A");
```

... (Other statements are the same.)

b. Add users' educational backgrounds and titles to the **user\_info** table.

For example, to add educational background and title information about user 12005000201, run the following commands.

```
alter table user_info add columns(education string,technical string);
```

c. Query user names and addresses by user ID.

For example, to query the name and address of user 12005000201, run the following command:

```
select name,addr from user_info where id='12005000201';
```

- d. Delete the user information table:

```
drop table user_info;
```

- **Operations on external partition tables:**

Create an external partition table and import data.

- a. Create a path for storing external table data.

**kinit hive** (Run this command only in security mode.)

 **NOTE**

The user must have the hive administrator permissions.

**hdfs dfs -mkdir /hive**

**hdfs dfs -mkdir /hive/user\_info**

- b. Create a table.

**impala-shell**

 **NOTE**

By default, **impala-shell** attempts to connect to the Impala daemon on port 21000 of **localhost**. To connect to another host, use the **-i <host:port>** option, for example, **impala-shell -i xxx.xxx.xxx.xxx:21000**. To automatically connect to a specific Impala database, use the **-d <database>** option. For example, if all your Kudu tables are in the **impala\_kudu** database, **-d impala\_kudu** can use this database. To exit Impala Shell, run the **quit** command.

```
create external table user_info(id string,name string,gender string,age int,addr string)
partitioned by(year string) row format delimited fields terminated by ' ' lines terminated by '\n'
stored as textfile location '/hive/user_info';
```

 **NOTE**

- **fields terminated** indicates delimiters, for example, spaces.
- **lines terminated** indicates line breaks, for example, **\n**.
- **/hive/user\_info** indicates the path of the data file.

- c. Import data.

- i. Execute the **insert** statement to insert data.

```
insert into user_info partition(year="2018") values ("12005000201", "A", "Male", 19, "City A");
```

- ii. Run the **load data** command to import file data.

- 1) Create a file based on the data in [Table 14-1](#). For example, the file name is **txt.log**. Fields are separated by space, and the line feed characters are used as the line breaks.

- 2) Upload the file to HDFS.

**hdfs dfs -put txt.log /tmp**

- 3) Load data to the table.

**load data inpath '/tmp/txt.log' into table user\_info partition (year='2018');**

- d. Query the imported data:

```
select * from user_info;
```

- e. Delete the user information table:

```
drop table user_info;
```

----End

## 14.2 Common Impala Parameters

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

### Page Access

On Manager, choose **Cluster > Services > Impala**. On the page that is displayed, click the **Configuration** tab then the **All Configurations** sub-tab. Enter a parameter name in the search box.

### Parameter Description

 **NOTE**

The following table lists only some common parameters. The actual parameters are subject to the FusionInsight Manager page.

**Table 14-2** Common Impala parameters

Parameter	Description	Default Value	Value Range
impalad.customized.configs	Custom configuration item of the impalad process	N/A	N/A
--enable_ldap_auth	Whether to enable LDAP authentication	false	The value can be <b>true</b> or <b>false</b> .
--ldap_bind_pattern	ldap userDNPattern, for example, <b>cn=%s,ou=People,dc=,dc=com</b>	N/A	N/A
--ldap_passwords_in_clear_ok	If this parameter is set to <b>true</b> , LDAP passwords are sent in plaintext (excluding TLS and SSL) on the network.	false	The value can be <b>true</b> or <b>false</b> .
--ldap_uri-ip	ldap ip	N/A	N/A
--ldap_uri-port	ldap port	389	N/A
--max_log_files	Maximum number of process log files	10	N/A
--max_log_size	Maximum size of a process log file, in MB	200	N/A

Parameter	Description	Default Value	Value Range
statestored.customized.configs	Custom configuration item of the Statestored process	N/A	N/A
catalogd.customized.configs	Custom configuration item of the Catalogd process	N/A	N/A

## 14.3 Accessing the Impala Web UI

You can view Impala job information on the Impala web UI. Impala web UIs are classified into the following types based on instances:

- **StateStore WebUI:** used to manage nodes.
- **Catalog WebUI:** used to view metadata.
- **Impalad WebUI:** used to view details about each SQL statement.

### Prerequisites

Impala has been installed in a cluster.

### Accessing the StateStore Web UI

**Step 1** Access Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).

**Step 2** Choose **Services > Impala**.

**Step 3** In **StateStore WebUI** of **Impala Summary**, click **StateStore(Statestore)**. The StateStore web UI is displayed.

----End

### Accessing the Catalog Web UI

**Step 1** Access Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).

**Step 2** Choose **Services > Impala**.

**Step 3** In **Catalog WebUI** of **Impala Summary**, click **Catalog(Catalog)**. The Catalog web UI is displayed.

----End

### Accessing the Impalad Web UI

**Step 1** Access Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).

**Step 2** Choose **Services > Impala > Instance**.

**Step 3** Move the cursor to the Impalad instance in the **Role** column. The following link is displayed in the lower left corner of the page. Obtain the value after **null**, for example, **82** in this example.

**https://EIP:9022/mrsmanager/index.jsp?locale=zh-cn#/app/services/Impala/Impalad/null/82/EIP/STARTED/status/detail**

In the preceding command, **82** is an example. Change it based on the site requirements.

**Step 4** For details, see [Accessing the StateStore Web UI](#).

**Step 5** Change **StateStore/xx** in the URL of the StateStore web UI to **Impalad/xx** and access the new URL, where *xx* is the value obtained in **Step 3**.

----End

## 14.4 Using Impala to Operate Kudu

You can use the SQL statements of Impala to insert, query, update, and delete data in Kudu as an alternative to using Kudu APIs to build custom Kudu applications.

### Prerequisite

A complete cluster client has been installed. For example, the installation directory is **/opt/Bigdata/client**. The client directory in the following operations is only an example. Replace it with the actual installation directory.

### Impala on Kudu

**Step 1** Log in to the node where the client is installed.

**Step 2** Run the following command to initialize environment variables:

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

**Step 3** If Kerberos authentication is enabled for the cluster, perform the following operation to authenticate the user. If Kerberos authentication is not enabled for the cluster, skip this step.

```
kinit Service user
```

**Step 4** Run the following command to log in to the Impala client:

```
impala-shell
```

#### NOTE

By default, **impala-shell** attempts to connect to the Impala daemon on port 21000 of **localhost**. To connect to another host, use the **-i <host:port>** option. To automatically connect to a specific Impala database, use the **-d <database>** option. For example, if all your Kudu tables are in the **impala\_kudu** database, **-d impala\_kudu** can use this database. To exit the Impala shell, run the **quit** command.

**Step 5** Run the following commands to create an Impala table and import the prepared data, for example, data in the **/tmp/data10** directory:

```
create table dataorigin (name string,age string,pt string, date_p date) row
format delimited fields terminated by ',' stored as textfile;
```

```
load data inpath '/tmp/data10' overwrite into table dataorigin;
```

- Step 6** Run the following command to create a Kudu table. In the command, **kudu.master\_addresses** indicates the IP address of the KuduMaster instance. Set it to the actual IP address.

```
create table dataorigin2 (name string,age string,pt string, date_p date,
primary key(name)) stored as kudu
TBLPROPERTIES('kudu.master_addresses'='192.168.190.164:7051,192.168.204.1
78:7051,192.168.244.63:7051');
```

- Step 7** Perform the following operations on the Kudu table.

1. Insert data.  
**insert into dataorigin2 select \* from dataorigin;**
  2. Update data.  
**UPDATE dataorigin2 SET date\_p="2021-03-31" where age="73";**
  3. Upsert rows.  
**UPSERT INTO dataorigin2 VALUES ("spjted","75","28","2021-03-32");**  
**UPSERT INTO dataorigin2 VALUES ("kwhakb","92","29","2021-03-33");**  
**UPSERT INTO dataorigin2 VALUES ("oftrkf","13","30","2021-03-34");**  
**UPSERT INTO dataorigin2 VALUES ("kiewti","36","31","2021-03-35");**  
**UPSERT INTO dataorigin2 VALUES ("rknmql","98","32","2021-03-36");**  
**UPSERT INTO dataorigin2 VALUES ("fwcoij","52","33","2021-03-37");**  
**UPSERT INTO dataorigin2 VALUES ("pgvpdo","37","34","2021-03-35");**
  4. Delete a row.  
**DELETE FROM dataorigin2 WHERE date\_p="2021-03-31";**
- End

## 14.5 Interconnecting Impala with External LDAP

This section applies to MRS 3.1.0 or later.

- Step 1** Log in to Manager.
- Step 2** On Manager, choose **Cluster > Name of the desired cluster > Services > Impala > Configurations > All Configurations > Impalad(Role) > LDAP**.
- Step 3** Set the following parameters.

**Table 14-3** Parameter configuration

Parameter	Description	Remarks
--enable_ldap_auth	Whether to enable LDAP authentication	Value: <b>true</b> or <b>false</b>

Parameter	Description	Remarks
--ldap_bind_pattern	LDAP user DN pattern	Example: <b>cn=#UID,ou=People,dc= ,dc=com</b> or <b>cn=%s,ou=People,dc= ,dc=com</b>
--ldap_passwords_in_clear_ok	Whether the LDAP password is sent in plaintext	If this parameter is set to <b>true</b> , the LDAP password can be sent in plaintext. Value: <b>true</b> or <b>false</b> <b>NOTE</b> If <b>--enable_ldap_auth</b> is set to <b>true</b> , the LDAP TLS protocol is disabled by default during authentication. Therefore, you need to set <b>--ldap_passwords_in_clear_ok</b> to <b>true</b> . Otherwise, the Impala role will fail to be started. To enable the Ldap TLS protocol, set <b>--ldap_tls</b> to <b>true</b> in the customized configuration of the Impala role. After the configuration, the password can be sent in ciphertext.
--ldap_uri-ip	LDAP IP address	-
--ldap_uri-port	LDAP port number	Default value: <b>389</b>

**Step 4** After the modification, click **Save** in the upper left corner. In the displayed dialog box, click **OK**.

**Step 5** Choose **Cluster > Name of the desired cluster > Services > Impala > Instance**. On the displayed page, select the instances whose **Configuration Status** is **Expired**, choose **More > Restart Instance**, and restart the instance.

----End

## 14.6 Enabling and Configuring a Dynamic Resource Pool for Impala

This section describes how to enable and configure a dynamic resource pool to control Impala concurrency.

### Background

Use a dynamic resource pool to control Impala concurrency.

### Pool Config

Property	Value
Max memory (cluster wide)	1048576
Max concurrent queries	-1
Max queue size	200
Queue Timeout (ms)	60000
Min Query MEM_LIMIT range	0
Max Query MEM_LIMIT range	0
Clamp MEM_LIMIT query option	true

1. Log in to the master1 node of the cluster, switch to user **omm**, and create the **fair-scheduler.xml** and **llama-site.xml** files in the **/home/omm** directory.

```
[omm@node-master1IoKo impala]$ ll
total 16
-rw-----. 1 omm wheel 708 May 11 23:40 fair-scheduler.xml
-rw-----. 1 omm wheel 1062 May 11 23:53 llama-site.xml
-rw-----. 1 omm wheel 1118 May 11 23:12 llama-site.xml.bak
-rw-----. 1 omm wheel 572 May 11 23:32 update_config.sh
[omm@node-master1IoKo impala]$
```

2. Open the **fair-scheduler.xml** file and add the following configurations:

```
<allocations>
 <queue name="root">
 <aclSubmitApps> </aclSubmitApps>
 <queue name="default">
 <maxResources> 1mb, 0 vcores</maxResources><!--This parameter is for reference only-->
 <aclSubmitApps>*</aclSubmitApps>
 </queue>
 <queue name="development">
 <maxResources>2048 mb, 0 vcores</maxResources><!--This parameter is for reference only-->
 <aclSubmitApps>admin</aclSubmitApps>
 </queue>
 <queue name="production">
 <maxResources>7168 mb, 0 vcores</maxResources><!--This parameter is for reference only-->
 <aclSubmitApps>omm</aclSubmitApps>
 </queue>
 </queue>
 <queuePlacementPolicy>
 <rule name="specified" create="false"/>
 <rule name="default" />
 </queuePlacementPolicy>
</allocations>
```

3. Open the **llama-site.xml** file and add the following configurations:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
 <property>
 <name>llama.am.throttling.maximum.placed.reservations.root.default</name>
 <value>1</value>
 </property>
 <property>
 <name>llama.am.throttling.maximum.queued.reservations.root.default</name>
 <value>2</value><!--This parameter is for reference only-->
 </property>
 <property>
 <name>impala.admission-control.pool-default-query-options.root.default</name>
 <value>mem_limit=128m,query_timeout_s=20,max_io_buffers=10</value>
 </property>
 <property>
 <name>impala.admission-control.pool-queue-timeout-ms.root.default</name>
 <value>30000</value><!--This parameter is for reference only-->
 </property>
</configuration>
```



```
<property>
 <name>impala.admission-control.max-query-mem-limit.root.default</name>
 <value>1024000000</value><!--1GB--><!--This parameter is for reference only.-->
</property>
<property>
 <name>impala.admission-control.min-query-mem-limit.root.default</name>
 <value>2048000000</value><!--2GB-->
</property>
<property>
 <name>impala.admission-control.clamp-mem-limit-query-option.root.default.regularPool</name>
 <value>true</value>
</property>
</configuration>
```

- Run the following commands to synchronize **fair-scheduler.xml** and **llama-site.xml** to the **etc** folder in the installation directory on all Impalad nodes, respectively:

```
scp fair-scheduler.xml {IP address of the Impalad instance}:/opt/Bigdata/
FusionInsight_Impala_***/***_Impalad/etc/
```

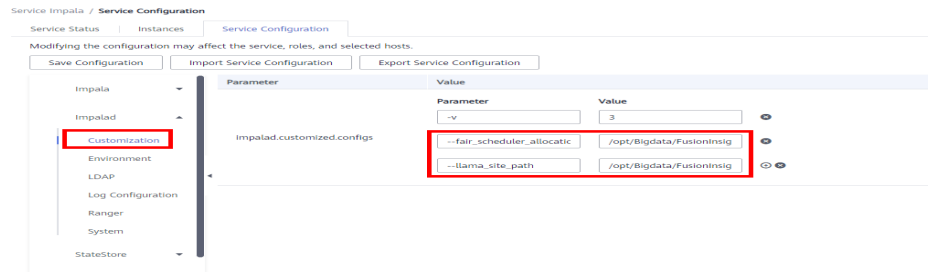
```
scp llama-site.xml {IP address of the Impalad instance}:/opt/Bigdata/
FusionInsight_Impala_***/***_Impalad/etc/
```

```
+ scp fair-scheduler.xml 192.168.1.19:/opt/Bigdata/FusionInsight_Impala_8.1.0.1/1_21_Impalad/etc/
Warning: Permanently added '192.168.1.19' (ED25519) to the list of known hosts.
```


- Log in to FusionInsight Manager, find the Impala component, and add the following custom configuration items and values to the Impala instance:

--fair\_scheduler\_allocation\_path The value is, for example, /opt/Bigdata/FusionInsight\_Impala\_\*\*\*/\*\*\*\_Impalad/etc/fair-scheduler.xml

--llama\_site\_path The value is, for example, /opt/Bigdata/FusionInsight\_Impala\_\*\*\*/\*\*\*\_Impalad/etc/llama-site.xml



- Restart the Impalad instance.



- Log in to the node where the Impala client is installed, run the **source** command to obtain the environment variables, and run the following command:

```
impala-shell -i {IP address of the Impalad instance:Port number} -Q
request_pool=root.default (resource pool configured in fair-scheduler.xml
and llama-site.xml)
```

```
[root@node-master1ioKo ~]# impala-shell -i 192.168.1.19:21000 -Q request_pool=root.default
Starting Impala Shell without Kerberos authentication
Opened TCP connection to 192.168.1.19:21000
Connected to 192.168.1.19:21000
Server version: impalad version 3.4.0-RELEASE RELEASE (build ac0f95df4baa94cfd36ef370f6a432d582ac1f)

Welcome to the Impala shell.
(Impala Shell v3.4.0-RELEASE (f68c12e) built on Sat Jun 26 17:16:01 CST 2021)

The '-B' command line flag turns off pretty-printing for query results. Use this
flag to remove formatting from results you want to save for later, or to benchmark
Impala.

[192.168.1.19:21000] default>
```

Execute SQL statements.

```
192.168.1.19:21000] default> select * from test1;
Query: select * from test1
Query submitted at: 2022-05-12 10:01:01 (Coordinator: http://192.168.1.19:25000)
Query progress can be monitored at: http://192.168.1.19:25000/query_plan?query_id=97440454dbab28ea35bd90d600000000
ERROR: Rejected query from pool root.default: Invalid pool config: the min_query_mem_limit 2048000 is greater than the max_mem_r
esources 1048576 (configured statically)
```

- Log in to the Impalad WebUI to check the resource pool usage and ensure that the configuration has taken effect.

<https://{{Cluster console address}}:9022/component/Impala/Impalad/95/>

**Admission Controller** Reset informational stats for all pools

This page lists all resource pools to which queries have been submitted at least once and their corresponding state and statistics. See the [backends](#) debug page for memory admitted and reserved per backend.

Time since last statestore update containing admission control topic state (ms): 24

**root.default**

**Pool Config**

Property	Value
Max memory (cluster wide)	1048576
Max concurrent queries	1
Max queue size	2
Queue Timeout (ms)	30000
<b>Min Query MEM_LIMIT range</b>	<b>2048000</b>
<b>Max Query MEM_LIMIT range</b>	<b>1024000</b>
Clamp MEM_LIMIT query option	true

# 15 Using Kafka

---

## 15.1 Using Kafka from Scratch

### Scenario

You can create, query, and delete topics on a cluster client.

### Prerequisites

The client has been installed in a directory, for example, `/opt/client`. The client directory in the following operations is only an example. Change it based on site requirements.

### Using the Kafka Client (Versions Earlier Than MRS 3.x)

**Step 1** Access the ZooKeeper instance page.

Click the cluster name to go to the cluster details page and choose **Components > ZooKeeper > Instances**.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

**Step 2** View the IP addresses of the ZooKeeper role instance.

Record any IP address of the ZooKeeper instance.

**Step 3** Log in to the node where the client is installed.

**Step 4** Run the following command to switch to the client installation directory, for example, `/opt/client/Kafka/kafka/bin`.

```
cd /opt/client/Kafka/kafka/bin
```

**Step 5** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

- Step 6** If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit Kafka user
```

- Step 7** Create a topic.

```
sh kafka-topics.sh --create --topic Topic name --partitions Number of partitions occupied by the topic --replication-factor Number of replicas of the topic --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

Example: **sh kafka-topics.sh --create --topic TopicTest --partitions 3 --replication-factor 3 --zookeeper 10.10.10.100:2181/kafka**

- Step 8** Run the following command to view the topic information in the cluster:

```
sh kafka-topics.sh --list --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

Example: **sh kafka-topics.sh --list --zookeeper 10.10.10.100:2181/kafka**

- Step 9** Delete the topic created in [Step 7](#).

```
sh kafka-topics.sh --delete --topic Topic name --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

Example: **sh kafka-topics.sh --delete --topic TopicTest --zookeeper 10.10.10.100:2181/kafka**

Type **y** and press **Enter**.

```
----End
```

## Using the Kafka Client (MRS 3.x or Later)

- Step 1** Access the ZooKeeper instance page.

Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Instance**.

- Step 2** View the IP addresses of the ZooKeeper role instance.

Record any IP address of the ZooKeeper instance.

- Step 3** Log in to the node where the client is installed.

- Step 4** Run the following command to switch to the client installation directory, for example, **/opt/client/Kafka/kafka/bin**.

```
cd /opt/client/Kafka/kafka/bin
```

- Step 5** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

- Step 6** If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step.

**kinit** *Kafka user*

**Step 7** Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > ZooKeeper**, and click the **Configurations** tab and then **All Configurations**. On the displayed page, search for the **clientPort** parameter and record its value.

**Step 8** Create a topic.

```
sh kafka-topics.sh --create --topic Topic name --partitions Number of partitions occupied by the topic --replication-factor Number of replicas of the topic --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

Example: **sh kafka-topics.sh --create --topic TopicTest --partitions 3 --replication-factor 3 --zookeeper 10.10.10.100:2181/kafka**

**Step 9** Run the following command to view the topic information in the cluster:

```
sh kafka-topics.sh --list --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

Example: **sh kafka-topics.sh --list --zookeeper 10.10.10.100:2181/kafka**

**Step 10** Delete the topic created in [Step 8](#).

```
sh kafka-topics.sh --delete --topic Topic name --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

Example: **sh kafka-topics.sh --delete --topic TopicTest --zookeeper 10.10.10.100:2181/kafka**

----End

## 15.2 Managing Kafka Topics

### Scenario

You can manage Kafka topics on a cluster client based on service requirements. Management permission is required for clusters with Kerberos authentication enabled.

### Prerequisites

You have installed the Kafka client.

### Procedure

**Step 1** Access the ZooKeeper instance page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components > ZooKeeper > Instances**.

#### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Instance**.

**Step 2** View the IP addresses of the ZooKeeper role instance.

Record any IP address of the ZooKeeper instance.

**Step 3** Prepare the client based on service requirements. Log in to the node where the client is installed.

Log in to the node where the client is installed.. )

**Step 4** Run the following command to switch to the client installation directory, for example, `/opt/client/Kafka/kafka/bin`.

```
cd /opt/client/Kafka/kafka/bin
```

**Step 5** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

**Step 6** Run the following command to perform user authentication (skip this step in normal mode):

```
kinit Component service user
```

**Step 7** For versions earlier than MRS 3.x, run the following commands to manage Kafka topics:

- Creating a topic

```
sh kafka-topics.sh --create --topic Topic name --partitions Number of partitions occupied by the topic --replication-factor Number of replicas of the topic --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

- Deleting a topic

```
sh kafka-topics.sh --delete --topic Topic name --zookeeper IP address of the node where the ZooKeeper instance resides:clientPort/kafka
```

 **NOTE**

- The number of topic partitions or topic backup replicas cannot exceed the number of Kafka instances.
- By default, the value of **clientPort** of ZooKeeper is **2181**.
- There are three ZooKeeper instances. Use the IP address of any one.
- For details about managing messages in Kafka topics, see [Managing Messages in Kafka Topics](#).

**Step 8** MRS 3.x and later versions: Use **kafka-topics.sh** to manage Kafka topics.

- Creating a topic:

By default, partitions of a topic are distributed based on the number of partitions on the node and disk. To distribute partitions based on the disk capacity, set **log.partition.strategy** to **capacity** for the Kafka service.

When a topic is created in Kafka, partitions and copies can be generated based on the combination of rack awareness and cross-AZ feature. The **--zookeeper** and **--bootstrap-server** modes are supported.

- Disable the rack policy and cross-AZ feature (default policy).  
Copies of topics created based on this policy are randomly allocated to any node in the cluster.  
**`./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --zookeeper IP address of any ZooKeeper node:clientPort/kafka`**  
**`./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --bootstrap-server IP address of the Kafkacluster:21007 --command-config ../config/client.properties`**  
If you use `--bootstrap-server` to create a topic, set `rack.aware.enable` and `az.aware.enable` to **false**.
- Enable the rack policy and disable the cross-AZ feature.  
The leader of each partition of the topic created based on this policy is randomly allocated on the cluster node. However, different replicas of the same partition are allocated to different racks. Therefore, when this policy is used, ensure that the number of nodes in each rack is the same, otherwise, the load of nodes in the rack with fewer nodes is much higher than the average load of the cluster.  
**`./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --zookeeper IP address of any ZooKeeper node:clientPort/kafka --enable-rack-aware`**  
**`./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --bootstrap-server IP address of the Kafkacluster:21007 --command-config ../config/client.properties`**  
If you use `--bootstrap-server` to create a topic, set `rack.aware.enable` to **true** and `az.aware.enable` to **false**.
- Disable the rack policy and enable the cross-AZ feature.  
The leader of each partition of the topic created based on this policy is randomly allocated on the cluster node. However, different replicas of the same partition are allocated to different AZs. Therefore, when this policy is used, ensure that the number of nodes in each AZ is the same, otherwise, the load of nodes in the AZ with fewer nodes is much higher than the average load of the cluster.  
**`./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --zookeeper IP address of any ZooKeeper node:clientPort/kafka --enable-az-aware`**  
**`./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --bootstrap-server IP address of the Kafkacluster:21007 --command-config ../config/client.properties`**  
If you use `--bootstrap-server` to create a topic, set `rack.aware.enable` to **false** and `az.aware.enable` to **true**.
- Enable the rack policy and cross-AZ feature.  
The leader of each partition of the topic created based on this policy is randomly allocated on the cluster node. However, different replicas of the

same partition are allocated to different racks in different AZs. This policy ensures that the number of nodes on each rack in each AZ is the same, otherwise, the load in the cluster is unbalanced.

```
./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --zookeeper IP address of any ZooKeeper node:clientPort/kafka --enable-rack-aware --enable-az-aware
```

```
./kafka-topics.sh --create --topic topic name --partitions number of partitions occupied by the topic --replication-factor number of replicas of the topic --bootstrap-server IP address of the Kafkacluster:21007 --command-config ../config/client.properties
```

If you use **--bootstrap-server** to create a topic, set **rack aware.enable** and **az aware.enable** to **true**.

 **NOTE**

- Kafka supports topic creation in either of the following modes:
  - In **--zookeeper** mode, the client generates a copy allocation scheme. The community supports this mode from the beginning. To reduce the dependency on the ZooKeeper component, the community will delete the support for this mode in later versions. When creating a topic in this mode, you can select a copy allocation policy by combining the **--enable-rack-aware** and **--enable-az-aware** options. Note: The **--enable-az-aware** option can be used only when the cross-AZ feature is enabled on the server, that is, **az aware.enable** is set to **true**. Otherwise, the execution fails.
  - In **--bootstrap-server** mode, the server generates a copy allocation solution. In later versions, the community supports only this mode for topic management. When a topic is created in this mode, the **--enable-rack-aware** and **--enable-az-aware** options cannot be used to control the copy allocation policy. The **rack aware.enable** and **az aware.enable** parameters can be used together to control the copy allocation policy. Note that the **az aware.enable** parameter cannot be modified; if the cross-AZ feature is enabled during cluster creation, this parameter is automatically set to **true**; the **rack aware.enable** parameter can be customized.
- List of topics:
  - **./kafka-topics.sh --list --zookeeper** *service IP address of any ZooKeeper node:clientPort/kafka*
  - **./kafka-topics.sh --list --bootstrap-server** *IP address of the Kafkacluster:21007* **--command-config** *../config/client.properties*
- Viewing the topic:
  - **./kafka-topics.sh --describe --zookeeper** *service IP address of any ZooKeeper node:clientPort/kafka* **--topic** *topic name*
  - **./kafka-topics.sh --describe --bootstrap-server** *IP address of the Kafkacluster:21007* **--command-config** *../config/client.properties* **--topic** *topic name*
- Modifying a topic:
  - **./kafka-topics.sh --alter --topic** *topic name* **--config** *configuration item=configuration value* **--zookeeper** *service IP address of any ZooKeeper node:clientPort/kafka*
- Expanding partitions:



- `./kafka-topics.sh --alter --topic topic name --zookeeper service IP address of any ZooKeeper node:clientPort/kafka --command-config Kafka/kafka/config/client.properties --partitions number of partitions after the expansion`
  - `./kafka-topics.sh --alter --topic topic name --bootstrap-server IP address of the Kafka cluster:21007 --command-config Kafka/kafka/config/client.properties --partitions number of partitions after the expansion`
  - Deleting a topic
    - `./kafka-topics.sh --delete --topic topic name --zookeeper Service IP address of any ZooKeeper node:clientPort/kafka`
    - `./kafka-topics.sh --delete --topic topic name --bootstrap-server IP address of the Kafka cluster:21007 --command-config ./config/client.properties`
- End

## 15.3 Querying Kafka Topics

### Scenario

You can query existing Kafka topics on MRS.

### Procedure

**Step 1** Go to the Kafka service page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components > Kafka**.

#### NOTE

- If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)
- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Name of the desired cluster > Services > Kafka**.

**Step 2** Click **KafkaTopicMonitor**.

All topics are displayed in the list by default. You can view the number of partitions and replicas of the topics.

**Step 3** Click the desired topic in the list to view its details.

 NOTE

If the following operations are performed, Kafka topic monitoring may not be displayed:

- Capacity expansion or reduction has been performed on Kafka or ZooKeeper.
- Instances have been added to or deleted from Kafka or ZooKeeper.
- The Elasticsearch service is reinstalled.
- Kafka is switched to another ZooKeeper service.

Perform the following steps to rectify the fault:

1. Log in to the active OMS node of the cluster and run the following command to switch to user **omm**:

```
su - omm
```

2. Restart the CEP service.

```
restart_app cep
```

Wait for 3 minutes and check the Kafka topic monitoring again.

----End

## 15.4 Managing Kafka User Permissions

### Scenario

For clusters with Kerberos authentication enabled, using Kafka requires relevant permissions. MRS clusters can grant the use permission of Kafka to different users.

**Table 15-1** lists the default Kafka user groups.

 NOTE

In MRS 3.x or later, Kafka supports two types of authentication plug-ins: Kafka open source authentication plug-in and Ranger authentication plug-in.

This section describes the user permission management based on the Kafka open source authentication plug-in. For details about how to use the Ranger authentication plug-in, see [Adding a Ranger Access Permission Policy for Kafka](#).

**Table 15-1** Default Kafka user groups

User Group	Description
kafkaadmin	Kafka administrator group. Users in this group have the permissions to create, delete, read, and write all topics, and authorize other users.
kafkasuperuser	Kafka super user group. Users in this group have the permissions to read and write all topics.
kafka	Kafka common user group. Users in this group can access a topic only when they are granted with the read and write permissions of the topic by a user in the <b>kafkaadmin</b> group.

## Prerequisites

- You have installed the Kafka client.
- A user in the **kafkaadmin** group, for example **admin**, has been prepared.

## Procedure

**Step 1** Access the ZooKeeper instance page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components > ZooKeeper > Instances**.

### NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Name of the desired cluster > Services > ZooKeeper > Instance**.

**Step 2** View the IP addresses of the ZooKeeper role instance.

Record the IP address of any ZooKeeper instance.

**Step 3** Prepare the client based on service requirements. Log in to the node where the client is installed.

Log in to the node where the client is installed.. )

**Step 4** Run the following command to switch to the client installation directory, for example, **/opt/client/Kafka/kafka/bin**.

```
cd /opt/client/Kafka/kafka/bin
```

**Step 5** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

**Step 6** Run the following command to authenticate the user(skip this step in normal mode):

```
kinit Component service user
```

**Step 7** Versions earlier than MRS 3.x: Select the scenario required by the service and manage Kafka user permissions.

- Querying the permission list of a topic  

```
sh kafka-acls.sh --authorizer-properties zookeeper.connect=IP address of the node where the ZooKeeper instance resides:2181/kafka --list --topic Topic name
```
- Adding producer permission to a user  

```
sh kafka-acls.sh --authorizer-properties zookeeper.connect=IP address of the node where the ZooKeeper instance resides:2181/kafka --add --allow-principal User:Username --producer --topic Topic name
```
- Removing producer permission of a user

```
sh kafka-acls.sh --authorizer-properties zookeeper.connect=IP address of the node where the ZooKeeper instance resides:2181/kafka --remove --allow-principal User:Username --producer --topic Topic name
```

- Adding consumer permission to a user

```
sh kafka-acls.sh --authorizer-properties zookeeper.connect=IP address of the node where the ZooKeeper instance resides:2181/kafka --add --allow-principal User:Username --consumer --topic Topic name --group Consumer group name
```

- Removing consumer permission of a user

```
sh kafka-acls.sh --authorizer-properties zookeeper.connect=IP address of the node where the ZooKeeper instance resides:2181/kafka --remove --allow-principal User:Username --consumer --topic Topic name --group Consumer group name
```

 NOTE

You need to enter **y** twice to confirm the removal of permission.

**Step 8** MRS 3.x and later versions: The following lists the common **kafka-acl.sh** commands used for user authorization:

- View the permission control list of a topic:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node>:2181/kafka > --list --topic <Topic name>
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster>:21007 --command-config ../config/client.properties --list --topic <topic name>
```

- Add the Producer permission for a user:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node>:2181/kafka > --add --allow-principal User:<Username> --producer --topic <Topic name>
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster>:21007 --command-config ../config/client.properties --add --allow-principal User:<username> --producer --topic <topic name>
```

- Assign the Producer permission to a user in batches.

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node>:2181/kafka > --add --allow-principal User:<Username> --producer --topic <Topic name> --resource-pattern-type prefixed
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster>:21007 --command-config ../config/client.properties --add --allow-principal User:<username> --producer --topic <topic name> --resource-pattern-type prefixed
```

- Remove the Producer permission from a user:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node>:2181/kafka > --remove --allow-principal User:<Username> --producer --topic <Topic name>
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster>:21007 --command-config ../config/client.properties --remove --allow-principal User:<username> --producer --topic <topic name>
```

- Delete the Producer permission of a user in batches:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node:2181/kafka > --remove --allow-principal User:<Username> --producer --topic <Topic name> --resource-pattern-type prefixed
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster:21007> --command-config ../config/client.properties --remove --allow-principal User:<username> --producer --topic <topic name>--resource-pattern-type prefixed
```

- Add the Consumer permission for a user:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node:2181/kafka > --add --allow-principal User:<Username> --consumer --topic <Topicname> --group <Consumer group name>
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster:21007> --command-config ../config/client.properties --add --allow-principal User:<username> --consumer --topic <topicname> --group <consumer group name>
```

- Add consumer permissions to a user in batches:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node:2181/kafka > --add --allow-principal User:<Username> --consumer --topic <Topic name> --group <Consumer group name> --resource-pattern-type prefixed
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster:21007> --command-config ../config/client.properties --add --allow-principal User:<username> --consumer --topic <topicname> --group <consumer group name> --resource-pattern-type prefixed
```

- Remove the consumer permission from a user:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node:2181/kafka > --remove --allow-principal User:<Username> --consumer --topic <Topic name> --group <Consumer group name>
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster:21007> --command-config ../config/client.properties --remove --allow-principal User:<username> --consumer --topic <topic name> --group <consumer group name>
```

- Delete the consumer permission of a user in batches:

```
./kafka-acls.sh --authorizer-properties zookeeper.connect=<Service IP address of any ZooKeeper node:2181/kafka > --remove --allow-principal User:<Username> --consumer --topic <Topic name> --group <Consumer group name> --resource-pattern-type prefixed
```

```
./kafka-acls.sh --bootstrap-server <IP address of the Kafkacluster:21007> --command-config ../config/client.properties --remove --allow-principal User:<username> --consumer --topic <topicname> --group <consumer group name> --resource-pattern-type prefixed
```

----End

## 15.5 Managing Messages in Kafka Topics

### Scenario

You can produce or consume messages in Kafka topics using the MRS cluster client.

### Prerequisites

- The cluster client has been installed.
- For clusters with Kerberos authentication enabled, you need to create a service user on Manager in advance. The user has the permission to perform operations in Kafka topics.

### Procedure

**Step 1** Go to the Kafka service page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components > Kafka**.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > Kafka**.

**Step 2** Click **instance**. Query the IP addresses of the Kafka broker instances.

Record the IP address of any Kafka instance.

**Step 3** Prepare the client based on service requirements. Log in to the node where the client is installed.

Log in to the node where the client is installed.. )

**Step 4** Run the following command to switch to the client installation directory, for example, **/opt/client/Kafka/kafka/bin**.

```
cd /opt/client/Kafka/kafka/bin
```

**Step 5** Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
```

**Step 6** For clusters with Kerberos authentication enabled, run the following command to authenticate the user. For clusters with Kerberos authentication disabled, skip this step.

```
kinit Kafka user
```

**Step 7** Manage messages in Kafka topics using the following commands:

- Producing messages

```
sh kafka-console-producer.sh --broker-list IP address of the node where the broker instance is located:9092 --topic Topic name --producer.config /opt/client/Kafka/kafka/config/producer.properties
```

A topic must be created in advance. You can input specified information as the messages produced by the producer and then press **Enter** to send the messages. To end message producing, press **Ctrl + C** to exit.

- Consuming messages

Start another client connection and run the following commands to consume messages in the topic:

```
cd /opt/client/Kafka/kafka/bin
```

```
source bigdata_env
```

```
sh kafka-console-consumer.sh --topic Topic name --bootstrap-server IP address of the node where the broker instance is located:9092 --consumer.config /opt/client/Kafka/kafka/config/consumer.properties
```

In the configuration file, **group.id** (indicating the consumer group) is set to **example-group1** by default. Users can change the value as required. The value takes effect each time consumption occurs.

By default, the system reads unprocessed messages in the current consumer group when the command is executed. If a new consumer group is specified in the configuration file and the **--from-beginning** parameter is added to the command, the system reads all messages that have not been automatically deleted in Kafka.

 **NOTE**

----End

## 15.6 Synchronizing Binlog-based MySQL Data to the MRS Cluster

This section describes how to use the Maxwell data synchronization tool to migrate offline binlog-based data to an MRS Kafka cluster.

Maxwell is an open-source application that reads MySQL binlogs, converts operations, such as addition, deletion, and modification, into a JSON format, and sends them to an output end, such as a console, a file, and Kafka. Maxwell can be deployed on a MySQL server or on other servers that can communicate with MySQL.

Maxwell runs on a Linux server, including EulerOS, Ubuntu, Debian, CentOS, and OpenSUSE. Java 1.8+ must be supported.

The following provides details about data synchronization.

1. [Configuring MySQL](#)
2. [Installing Maxwell](#)
3. [Configuring Maxwell](#)
4. [Starting Maxwell](#)
5. [Verifying Maxwell](#)

6. [Stopping Maxwell](#)
7. [Format of the Maxwell Generated Data and Description of Common Fields](#)

## Configuring MySQL

- Step 1** Start the binlog, open the **my.cnf** file in MySQL, and check whether **server\_id**, **log-bin**, and **binlog\_format** are configured in the **[mysqld]** block. If they are not configured, run the following command to add configuration items and restart MySQL. If they are configured, skip this step.

```
$ vi my.cnf

[mysqld]
server_id=1
log-bin=master
binlog_format=row
```

- Step 2** Maxwell needs to connect to MySQL, create a database named **maxwell** for storing metadata, and access the database to be synchronized. Therefore, you are advised to create a MySQL user for Maxwell to use. Log in to MySQL as user **root** and run the following commands to create a user named **maxwell** (**XXXXXX** indicates the password and needs to be replaced with actual one).

- If Maxwell is deployed on a non-MySQL server, the created user **maxwell** must have a permission to remotely log in to the database. In this case, run the following command to create the user:  

```
mysql> GRANT ALL on maxwell.* to 'maxwell'@'%' identified by 'XXXXXX';
```

```
mysql> GRANT SELECT, REPLICATION CLIENT, REPLICATION SLAVE on *.* to 'maxwell'@'%';
```
  - If Maxwell is deployed on the MySQL server, the created user **maxwell** can be configured to log in to the database only on the local host. In this case, run the following command:  

```
mysql> GRANT SELECT, REPLICATION CLIENT, REPLICATION SLAVE on *.* to 'maxwell'@'localhost' identified by 'XXXXXX';
```

```
mysql> GRANT ALL on maxwell.* to 'maxwell'@'localhost';
```
- End

## Installing Maxwell

- Step 1** Download the installation package at <https://github.com/zendesk/maxwell/releases> and select the **maxwell-XXX.tar.gz** binary file for download. In the file name, **XXX** indicates a version number.
- Step 2** Upload the **tar.gz** package to any directory (the **/opt** directory of the Master node used as an example here).
- Step 3** Log in to the server where Maxwell is deployed and run the following command to go to the directory where the **tar.gz** package is stored.
- ```
cd /opt
```
- Step 4** Run the following commands to decompress the **maxwell-XXX.tar.gz** package and go to the **maxwell-XXX** directory:


```
tar -zxvf maxwell-XXX.tar.gz
cd maxwell-XXX
----End
```

Configuring Maxwell

If the **conf** directory exists in the **maxwell-XXX** folder, configure the **config.properties** file. For details about the configuration items, see [Table 15-2](#). If the **conf** directory does not exist, change **config.properties.example** in the **maxwell-XXX** folder to **config.properties**.

Table 15-2 Maxwell configuration item description

| Parameter | Mandatory | Description | Default Value |
|-----------------------|-----------|--|---------------|
| user | Yes | Name of the user for connecting to MySQL, that is, the user created in Step 2 . | - |
| password | Yes | Password for connecting to MySQL | - |
| host | No | MySQL address | localhost |
| port | No | MySQL port | 3306 |
| log_level | No | Log print level. The options are as follows: <ul style="list-style-type: none"> • debug • info • warn • error | info |
| output_ddl | No | Whether to send a DDL (modified based on definitions of the database and data table) event <ul style="list-style-type: none"> • true: Send DDL events. • false: Do not send DDL events. | false |
| producer | Yes | Producer type. Set this parameter to kafka . <ul style="list-style-type: none"> • stdout: Log the generated events. • kafka: Send the generated events to Kafka. | stdout |
| producer_partition_by | No | Partition policy used to ensure that data of the same type is written to the same partition of Kafka. <ul style="list-style-type: none"> • database: Events of the same database are written to the same partition of Kafka. • table: Events of the same table are written to the same partition of Kafka. | database |

| Parameter | Mandatory | Description | Default Value |
|-------------------------|-----------|---|---------------|
| ignore_producer_error | No | Specifies whether to ignore the error that the producer fails to send data. <ul style="list-style-type: none"> true: The error information is logged and the error data is skipped. The program continues to run. false: The error information is logged and the program is terminated. | true |
| metrics_slf4j_interval | No | Interval for outputting statistics on data successfully uploaded or failed to be uploaded to Kafka in logs. The unit is second. | 60 |
| kafka.bootstrap.servers | Yes | Address of the Kafka proxy node. The value is in the format of HOST:PORT[,HOST:PORT] . | - |
| kafka_topic | No | Name of the topic that is written to Kafka | maxwell |
| dead_letter_topic | No | Kafka topic used to record the primary key of the error log record when an error occurs when the record is sent | - |
| kafka_version | No | Kafka producer version used by Maxwell, which cannot be configured in the config.properties file. You need to use the --kafka_version xxx parameter to import the version number when starting the command. | - |
| kafka_partition_hash | No | Kafka topic partitioning algorithm. The value can be default or murmur3 . | default |
| kafka_key_format | No | Key generation method of the Kafka record. The value can be array or Hash . | Hash |
| ddl_kafka_topic | No | Topic that is written to the DDL operation when output_ddl is set to true | {kafka_topic} |

| Parameter | Mandatory | Description | Default Value |
|-----------|-----------|--|---------------|
| filter | No | <p>Used to filter databases or tables.</p> <ul style="list-style-type: none"> If only the mydatabase database needs to be collected, set this parameter to the following:
exclude: *.*;include: mydatabase.* If only the mydatabase.mytable table needs to be collected, set this parameter to the following:
exclude: *.*;include: mydatabase.mytable If only the mytable, mydate_123, and mydate_456 tables in the mydatabase database need to be collected, set this parameter to the following:
exclude: *.*;include: mydatabase.mytable, include: mydatabase./mydate_\\d*/ | - |

Starting Maxwell

Step 1 Log in to the server where Maxwell is deployed.

Step 2 Run the following command to go to the Maxwell installation directory:

```
cd /opt/maxwell-1.21.0/
```

NOTE

For the first time to use Maxwell, you are advised to change **log_level** in **conf/config.properties** to **debug** (debug level) so that you can check whether data can be obtained from MySQL and sent to Kafka after startup. After the entire process is debugged, change **log_level** to **info**, and then restart Maxwell for the modification to take effect.

```
# log level [debug | info | warn | error]
```

```
log_level=debug
```

Step 3 Run the following commands to start Maxwell:

```
source /opt/client/bigdata_env
```

```
bin/Maxwell
```

```
bin/maxwell --user='maxwell' --password='XXXXXX' --host='127.0.0.1' \
```

```
--producer=kafka --kafka.bootstrap.servers=kafkahost:9092 --  
kafka_topic=Maxwell
```

In the preceding commands, **user**, **password**, and **host** indicate the username, password, and IP address of MySQL, respectively. You can configure the three parameters by modifying configurations of the configuration items or using the preceding commands. **kafkahost** indicates the IP address of the Core node in the streaming cluster.

If information similar to the following appears, Maxwell has started successfully:

```
Success to start Maxwell [78092].
```

----End

Verifying Maxwell

Step 1 Log in to the server where Maxwell is deployed.

Step 2 View the logs. If the log file does not contain an ERROR log and the following information is displayed, the connection between Maxwell and MySQL is normal:

```
BinlogConnectorLifecycleListener - Binlog connected.
```

Step 3 Log in to the MySQL database and update, create, or delete test data. The following provides operation statement examples for your reference.

```
--Creating a database
create database test;
--Creating a table
create table test.e (
  id int(10) not null primary key auto_increment,
  m double,
  c timestamp(6),
  comment varchar(255) charset 'latin1'
);
-- Adding a record
insert into test.e set m = 4.2341, c = now(3), comment = 'I am a creature of light.';
--Updating a record
update test.e set m = 5.444, c = now(3) where id = 1;
--Deleting a record
delete from test.e where id = 1;
--Modifying a table
alter table test.e add column torvalds bigint unsigned after m;
--Deleting a table
drop table test.e;
-- Deleting a database
drop database test;
```

Step 4 Check the Maxwell logs. If no WARN/ERROR is displayed, Maxwell is installed and configured properly.

To check whether the data is successfully uploaded, set **log_level** in the **config.properties** file to **debug**. When the data is successfully uploaded, the following JSON data is printed immediately. For details about the fields, see [Format of the Maxwell Generated Data and Description of Common Fields](#).

```
{"database":"test","table":"e","type":"insert","ts":1541150929,"xid":60556,"commit":true,"data":
{"id":1,"m":4.2341,"c":"2018-11-02 09:28:49.297000","comment":"I am a creature of light."}}
.....
```

NOTE

After the entire process is debugged, you can change the value of **log_level** in the **config.properties** file to **info** to reduce the number of logs to be printed and restart Maxwell for the modification to take effect.

```
# log level [debug | info | warn | error]
log_level=info
```

----End

Stopping Maxwell

Step 1 Log in to the server where Maxwell is deployed.

Step 2 Run the command to obtain the Maxwell process ID (PID). The second field in the command output is PID.

```
ps -ef | grep Maxwell | grep -v grep
```

Step 3 Run the following command to forcibly stop the Maxwell process:

```
kill -9 PID
```

```
----End
```

Format of the Maxwell Generated Data and Description of Common Fields

The data generated by Maxwell is in JSON format. The common fields are described as follows:

- **type**: operation type. The options are **database-create**, **database-drop**, **table-create**, **table-drop**, **table-alter**, **insert**, **update**, and **delete**.
- **database**: name of the database to be operated
- **ts**: operation time, which is a 13-digit timestamp
- **table**: name of the table to be operated
- **data**: content after data is added, deleted, or modified
- **old**: content before data is modified or schema definition before a table is modified
- **sql**: SQL statement for DDL operations
- **def**: schema definition for table creation and modification
- **xid**: unique ID of an object
- **commit**: check whether such operations as data addition, deletion, and modification have been submitted.

15.7 Creating a Kafka Role

Scenario

Create and configure a Kafka role as an MRS cluster administrator.

This section applies to MRS 3.x or later.

NOTE

Users can create Kafka roles only in security mode.

If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for Kafka](#).

Prerequisites

The MRS cluster administrator has understood service requirements.

Procedure

Step 1 Log in to FusionInsight Manager and choose **System > Permission > Role**.

Step 2 On the displayed page, click **Create Role** and enter a **Role Name** and **Description**.

- Step 3** On the **Configure Resource Permission** page, choose *Name of the desired cluster* > **Kafka**.
- Step 4** Select permissions based on service requirements. For details about configuration items, see [Table 15-3](#).

Table 15-3 Description

| Scenario | Role Authorization |
|---|--|
| Setting the Kafka administrator permissions | In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Kafka > Kafka Manager Privileges .

NOTE
This permission allows you to create and delete topics, but does not allow you to produce or consume any topics. |
| Setting the production permission of a user on a topic | <ol style="list-style-type: none"> In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Kafka > Kafka Topic Producer And Consumer Privileges. In the Permission column of the specified topic, select Kafka Producer Permission. |
| Setting the consumption permission of a user on a topic | <ol style="list-style-type: none"> In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Kafka > Kafka Topic Producer And Consumer Privileges. In the Permission column of the specified topic, select Kafka Consumer Privileges. |

- Step 5** Click **OK**, and return to the **Role** page.
- End

15.8 Kafka Common Parameters

This section applies to MRS 3.x or later.

Navigation path for setting parameters:

For details about how to set parameters, see [Modifying Cluster Service Configuration Parameters](#).

Common Parameters

Table 15-4 Parameter description

| Parameter | Description | Default Value |
|----------------------------|--|--|
| log.dirs | List of Kafka data storage directories. Use commas (,) to separate multiple directories. | %
{@auto.detect.datapart.b
k.log.logs} |
| KAFKA_HEAP_OPTS | Specifies the JVM option used for Kafka to start broker. It is recommended that you set this parameter based on service requirements. | -Xmx6G -Xms6G |
| auto.create.topics.enable | Indicates whether a topic is automatically created. If this parameter is set to false , you need to run a command to create a topic before sending a message. | true |
| default.replication.factor | Default number of replicas of a topic is automatically created. | 2 |
| monitor.preInitDelay | Delay of the first health check after the server is started. If the startup takes a long time, increase the value of the parameter. Unit: millisecond | 600,000 |

Timeout Parameters

Table 15-5 Broker-related timeout parameters

| Parameter | Description | Default Value | Impact |
|------------------------------|--|---------------|--|
| controller.socket.timeout.ms | Specifies the timeout for connecting controller to broker. Unit: millisecond | 30,000 | Generally, retain the default value of this parameter. |

| Parameter | Description | Default Value | Impact |
|------------------------------|---|---------------|--|
| group.max.session.timeout.ms | Specifies the maximum session timeout during the consumer registration. Unit: millisecond | 1800000 | The configured value must be less than the value of this parameter. |
| group.min.session.timeout.ms | Specifies the minimum session timeout during the consumer registration. Unit: millisecond | 6000 | The configured value must be greater than the value of this parameter. |
| offsets.commit.timeout.ms | Specifies the timeout for the Offset to submit requests. Unit: millisecond | 5000 | This parameter specifies the maximum delay for processing an Offset request. |
| replica.socket.timeout.ms | Specifies the timeout of the request for synchronizing replica data. Its value must be greater than or equal to that of the replica.fetch.wait.max.ms parameter. Unit: millisecond | 30000 | Specifies the maximum timeout for establishing a channel before the synchronization thread sends a synchronization request. The value must be greater than that of the replica.fetch.wait.max.ms parameter. |
| request.timeout.ms | Specifies the timeout for waiting for a response after the client sends a connection request. If no response is received within the timeout, the client resends the request. A request failure is returned after the maximum retry times is reached. Unit: millisecond | 30000 | This parameter is configured when the networkclient connection is transferred in the controller and replica threads on the broker node. |
| transaction.max.timeout.ms | Specifies the maximum timeout allowed by the transaction. If the client request time exceeds the value of this parameter, broker returns an error in InitProducerIdRequest. This prevents a long client request timeout, ensuring that consumer can receive topics. Unit: millisecond | 900000 | Specifies the maximum timeout for transactions. |

| Parameter | Description | Default Value | Impact |
|----------------------------------|---|---------------|---|
| user.group.cache.timeout.seconds | Specifies the time when the user group information is stored in the cache. Unit: second | 300 | Specifies the time for caching the mapping between users and user groups. If time exceeds the threshold, the system automatically runs the id -Gn command to query the user information. During this period, the mapping in the cache is used. |
| zookeeper.connection.timeout.ms | Specifies the timeout for connecting to ZooKeeper. Unit: millisecond | 45,000 | This parameter specifies the duration for connecting the ZooKeeper and zkclient for the first time. If the duration exceeds the value of this parameter, the zkclient automatically disconnects the connection. |

| Parameter | Description | Default Value | Impact |
|------------------------------|---|---------------|--|
| zookeeper.session.timeout.ms | Specifies the ZooKeeper session timeout duration. During this period, ZooKeeper disconnects the connection if broker does not report its heartbeats to ZooKeeper. Unit: millisecond | 45,000 | ZooKeeper session timeout has the following functions:
1) Based on value of this parameter and the number of ZooKeeper URLs in ZKURL, if the connection duration exceeds the node timeout value (sessionTimeout/ Number of transferred ZooKeeper URLs), the connection fails and the system attempts to connect to the next node.
2) After the connection is established, a session (for example, the temporary BrokerId node registered on the ZooKeeper) is cleared by the ZooKeeper a session timeout later if the broker is stopped. |

Table 15-6 Producer-related timeout parameters

| Parameter | Description | Default Value | Impact |
|--------------------|---|---------------|--|
| request.timeout.ms | Specifies the timeout of a message request. | 30,000 | If a network fault occurs, increase the value of this parameter. If the value is too small, the Batch Expire occurs. |

Table 15-7 Consumer-related timeout parameters

| Parameter | Description | Default Value | Impact |
|-------------------------|--|---------------|---|
| connections.max.idle.ms | Specifies the maximum retention period for idle connections. | 600,000 | If the idle connection time is greater than this parameter value, this connection is disconnected. If necessary, a new connection is created. |
| request.timeout.ms | Specifies the timeout for consumer requests. | 30,000 | If the request times out, the request will fail and be sent again. |

15.9 Safety Instructions on Using Kafka

This section applies to MRS 3.x or later.

Brief Introduction to Kafka APIs

- **Producer API**
Indicates the API defined in **org.apache.kafka.clients.producer.KafkaProducer**. When **kafka-console-producer.sh** is used, the API is used by default.
- **Consumer API**
Indicates the API defined in **org.apache.kafka.clients.consumer.KafkaConsumer**. When **kafka-console-consumer.sh** is used, the API is used by default.

 **NOTE**

In MRS 3.x or later, Kafka no longer support old Producer or Consumer APIs.

Protocol Description for Accessing Kafka

The protocols used to access Kafka are as follows: PLAINTEXT, SSL, SASL_PLAINTEXT, and SASL_SSL.

When Kafka service is started, the listeners using the PLAINTEXT and SASL_PLAINTEXT protocols are started. You can set **ssl.mode.enable** to **true** in Kafka service configuration to start listeners using SSL and SASL_SSL protocols. The following table describes the four protocols:

| Protocol | Description | Default Port |
|-----------|---|--------------|
| PLAINTEXT | Supports plaintext access without authentication. | 9092 |

| Protocol | Description | Default Port |
|----------------|---|--------------|
| SASL_PLAINTEXT | Supports plaintext access with Kerberos authentication. | 21007 |
| SSL | Supports SSL-encrypted access without authentication. | 9093 |
| SASL_SSL | Supports SSL-encrypted access with Kerberos authentication. | 21009 |

ACL Settings for a Topic

To view and set topic permission information, run the `kafka-acls.sh` script on the Linux client. For details, see [Managing Kafka User Permissions](#).

Use of Kafka APIs in Different Scenarios

- Scenario 1: accessing the topic with an ACL

| Used API | User Group | Client Parameter | Server Parameter | Accessed Port |
|----------|--|---|--|---|
| API | Users need to meet one of the following conditions: <ul style="list-style-type: none"> Assigned the System Administrator role In the kafkaadmin group In the kafka-supervisor group In the kafka group and be authorized | security.inter.broker.protocol=SASL_PLAINTEXT
sasl.kerberos.service.name = kafka | - | sasl.port
(The default number is 21007.) |
| | | security.protocol=SASL_SSL
sasl.kerberos.service.name = kafka | Set ssl.mode.enabled to true . | sasl-ssl.port
(The default number is 21009.) |

- Scenario 2: accessing the topic without an ACL

| Used API | User Group | Client Parameter | Server Parameter | Accessed Port |
|----------|--|--|--|---|
| API | <p>Users need to meet one of the following conditions:</p> <ul style="list-style-type: none"> Assigned the System_administrator role In the kafkaadmin group In the kafkasuperuser group | <p>security.protocol=SASL_PLAINTEXT
sasl.kerberos.service.name = kafka</p> | - | sasl.port
(The default number is 21007.) |
| | <p>Users are in the kafka group.</p> | | <p>Set allow.everyone.if.no.acl.found to true.</p> <p>NOTE
In normal mode, the server parameter allow.everyone.if.no.acl.found does not need to be modified.</p> | sasl.port
(The default number is 21007.) |
| | <p>Users need to meet one of the following conditions:</p> <ul style="list-style-type: none"> Assigned the System_administrator role In the kafkaadmin group In the kafkasuperuser group | <p>security.protocol=SASL_SSL
sasl.kerberos.service.name = kafka</p> | <p>Set ssl.mode.enable to true.</p> | sasl-ssl.port
(The default number is 21009.) |

| Used API | User Group | Client Parameter | Server Parameter | Accessed Port |
|----------|--------------------------------------|-----------------------------|---|--|
| | Users are in the kafka group. | | <ol style="list-style-type: none"> 1. Set allow.everyone.if.no.acl.found to true. 2. Set ssl.mode.enable to true. | sasl-ssl.port (The default number is 21009.) |
| | - | security.protocol=PLAINTEXT | Set allow.everyone.if.no.acl.found to true . | port (The default number is 9092.) |
| | - | security.protocol=SSL | <ol style="list-style-type: none"> 1. Set allow.everyone.if.no.acl.found to true. 2. Set ssl.mode.enable to true. | ssl.port (The default number is 9063.) |

15.10 Kafka Specifications

This section applies to MRS 3.x or later.

Upper Limit of Topics

The maximum number of topics depends on the number of file handles (mainly used by data and index files on site) opened in the process.

1. Run the **ulimit -n** command to view the maximum number of file handles that can be opened in the process.
2. Run the **lsof -p <Kafka PID>** command to view the file handles (which may keep increasing) that are opened in the Kafka process on the current single node.
3. Determine whether the maximum number of file handles will be reached and whether the running of Kafka is affected after required topics are created, and estimate the maximum size of data that each partition folder can store and the number of data (*.log file, whose default size is 1 GB and can be adjusted by modifying **log.segment.bytes**) and index (*.index file, whose default size is 10 MB and can be adjusted by modifying **log.index.size.max.bytes**) files that will be produced after required topics are created.

Number of Concurrent Consumers

In an application, it is recommended that the number of concurrent consumers in a group be the same as the number of partitions in a topic, ensuring that a consumer consumes data in only a specified partition. If the number of concurrent consumers is more than the number of partitions, the redundant consumers have no data to consume.

Relationship Between Topic and Partition

- If K Kafka nodes are deployed in the cluster, each node is configured with N disks, the size of each disk is M , the cluster contains n topics (named as T_1, T_2, \dots, T_n), the data input traffic per second of the m topic is $X(T_m)$ MB/s, the number of configured replicas is $R(T_m)$, and the configured data retention time is $Y(T_m)$ hour, the following requirement must be met:

$$M \times N \times K > \sum_{i=T_1}^{T_n} (X(i)R(i)Y(i) \times 3600)$$

- If the size of a disk is M , the disk has n partitions (named as P_0, P_1, \dots, P_n), the data write traffic per second of the m partition is $Q(P_m)$ MB/s (calculation method: data traffic of the topic to which the m partition belongs divided by the number of partitions), and the data retention time is $T(P_m)$ hours, the following requirement must be met for the disk:

$$M > \sum_{i=P_0}^{P_n} (Q(i)T(i) \times 3600)$$

- Based on the throughput, if the throughput that can be reached by the producer is P , the throughput that can be reached by the consumer is C , and the expected throughput of Kafka is T , it is recommended that the number of partitions of the topic be set to $\text{Max}(T/P, T/C)$.

NOTE

- In a Kafka cluster, more partitions mean higher throughput. However, too many partitions also pose potential impacts, such as a file handle increase, unavailability increase (for example, if a node is faulty, the time window becomes large after the leader is reselected in some partitions), and end-to-end latency increase.
- Suggestion: The disk usage of a partition is smaller than or equal to 100 GB; the number of partitions on a node is smaller than or equal to 3,000; the number of partitions in the entire cluster is smaller than or equal to 10,000.

15.11 Using the Kafka Client

Scenario

This section guides users to use a Kafka client in an O&M or service scenario.

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

Prerequisites

- The client has been installed. For example, the installation directory is **/opt/client**.
- Service component users have been created by the MRS cluster administrator. Machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login. (Not involved in normal mode)
- After changing the domain name of a cluster, redownload the client to ensure that the **kerberos.domain.name** value in the configuration file of the client is set to the correct server domain name.

Procedure

Step 1 Log in to the node where the client is installed as the client installation user.

Step 2 Run the following command to go to the client installation directory:

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 Run the following command to perform user authentication (skip this step in normal mode):

```
kinit Component service user
```

Step 5 Run the following command to switch to the Kafka client installation directory:

```
cd Kafka/kafka/bin
```

Step 6 Run the following command to use the client tool to view and use the help information:

- **./kafka-console-consumer.sh**: Kafka message reading tool
- **./kafka-console-producer.sh**: Kafka message publishing tool
- **./kafka-topics.sh**: Kafka topic management tool

----End

15.12 Configuring Kafka HA and High Reliability Parameters

Scenario

For the Kafka message transmission assurance mechanism, different parameters are available for meeting different performance and reliability requirements. This section describes how to configure Kafka high availability (HA) and high reliability parameters.

This section applies to MRS 3.x or later.

Impact on the System

- Impact of HA and high performance configurations:

NOTICE

After HA and high performance are configured, the data reliability decreases. Specifically, data may be lost if disks or nodes are faulty.

- Impact of high reliability configurations:
 - Deteriorated performance
If **ack** is set to **-1**, data written is considered as successful only when data is written to multiple replicas. As a result, the delay of a single message increases and the client processing capability decreases. The impact is subject to the actual test data.
 - Reduced availability
A replica that is not in the ISR list cannot be elected as a leader. If the leader goes offline and other replicas are not in the ISR list, the partition remains unavailable until the leader node recovers. When the node where a replica of a partition is located is faulty, the minimum number of successful replicas cannot be met. As a result, service writing fails.
- If parameters are at the service level, Kafka needs to be restarted. You are advised to modify the service-level configuration in the change window.

Parameter Description

- If services require high availability and high performance, set the parameters listed in [Table 15-8](#) on the server. For details about the parameter configuration entry, see [Modifying Cluster Service Configuration Parameters](#).

Table 15-8 Server HA and high performance parameters

| Parameter | Default Value | Description |
|--------------------------------|---------------|--|
| unclean.leader.election.enable | true | Specifies whether a replica that is not in the ISR can be selected as the leader. If this parameter is set to true , data may be lost. |
| auto.leader.rebalance.enable | true | Specifies whether the leader automated balancing function is used.
If this parameter is set to true , the controller periodically balances the leader of each partition on all nodes and assigns the leader to a replica with a higher priority. |

| Parameter | Default Value | Description |
|---------------------|---------------|--|
| min.insync.replicas | 1 | Specifies the minimum number of replicas to which data is written when acks is set to -1 for the Producer. |

Set the parameters listed in [Table 15-9](#) in the client configuration file **producer.properties**. The path for storing **producer.properties** is **/opt/client/Kafka/kafka/config/producer.properties**, where **/opt/client** indicates the installation directory of the Kafka client.

Table 15-9 Client HA and high performance parameters

| Parameter | Default Value | Description |
|-----------|---------------|--|
| acks | 1 | <p>The leader needs to check whether the message has been received and determine whether the required operation has been processed. This parameter affects message reliability and performance.</p> <ul style="list-style-type: none"> • If this parameter is set to 0, the producer does not wait for any response from the server, and the message is considered successful. • If this parameter is set to 1, when the leader of the replica verifies that data has been written into the cluster, the leader returns a response without waiting for data to be written to all replicas. In this case, if the leader is abnormal when the leader makes the confirmation but replica synchronization is not complete, data will be lost. • If this parameter is set to -1, the message is considered to be successfully received only when all synchronized replicas are confirmed. If the min.insync.replicas |

| Parameter | Default Value | Description |
|-----------|---------------|---|
| | | parameter is also configured, data can be written into multiple replicas. In this case, records will not be lost as long as one replica remains active. |

- To ensure high data reliability for services, set the parameters listed in [Table 15-10](#) on the server. For details about the parameter configuration entry, see [Modifying Cluster Service Configuration Parameters](#).

Table 15-10 Server HA parameters

| Parameter | Recommended Value | Description |
|--------------------------------|-------------------|--|
| unclean.leader.election.enable | false | A replica that is not in the ISR list cannot be elected as a leader. |
| min.insync.replicas | 2 | Specifies the minimum number of replicas to which data is written when acks is set to -1 for the Producer.
Ensure that the value of min.insync.replicas is equal to or less than that of replication.factor . |

Set the parameters listed in [Table 15-11](#) in the client configuration file **producer.properties**. The path for storing **producer.properties** is **/opt/client/Kafka/kafka/config/producer.properties**, where **/opt/client** indicates the installation directory of the Kafka client.

Table 15-11 Server HA parameters

| Parameter | Recommended Value | Description |
|-----------|-------------------|--|
| acks | -1 | <p>The leader needs to check whether the message has been received and determine whether the required operation has been processed.</p> <p>If this parameter is set to -1, the message is considered to be successfully received only when all replicas in the ISR list have confirmed to receive the message. This parameter is used along with min.insync.replicas to ensure that multiple copies are successfully written. As long as one copy is active, the record will not be lost. If this parameter is set to -1, the production performance deteriorates. Therefore, you need to set this parameter based on the actual situation.</p> |

Configuration Suggestions

Configure parameters based on requirements on reliability and performance in the following service scenarios:

- For valued data, you are advised to configure RAID1 or RAID5 for Kafka data directory disks to improve data reliability when a single disk is faulty.
- For parameters that can be modified at the topic level, the service level configurations are used by default.

These parameters can be separately configured based on topic reliability requirements. For example, log in to the Kafka client as user **root**, and run the following command to configure the reliability parameter with topic named test in the client installation directory:

```
cd Kafka/kafka/bin
```

```
kafka-topics.sh --zookeeper 192.168.1.205:2181/kafka --alter --topic test
--config unclean.leader.election.enable=false --config
min.insync.replicas=2
```

192.168.1.205 indicates the ZooKeeper service IP address.

- If parameters are at the service level, Kafka needs to be restarted. You are advised to modify the service-level configuration in the change window.

15.13 Changing the Broker Storage Directory

Scenario

This section applies to MRS 3.x or later.

When a broker storage directory is added, the MRS cluster administrator needs to change the broker storage directory on FusionInsight Manager, to ensure that the Kafka can work properly. The new topic partition will be generated in the directory that has fewest partitions. Changing the ZooKeeper storage directory includes the following scenarios:

NOTE

Because Kafka does not detect disk capacity, ensure that the disk quantity and capacity configured for each Broker instance are the same.

- Change the storage directory of the Broker role. In this way, the storage directories of all Broker instances are changed.
- Change the storage directory of a single Broker instance. In this way, only the storage directory of this Broker instance is changed, and the storage directories of other Broker instances remain the same.

Impact on the System

- Changing the Broker role storage directory requires the restart of services. The services cannot be accessed during the restart.
- The storage directory of a single Broker instance can be changed only after the instance is restarted. The instance cannot provide services during the restart.
- The directory for storing service parameter configurations must also be updated.

Prerequisites

- New disks have been prepared and installed on each data node, and the disks are formatted.
- The Kafka client has been installed.
- When you change the storage directory of a single Broker instance, the number of active Broker instances must be greater than the number of backups specified during topic creation.

Procedure

Changing the storage directory of the Kafka role

- Step 1** Log in as user **root** to each node on which the Kafka service is installed, and perform the following operations:

1. Create a target directory.
For example, to create the target directory `${BIGDATA_DATA_HOME}/kafka/data2`, run the following command:

```
mkdir ${BIGDATA_DATA_HOME}/kafka/data2
```
2. Mount the directory to the new disk. For example, mount `${BIGDATA_DATA_HOME}/kafka/data2` to the new disk.
3. Modify permissions on the new directory.
For example, to modify permissions on the `${BIGDATA_DATA_HOME}/kafka/data2` directory, run the following commands:

```
chmod 700 ${BIGDATA_DATA_HOME}/kafka/data2 -R and chown omm:wheel ${BIGDATA_DATA_HOME}/kafka/data2 -R
```

Step 2 Log in to FusionInsight Manager for clusters of MRS 3.x or later and choose **Cluster > Services > Kafka > Configurations**.

Step 3 Add a new directory to the end of the default value of **log.dirs**.

Enter **log.dirs** in the search box and add the new directory to the end of the default value of the **log.dirs** configuration item. Use commas (,) to separate multiple directories. For example:

```
${BIGDATA_DATA_HOME}/kafka/data1/kafka-logs,${BIGDATA_DATA_HOME}/kafka/data2/kafka-logs
```

Step 4 Click **Save**, and then click **OK**. When **Operation succeeded** is displayed, click **Finish**.

Step 5 Choose **Cluster > Services > Kafka**. In the upper right corner, choose **More > Restart Service** to restart the Kafka service.

Changing the storage directory of a single Kafka instance

Step 6 Log in to the Broker node as user **root** and perform the following operations:

1. Create a target directory.
For example, to create the target directory `${BIGDATA_DATA_HOME}/kafka/data2`, run the following command:

```
mkdir ${BIGDATA_DATA_HOME}/kafka/data2
```
2. Mount the directory to the new disk. For example, mount `${BIGDATA_DATA_HOME}/kafka/data2` to the new disk.
3. Modify permissions on the new directory.
For example, to modify permissions on the `${BIGDATA_DATA_HOME}/kafka/data2` directory, run the following commands:

```
chmod 700 ${BIGDATA_DATA_HOME}/kafka/data2 -R and chown omm:wheel ${BIGDATA_DATA_HOME}/kafka/data2 -R
```

Step 7 Log in to FusionInsight Manager for MRS 3.x or later, and choose **Cluster > Services > Kafka > Instance**.

Step 8 Click the specified broker instance and switch to **Instance Configurations**.

Enter **log.dirs** in the search box and add the new directory to the end of the default value of the **log.dirs** configuration item. Use commas (,) to separate

multiple directories, for example, `${BIGDATA_DATA_HOME}/kafka/data1/kafka-logs`, `${BIGDATA_DATA_HOME}/kafka/data2/kafka-logs`.

Step 9 Click **Save**, and then click **OK**. A message is displayed, indicating that the operation is successful. Click **Finish**.

Step 10 On the Broker instance page, choose **More > Restart Instance** to restart the Broker instance.

----End

15.14 Checking the Consumption Status of Consumer Group

Scenario

View the current consumption information on the client as an MRS cluster administrator based on service requirements.

This section applies to MRS 3.x or later.

Prerequisites

- The MRS cluster administrator has understood service requirements and prepared a system user.
- The Kafka client has been installed.

Procedure

Step 1 Log in as a client installation user to the node on which the Kafka client is installed.

Step 2 Switch to the Kafka client installation directory, for example, `/opt/client`.

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 Run the following command to perform user authentication (skip this step in normal mode):

```
kinit Component service user
```

Step 5 Run the following command to switch to the Kafka client installation directory:

```
cd Kafka/kafka/bin
```

Step 6 Run the `kafka-consumer-groups.sh` command to check the current consumption status.

- Check the Consumer Group list on Kafka saved by Offset:

```
./kafka-consumer-groups.sh --list --bootstrap-server <Service IP address of any node where a broker instances is located>.Port number of the Kafka cluster> --command-config ../config/consumer.properties
```


Example: `./kafka-consumer-groups.sh --bootstrap-server 192.168.1.1:21007 --list --command-config ../config/consumer.properties`

- Check the consumption status of Consumer Group on Kafka saved by Offset:
`./kafka-consumer-groups.sh --describe --bootstrap-server <Service IP address of any node where a broker instances is located>:Port number of the Kafka cluster> --group Consumer group name --command-config ../config/consumer.properties`

Example: `./kafka-consumer-groups.sh --describe --bootstrap-server 192.168.1.1:21007 --group example-group --command-config ../config/consumer.properties`

NOTICE

1. Ensure that the current consumer is online and consumes data.
2. Configure the **group.id** in the **consumer.properties** configuration file and **--group** in the command to the group to be queried.
3. The Kafka cluster's IP port number is 21007 in security mode and 9092 in normal mode.

----End

15.15 Kafka Balancing Tool Instructions

Scenario

This section describes how to use the Kafka balancing tool on a client to balance the load of the Kafka cluster based on service requirements in scenarios such as node decommissioning, node recommissioning, and load balancing.

This section applies to MRS 3.x or later. For versions earlier than MRS 3.x, see [Balancing Data After Kafka Node Scale-Out](#).

Prerequisites

- The MRS cluster administrator has understood service requirements and prepared a Kafka administrator (belonging to the **kafkaadmin** group. It is not required for the normal mode.).
- The Kafka client has been installed.

Procedure

- Step 1** Log in as a client installation user to the node on which the Kafka client is installed.
- Step 2** Switch to the Kafka client installation directory, for example, `/opt/client`.
`cd /opt/client`
- Step 3** Run the following command to configure environment variables:
`source bigdata_env`

Step 4 Run the following command to authenticate the user (skip this step in normal mode):

```
kinit Component service user
```

Step 5 Run the following command to switch to the Kafka client installation directory:

```
cd Kafka/kafka
```

Step 6 Run the **kafka-balancer.sh** command to balance user cluster. The commonly used commands are:

- Run the **--run** command to perform cluster balancing:

```
./bin/kafka-balancer.sh --run --zookeeper <ZooKeeper service IP address of any ZooKeeper node:zkPort/kafka> --bootstrap-server <Kafka cluster IP:port> --throttle 1000000 --consumer-config config/consumer.properties --enable-az-aware --show-details
```

This command consists of generation and execution of the balancing solution. **--show-details** is optional, indicating whether to print the solution details. **--throttle** indicates the bandwidth limit during the execution of the balancing solution. The unit is bytes per second (bytes/sec). **--enable-az-aware** indicates that the cross-AZ feature is enabled when the balancing solution is generated. When this parameter is used, ensure that the cross-AZ feature has been enabled for the cluster.

- Run the **--run** command to decommission a node:

```
./bin/kafka-balancer.sh --run --zookeeper <Service IP address of any ZooKeeper node:zkPort/kafka> --bootstrap-server <Kafka cluster IP address:port> --throttle 1000000 --consumer-config config/consumer.properties --remove-brokers <BrokerId list> --enable-az-aware --force
```

In the command, **--remove-brokers** indicates the list of broker IDs to be deleted. Multiple broker IDs are separated by commas (.). **--force** is optional, indicating that the disk usage alarm is ignored and the migration solution is forcibly generated. **--enable-az-aware** is optional, indicating that the cross-AZ feature is enabled when the balancing solution is generated. When this parameter is used, ensure that the cross-AZ feature has been enabled for the cluster.

NOTE

This command migrates data on the Broker nodes to be decommissioned to other Broker nodes.

- Run the following command to view the execution status:

```
./bin/kafka-balancer.sh --status --zookeeper <Service IP address of any ZooKeeper node:zkPort/kafka>
```

- Run the following command to generate a balancing solution:

```
./bin/kafka-balancer.sh --generate --zookeeper <Service IP address of any ZooKeeper node:zkPort/kafka> --bootstrap-server <Kafka cluster IP address:port> --consumer-config config/consumer.properties --enable-az-aware
```

This command is used to generate a migration solution based on the current cluster status and print the solution to the console. **--enable-az-aware** is optional, indicating that the cross-AZ feature is enabled when a migration solution is generated. If this parameter is used, ensure that the cross-AZ feature has been enabled for the cluster.

- Clearing the intermediate status
`./bin/kafka-balancer.sh --clean --zookeeper <Service IP address of any ZooKeeper node:zkPort/kafka>`

This command is used to clear the intermediate status information on the ZooKeeper when the migration is not complete.

NOTICE

The port number of the Kafka cluster's IP address is 21007 in security mode and 9092 in normal mode.

----End

Troubleshooting

During partition migration using the Kafka balancing tool, if the execution progress of the balancing tool is blocked due to a Broker fault in the cluster, you need to manually rectify the fault. The scenarios are as follows:

- The Broker is faulty because the disk usage reaches 100%.
 - a. Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**, stop the Broker instance in the **Restoring** state, and record the management IP address of the node where the instance resides and the corresponding **broker.id**. You can click the role name to view the value, on the **Instance Configurations** page, select **All Configurations** and search for the **broker.id** parameter.
 - b. 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`.
 - c. Go to the directory, run the **du -sh *** command to view the size of each file in the directory, Check whether files other than files in the **kafka-logs** directory exist, and determine whether these files can be deleted or migrated.
 - If yes, delete or migrate the related data and go to **8**.
 - If no, go to **4**.
 - d. 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.
 - e. Modify the **recovery-point-offset-checkpoint** and **replication-offset-checkpoint** files in the **kafka-logs** directory in the same way.
 - i. 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.)
 - ii. 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.)

- f. 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 5.)
- g. 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.
- h. Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance** to start the stopped Broker instance.
- i. Wait for 5 to 10 minutes and check whether the health status of the Broker instance is **Good**.
 - If yes, resolve the disk capacity insufficiency problem according to the handling method of "ALM-38001 Insufficient Kafka Disk Capacity" after the alarm is cleared.
 - If no, contact O&M support.

After the faulty Broker is recovered, the blocked balancing task continues. You can run the **--status** command to view the task execution progress.

- The Broker fault occurs because of other causes, the fault scenario is clear, and the fault can be rectified within a short period of time.
 - a. Restore the faulty Broker according to the root cause.
 - b. After the faulty Broker is recovered, the blocked balancing task continues. You can run the **--status** command to view the task execution progress.
- The Broker fault occurs because of other causes, the fault scenario is complex, and the fault cannot be rectified within a short period of time.
 - a. Run the **kinit Kafka administrator account** command (skip this step in normal mode).
 - b. Run the **zkCli.sh -server <ZooKeeper cluster service IP address.zkPort/ kafka>** command to log in to ZooKeeper Shell.
 - c. Run the **addauth krbgroup** command (skip this step in normal mode).
 - d. Delete the **/admin/reassign_partitions** and **/controller** directories.
 - e. Perform the preceding steps to forcibly stop the migration. After the cluster recovers, run the **kafka-reassign-partitions.sh** command to delete redundant copies generated during the intermediate process.

15.16 Balancing Data After Kafka Node Scale-Out

Scenario

This section describes how to use the Kafka balancing tool on the client to balance the load of the Kafka cluster after Kafka nodes are scaled out.

This section applies to versions earlier than MRS 3.x. For MRS 3.x or later, see [Kafka Balancing Tool Instructions](#).

Prerequisites

- The MRS cluster administrator has understood service requirements and prepared a Kafka administrator (belonging to the **kafkaadmin** group. It is not required for the normal mode.).
- The Kafka client has been installed in a directory, for example, **/opt/client**.
- Two topics named **test_2** and **test_3** has been created by referring to [Step 7](#). The **move-kafka-topic.json** file has been created in the **/opt/client/Kafka/kafka** directory. The topic format is as follows:

```
{
  "topics":
  [{"topic":"test_2"}, {"topic":"test_3"}],
  "version":1
}
```

Procedure

Step 1 Log in to the node where the Kafka client is installed as the client installation user.

Step 2 Run the following command to switch to the client installation directory:

```
cd /opt/client
```

Step 3 Run the following command to set environment variables:

```
source bigdata_env
```

Step 4 Run the following command to perform user authentication (skip this step if the cluster is in normal mode):

```
kinit Component service user
```

Step 5 Run the following command to go to the **bin** directory of the Kafka client:

```
cd Kafka/kafka/bin
```

Step 6 Run the following command to generate an execution plan:

```
./kafka-reassign-partitions.sh --zookeeper 172.16.0.119:2181/kafka --topics-to-move-json-file ../move-kafka-topic.json --broker-list "1,2,3" --generate
```

NOTE

- **172.16.0.119**: service IP address of the ZooKeeper instance
- **--broker-list "1,2,3"**: list of broker instances. **1,2,3** indicates all broker IDs after a scale-out.

```
[root@node-master1SPXC bin]# ./kafka-reassign-partitions.sh --zookeeper 172.16.0.119:2181/kafka --topics-to-move-json-file ../m
ve-kafka-topic.json --broker-list "1,2,3" --generate
Current partition replica assignment
{"version":1,"partitions":[{"topic":"test_2","partition":3,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partitio
n":4,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_3","partition":5,"replicas":[2,1],"log_dirs":["any","any"]}, {"topi
c":"test_3","partition":3,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_2","partition":2,"replicas":[2,1],"log_dirs":
["any","any"]}, {"topic":"test_3","partition":0,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_3","partition":2,"replic
as":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partition":6,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_3"
,"partition":4,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partition":0,"replicas":[2,1],"log_dirs":["any","any
"]}, {"topic":"test_3","partition":1,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_2","partition":1,"replicas":[1,2],"l
og_dirs":["any","any"]}, {"topic":"test_2","partition":5,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_3","partition":
6,"replicas":[1,2],"log_dirs":["any","any"]}]}
Proposed partition reassignment configuration
{"version":1,"partitions":[{"topic":"test_3","partition":0,"replicas":[2,3],"log_dirs":["any","any"]}, {"topic":"test_2","partiti
on":1,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partition":6,"replicas":[3,1],"log_dirs":["any","any"]}, {"topi
c":"test_3","partition":2,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partition":3,"replicas":[3,2],"log_dirs":
["any","any"]}, {"topic":"test_3","partition":5,"replicas":[1,3],"log_dirs":["any","any"]}, {"topic":"test_2","partition":0,"replic
as":[3,1],"log_dirs":["any","any"]}, {"topic":"test_2","partition":5,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_3"
,"partition":4,"replicas":[3,2],"log_dirs":["any","any"]}, {"topic":"test_2","partition":2,"replicas":[2,3],"log_dirs":["any","any
"]}, {"topic":"test_3","partition":1,"replicas":[3,1],"log_dirs":["any","any"]}, {"topic":"test_3","partition":6,"replicas":[2,3],"l
og_dirs":["any","any"]}, {"topic":"test_3","partition":3,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_2","partiti
on":4,"replicas":[1,3],"log_dirs":["any","any"]}]}
[root@node-master1SPXC bin]#
```

Step 7 Run the `vim ../reassignment.json` command to create the `reassignment.json` file and save it to the `/opt/kafkaclient/Kafka/kafka` directory.

Copy the content under **Proposed partition reassignment configuration** generated in **Step 6** to the `reassignment.json` file, as shown in the follows:

```
{"version":1,"partitions":[{"topic":"test"}, {"topic":"test","partition":4,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test","partition":1,"replicas":[1,3],"log_dirs":["any","any"]}, {"topic":"test","partition":3,"replicas":[3,1],"log_dirs":["any","any"]}, {"topic":"test","partition":0,"replicas":[3,2],"log_dirs":["any","any"]}, {"topic":"test","partition":2,"replicas":[2,1],"log_dirs":["any","any"]}]}

```

Step 8 Run the following command to redistribute partitions:

```
./kafka-reassign-partitions.sh --zookeeper 172.16.0.119:2181/kafka --reassignment-json-file ../reassignment.json --execute --throttle 50000000
```

NOTE

`--throttle 50000000`: The maximum bandwidth is 50 MB/s. You can change the bandwidth based on the data volume and the customer's requirements on the balancing time. If the data volume is 5 TB, the bandwidth is 50 MB/s and the data balancing takes about 8 hours.

```
[root@node-master1SPXC bin]# vim ../reassignment.json
[root@node-master1SPXC bin]# ./kafka-reassign-partitions.sh --zookeeper 172.16.0.119:2181/kafka --reassignment-json-file ../reas
signment.json --execute --throttle 50000000
Current partition replica assignment
{"version":1,"partitions":[{"topic":"test_2","partition":3,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partitio
n":4,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_3","partition":5,"replicas":[2,1],"log_dirs":["any","any"]}, {"topi
c":"test_3","partition":3,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_2","partition":2,"replicas":[2,1],"log_dirs":
["any","any"]}, {"topic":"test_3","partition":0,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_3","partition":2,"replic
as":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partition":6,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_3"
,"partition":4,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_2","partition":0,"replicas":[2,1],"log_dirs":["any","any
"]}, {"topic":"test_3","partition":1,"replicas":[2,1],"log_dirs":["any","any"]}, {"topic":"test_2","partition":1,"replicas":[1,2],"l
og_dirs":["any","any"]}, {"topic":"test_2","partition":5,"replicas":[1,2],"log_dirs":["any","any"]}, {"topic":"test_3","partition":
6,"replicas":[1,2],"log_dirs":["any","any"]}]}
Save this to use as the --reassignment-json-file option during rollback
Warning: You must run Verify periodically, until the reassignment completes, to ensure the throttle is removed. You can also alter
the throttle by rerunning the Execute command passing a new value.
The inter-broker throttle limit was set to 50000000 B/s
Successfully started reassignment of partitions.
[root@node-master1SPXC bin]#
```

Step 9 Run the following command to check the data migration status:

```
./kafka-reassign-partitions.sh --zookeeper 172.16.0.119:2181/kafka --reassignment-json-file ../reassignment.json --verify
```

```

ds are typed to all terminals (use Ctrl+Shift+Insert to paste)
Multi-paste  Exit multi-execution mode
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-5
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-6
[root@node-str-coreRuzk0001 kafka-logs]# ll
total 56
-rw----- 1 omm wheel 4 Sep 14 21:30 cleaner-offset-check
-rw----- 1 omm wheel 4 Sep 14 21:31 log-start-offset-che
-rw----- 1 omm wheel 54 Sep 14 19:39 meta.properties
-rw----- 1 omm wheel 103 Sep 14 21:31 recovery-point-offse
-rw----- 1 omm wheel 103 Sep 14 21:32 replication-offset-c
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-0
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-1
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-2
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-3
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-4
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-5
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-6
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-1
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-2
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-3
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-4
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-5
[root@node-str-coreRuzk0001 kafka-logs]#
[Disable this terminal from "MultiExec" mode]

drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-5
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-6
[root@node-str-coreRuzk0002 kafka-logs]# ll
total 56
-rw----- 1 omm wheel 4 Sep 14 21:30 cleaner-offset-check
-rw----- 1 omm wheel 4 Sep 14 21:31 log-start-offset-che
-rw----- 1 omm wheel 54 Sep 14 19:39 meta.properties
-rw----- 1 omm wheel 103 Sep 14 21:31 recovery-point-offse
-rw----- 1 omm wheel 103 Sep 14 21:32 replication-offset-c
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-1
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-2
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-3
drwx----- 2 omm wheel 4096 Sep 14 21:11 test_2-5
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-0
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-2
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-3
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-4
drwx----- 2 omm wheel 4096 Sep 14 21:12 test_3-6
[root@node-str-coreRuzk0002 kafka-logs]#
[Disable this terminal from "MultiExec" mode]

[Disable this terminal from "MultiExec" mode]
[root@node-str-coreacDNo data1]# cd kafka-logs/
[root@node-str-coreacDNo kafka-logs]# ll
total 60
-rw----- 1 omm wheel 4 Sep 14 21:18 cleaner-offset-check
-rw----- 1 omm wheel 4 Sep 14 21:31 log-start-offset-che
-rw----- 1 omm wheel 115 Sep 14 21:31 recovery-point-offse
-rw----- 1 omm wheel 115 Sep 14 21:32 replication-offset-c
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_2-0
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_2-2
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_2-3
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_2-4
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_2-6
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_3-0
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_3-1
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_3-4
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_3-5
drwx----- 2 omm wheel 4096 Sep 14 21:30 test_3-6
[root@node-str-coreacDNo kafka-logs]#
[Disable this terminal from "MultiExec" mode]

[Disable this terminal from "MultiExec" mode]
[root@node-master1SPXC bin]# ./kafka-reassign-partitions.sh --
assignment.json --verify
Status of partition reassignment:
Reassignment of partition test_2-3 completed successfully
Reassignment of partition test_2-4 completed successfully
Reassignment of partition test_3-5 completed successfully
Reassignment of partition test_3-3 completed successfully
Reassignment of partition test_2-2 completed successfully
Reassignment of partition test_3-0 completed successfully
Reassignment of partition test_3-2 completed successfully
Reassignment of partition test_2-6 completed successfully
Reassignment of partition test_3-4 completed successfully
Reassignment of partition test_2-0 completed successfully
Reassignment of partition test_3-1 completed successfully
Reassignment of partition test_2-1 completed successfully
Reassignment of partition test_2-5 completed successfully
Reassignment of partition test_3-6 completed successfully
Throttle was removed.
[root@node-master1SPXC bin]#
[Disable this terminal from "MultiExec" mode]

```

----End

15.17 Kafka Token Authentication Mechanism Tool Usage

Scenario

Operations need to be performed on tokens when the token authentication mechanism is used.

This section applies to Kerberos authentication-enabled clusters of MRS 3.x or later.

Prerequisites

- The MRS cluster administrator has understood service requirements and prepared a system user.
- The Kafka client has been installed.

Procedure

Step 1 Log in as a client installation user to the node on which the Kafka client is installed.

Step 2 Switch to the Kafka client installation directory, for example, `/opt/client`.

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 Run the following command to perform user authentication:

```
kinit Component service user
```

Step 5 Run the following command to switch to the Kafka client installation directory:

```
cd Kafka/kafka/bin
```

Step 6 Use `kafka-delegation-tokens.sh` to perform operations on tokens.

- Generate a token for a user.

```
./kafka-delegation-tokens.sh --create --bootstrap-server <IP1:PORT,
IP2:PORT,...> --max-life-time-period <Long: max life period in milliseconds>
--command-config <config file> --renewer-principal User:<user name>
```

Example: `./kafka-delegation-tokens.sh --create --bootstrap-server 192.168.1.1:21007,192.168.1.2:21007,192.168.1.3:21007 --command-config ../config/producer.properties --max-life-time-period -1 --renewer-principal User:username`

- List information about all tokens of a specified user.

```
./kafka-delegation-tokens.sh --describe --bootstrap-server <IP1:PORT,
IP2:PORT,...> --command-config <config file> --owner-principal User:<user
name>
```

Example: `./kafka-delegation-tokens.sh --describe --bootstrap-server 192.168.1.1:21007,192.168.1.2:21007,192.168.1.3:21007 --command-config ../config/producer.properties --owner-principal User:username`

- Update the token validity period.

```
./kafka-delegation-tokens.sh --renew --bootstrap-server <IP1:PORT,
IP2:PORT,...> --renew-time-period <Long: renew time period in milliseconds>
--command-config <config file> --hmac <String: HMAC of the delegation
token>
```

Example: `./kafka-delegation-tokens.sh --renew --bootstrap-server 192.168.1.1:21007,192.168.1.2:21007,192.168.1.3:21007 --renew-time-period -1 --command-config ../config/producer.properties --hmac ABCDEFG`

- Destroy a token.

```
./kafka-delegation-tokens.sh --expire --bootstrap-server <IP1:PORT,
IP2:PORT,...> --expiry-time-period <Long: expiry time period in milliseconds>
--command-config <config file> --hmac <String: HMAC of the delegation
token>
```

Example: `./kafka-delegation-tokens.sh --expire --bootstrap-server 192.168.1.1:21007,192.168.1.2:21007,192.168.1.3:21007 --expiry-time-period -1 --command-config ../config/producer.properties --hmac ABCDEFG`

----End

15.18 Introduction to Kafka Logs

This section applies to MRS 3.x or later.

Log Description

Log paths: The default storage path of Kafka logs is `/var/log/Bigdata/kafka`. The default storage path of audit logs is `/var/log/Bigdata/audit/kafka`.

- Broker: `/var/log/Bigdata/kafka/broker` (run logs)

Log archive rule: The automatic Kafka log compression function is enabled. By default, when the size of logs exceeds 30 MB, logs are automatically compressed into a log file named in the following format: `<Original log file name>-<yyyy-mm-dd_hh-mm-ss>.[ID].log.zip`. A maximum of 20 latest compressed files are retained by default. You can configure the number of compressed files and the compression threshold.

Table 15-12 Kafka log list

| Type | Log File Name | Description |
|----------------|--|--|
| Run log | server.log | Server run log of the broker process |
| | controller.log | Controller run log of the broker process |
| | kafka-request.log | Request run log of the broker process |
| | log-cleaner.log | Cleaner run log of the broker process |
| | state-change.log | State-change run log of the broker process |
| | kafkaServer-<SSH_USER>-<DATE>-<PID>-gc.log | GC log of the broker process |
| | postinstall.log | Work log after broker installation |
| | prestart.log | Work log before broker startup |
| | checkService.log | Log that records whether broker starts successfully |
| | start.log | Startup log of the broker process |
| | stop.log | Stop log of the broker process |
| | checkavailable.log | Log that records the health check details of the Kafka service |
| | checkInstanceHealth.log | Log that records the health check details of broker instances |
| | kafka-authorizer.log | Broker authorization log |
| kafka-root.log | Broker basic log | |

| Type | Log File Name | Description |
|------|--------------------------------|---|
| | cleanup.log | Cleanup log of broker uninstallation |
| | metadata-backup-recovery.log | Broker backup and recovery log |
| | ranger-kafka-plugin-enable.log | Log that records the Ranger plug-ins enabled by brokers |
| | server.out | Broker JVM log |
| | audit.log | Authentication log of the Ranger authentication plug-in. This log is archived in the <code>/var/log/Bigdata/audit/kafka</code> directory. |

Log Level

Table 15-13 describes the log levels supported by Kafka.

Levels of run logs are ERROR, WARN, INFO, and DEBUG from the highest to the lowest priority. Run logs of equal or higher levels are recorded. The higher the specified log level, the fewer the logs recorded.

Table 15-13 Log levels

| Level | Description |
|-------|--|
| ERROR | Logs of this level record error information about system running. |
| WARN | Logs of this level record exception information about the current event processing. |
| INFO | Logs of this level record normal running status information about the system and events. |
| DEBUG | Logs of this level record the system information and system debugging information. |

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page. See [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level.

Step 4 Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.

----End

Log Format

The following table describes the Kafka log format.

Table 15-14 Log formats

| Type | Format | Example |
|---------|---|---|
| Run log | <yyyy-MM-dd
HH:mm:ss,SSS> <Log
Level> <Thread that
generates the log>
<Message in the log>
<Full name of the log
event invocation
class>(<Log file>:<Row>) | 2015-08-08 11:09:53,483
 INFO [main] Loading
logs.
kafka.log.LogManager
(Logging.scala:68) |
| | <yyyy-MM-dd
HH:mm:ss><HostName>
<Component
name><logLevel><Messa
ge> | 2015-08-08 11:09:51
10-165-0-83 Kafka INFO
Running kafka-start.sh. |

15.19 Performance Tuning

15.19.1 Kafka Performance Tuning

Scenario

You can modify Kafka server parameters to improve Kafka processing capabilities in specific service scenarios.

Parameter Tuning

Modify the service configuration parameters. For details, see [Modifying Cluster Service Configuration Parameters](#). For details about the tuning parameters, see [Table 15-15](#).

Table 15-15 Tuning parameters

| Parameter | Default Value | Scenario |
|-----------------------------------|---------------|--|
| num.recovery.threads.per.data.dir | 10 | During the Kafka startup process, if a large volume of data exists, you can increase the value of this parameter to accelerate the startup. |
| background.threads | 10 | Specifies the number of threads processed by a broker background task. If a large volume of data exists, you can increase the value of this parameter to improve broker processing capabilities. |
| num.replica.fetchers | 1 | Specifies the number of threads used when a replica requests to the Leader for data synchronization. If the value of this parameter is increased, the replica I/O concurrency increases. |
| num.io.threads | 8 | Specifies the number of threads used by the broker to process disk I/O. It is recommended that the number of threads be greater than or equal to the number of disks. |
| KAFKA_HEAP_OPTS | -Xmx6G -Xms6G | Specifies the Kafka JVM heap memory setting. If the data volume on the broker is large, adjust the heap memory size. |

15.20 Kafka Feature Description

Kafka Idempotent Feature

Feature description: The function of creating idempotent producers is introduced in Kafka 0.11.0.0. After this function is enabled, producers are automatically upgraded to idempotent producers. When producers send messages with the same field values, brokers automatically detect whether the messages are duplicate to avoid duplicate data. Note that this feature can only ensure idempotence in a single partition. That is, an idempotent producer can ensure that no duplicate messages exist in a partition of a topic. Only idempotence on a single session can be implemented. The session refers to the running of the producer process. That is, idempotence cannot be ensured after the producer process is restarted.

Method for enabling this feature:

1. Add `props.put("enable.idempotence", true)` to the secondary development code.
2. Add `enable.idempotence = true` to the client configuration file.

Kafka Transaction Feature

Feature description: Kafka 0.11 introduces the transaction feature. The Kafka transaction feature indicates that a series of producer message production and consumer offset submission operations are in the same transaction, or are regarded as an atomic operation. Message production and offset submission succeed or fail at the same time. This feature provides transactions at the Read Committed isolation level to ensure that multiple messages are written to the target partition atomically and that the consumer can view only the transaction messages that are successfully submitted. The transaction feature of Kafka is used in the following scenarios:

1. Multiple pieces of data sent by a producer can be encapsulated in a transaction to form an atomic operation. All messages are successfully sent or fail to be sent.
2. read-process-write mode: Message consumption and production are encapsulated in a transaction to form an atomic operation. In a streaming application, a service usually needs to receive messages from the upstream system, process the messages, and then send the processed messages to the downstream system. This corresponds to message consumption and production.

Example of secondary development code:

```
// Initialize the configuration and enable the transaction feature.
Properties props = new Properties();
props.put("enable.idempotence", true);
props.put("transactional.id", "transaction1");
...

KafkaProducer producer = new KafkaProducer<String, String>(props);

// init transaction
producer.initTransactions();
try {
    // Start a transaction.
    producer.beginTransaction();
    producer.send(record1);
    producer.send(record2);
    // Stop a transaction.
    producer.commitTransaction();
} catch (KafkaException e) {
    // Abort a transaction.
    producer.abortTransaction();
}
```

Nearby Consumption

Feature description: In versions earlier than Kafka 2.4.0, the production and consumption of the client are leader copies oriented to each partition. Follower copies are used only for data redundancy and do not provide services for external systems. As a result, the leader copy has high pressure. In addition, in cross-DC and cross-rack consumption scenarios, a large volume of data is transmitted between DCs and between racks. In Kafka 2.4.0 and later versions, the Kafka kernel can consume data from follower replicas, which greatly reduces the data

transmission volume and reduces the network bandwidth pressure in cross-DC and cross-rack scenarios. The community opens the ReplicaSelector API to support this feature. By default, MRS Kafka provides two methods to use this API.

1. **RackAwareReplicaSelector**: indicates that replicas in the same rack are preferentially consumed (nearby consumption in a rack).
2. **AzAwareReplicaSelector**: indicates that copies from nodes in the same AZ are preferentially consumed (nearby consumption in an AZ).

The following uses **RackAwareReplicaSelector** as an example to describe how to consume the closest replica.

```
public class RackAwareReplicaSelector implements ReplicaSelector {

    @Override
    public Optional<ReplicaView> select(TopicPartition topicPartition,
        ClientMetadata clientMetadata,
        PartitionView partitionView) {
        if (clientMetadata.rackId() != null && !clientMetadata.rackId().isEmpty()) {
            Set<ReplicaView> sameRackReplicas = partitionView.replicas().stream()
                // Filter the replicas that are in the same rack as the client.
                .filter(replicaInfo -> clientMetadata.rackId().equals(replicaInfo.endpoint().rack()))
                .collect(Collectors.toSet());
            if (sameRackReplicas.isEmpty()) {
                // If no replicas are in the same rack as the client, the leader replica is returned.
                return Optional.of(partitionView.leader());
            } else {
                // It shows that a replica that is in the same rack as the client exists.
                if (sameRackReplicas.contains(partitionView.leader())) {
                    // If the client and the leader replica are in the same rack, the leader replica returns first.
                    return Optional.of(partitionView.leader());
                } else {
                    // Otherwise, the latest replica synchronized with the leader is returned.
                    return sameRackReplicas.stream().max(ReplicaView.comparator());
                }
            }
        } else {
            // If the rack information is not contained in the client request, the leader replica is returned first.
            return Optional.of(partitionView.leader());
        }
    }
}
```

Method for enabling this feature:

1. Server: Update the **replica.selector.class** configuration item based on different features.
 - To enable "nearby consumption in a rack", set this parameter to **org.apache.kafka.common.replica.RackAwareReplicaSelector**.
 - To enable "nearby consumption in an AZ", set this parameter to **org.apache.kafka.common.replica.AzAwareReplicaSelector**.
2. Client: Add the **client.rack** configuration item to the **consumer.properties** file in the *{Client installation directory}/Kafka/kafka/config* directory.
 - If the "nearby consumption in a rack" is enabled on the server, add the information about the rack where the client is located, for example, **client.rack = /default0/rack1**.
 - If the "nearby consumption in an AZ" is enabled on the server, add the information about the rack where the client is located, for example, **client.rack = /AZ1/rack1**.

Ranger Unified Authentication

Feature description: In versions earlier than Kafka 2.4.0, Kafka supports only the SimpleAclAuthorizer authentication plugin provided by the community. In Kafka 2.4.0 and later versions, MRS Kafka supports both the Ranger authentication plugin and the authentication plugin provided by the community. Ranger authentication is used by default. Based on the Ranger authentication plugin, fine-grained Kafka ACL management can be performed.

NOTE

If the Ranger authentication plugin is used on the server and **allow.everyone.if.no.acl.found** is set to **true**, all actions are allowed when a non-secure port is used for access. You are advised to disable **allow.everyone.if.no.acl.found** for security clusters that use the Ranger authentication plugin.

15.21 Migrating Data Between Kafka Nodes

Scenario

This section describes how to use Kafka client commands to migrate partition data between disks on a node without stopping the Kafka service.

Prerequisites

- The MRS cluster administrator has understood service requirements and prepared a Kafka user (belonging to the **kafkaadmin** group. It is not required for the normal mode.).
- The Kafka client has been installed.
- The Kafka instance status and disk status are normal.
- Based on the current disk space usage of the partition to be migrated, ensure that the disk space will be sufficient after the migration.

Procedure

- Step 1** Log in as a client installation user to the node on which the Kafka client is installed.
- Step 2** Run the following command to switch to the Kafka client installation directory, for example, **/opt/kafkaclient**:
cd /opt/kafkaclient
- Step 3** Run the following command to set environment variables:
source bigdata_env
- Step 4** Run the following command to authenticate the user (skip this step in normal mode):
kinit Component service user
- Step 5** Run the following command to switch to the Kafka client directory:
cd Kafka/kafka/bin

Step 6 Run the following command to view the topic details of the partition to be migrated:

Security mode:

```
./kafka-topics.sh --describe --bootstrap-server IP address of the
Kafkacluster:21007 --command-config ../config/client.properties --topic topic
name
```

Normal mode:

```
./kafka-topics.sh --describe --bootstrap-server IP address of the Kafka
cluster:21005 --command-config ../config/client.properties --topic Topic name
```

| Topic | testws | PartitionCount | 24 | ReplicationFactor | 2 | Configs |
|---------------|---------------|----------------|---------------|-------------------|---|---------|
| Topic: testws | Partition: 0 | Leader: 4 | Replicas: 4,3 | Isr: 4,3 | | |
| Topic: testws | Partition: 1 | Leader: 5 | Replicas: 5,4 | Isr: 5,4 | | |
| Topic: testws | Partition: 2 | Leader: 6 | Replicas: 6,5 | Isr: 6,5 | | |
| Topic: testws | Partition: 3 | Leader: 3 | Replicas: 3,6 | Isr: 3,6 | | |
| Topic: testws | Partition: 4 | Leader: 4 | Replicas: 4,5 | Isr: 4,5 | | |
| Topic: testws | Partition: 5 | Leader: 5 | Replicas: 5,4 | Isr: 5,4 | | |
| Topic: testws | Partition: 6 | Leader: 6 | Replicas: 6,3 | Isr: 6,3 | | |
| Topic: testws | Partition: 7 | Leader: 3 | Replicas: 3,4 | Isr: 3,4 | | |
| Topic: testws | Partition: 8 | Leader: 4 | Replicas: 4,6 | Isr: 4,6 | | |
| Topic: testws | Partition: 9 | Leader: 5 | Replicas: 5,3 | Isr: 5,3 | | |
| Topic: testws | Partition: 10 | Leader: 6 | Replicas: 6,4 | Isr: 6,4 | | |
| Topic: testws | Partition: 11 | Leader: 3 | Replicas: 3,5 | Isr: 3,5 | | |
| Topic: testws | Partition: 12 | Leader: 4 | Replicas: 4,3 | Isr: 4,3 | | |
| Topic: testws | Partition: 13 | Leader: 5 | Replicas: 5,4 | Isr: 5,4 | | |
| Topic: testws | Partition: 14 | Leader: 6 | Replicas: 6,5 | Isr: 6,5 | | |
| Topic: testws | Partition: 15 | Leader: 3 | Replicas: 3,6 | Isr: 3,6 | | |
| Topic: testws | Partition: 16 | Leader: 4 | Replicas: 4,5 | Isr: 4,5 | | |
| Topic: testws | Partition: 17 | Leader: 5 | Replicas: 5,6 | Isr: 5,6 | | |
| Topic: testws | Partition: 18 | Leader: 6 | Replicas: 6,3 | Isr: 6,3 | | |
| Topic: testws | Partition: 19 | Leader: 3 | Replicas: 3,4 | Isr: 3,4 | | |
| Topic: testws | Partition: 20 | Leader: 4 | Replicas: 4,6 | Isr: 4,6 | | |
| Topic: testws | Partition: 21 | Leader: 5 | Replicas: 5,3 | Isr: 5,3 | | |
| Topic: testws | Partition: 22 | Leader: 6 | Replicas: 6,4 | Isr: 6,4 | | |

Step 7 Run the following command to query the mapping between **Broker_ID** and the IP address:

```
./kafka-broker-info.sh --zookeeper IP address of the ZooKeeper quorumpeer
instance.ZooKeeper port number/kafka
```

| Broker_ID | IP_Address |
|-----------|---------------|
| 4 | 192.168.0.100 |
| 5 | 192.168.0.101 |
| 6 | 192.168.0.102 |

NOTE

- IP address of the ZooKeeper quorumpeer instance
To obtain IP addresses of all ZooKeeper quorumpeer instances, log in to FusionInsight Manager and choose **Cluster > Services > ZooKeeper**. On the displayed page, click **Instance** and view the IP addresses of all the hosts where the quorumpeer instances locate.
- Port number of the ZooKeeper client
Log in to FusionInsight Manager and choose **Cluster > Service > ZooKeeper**. On the displayed page, click **Configurations** and check the value of **clientPort**. The default value is **24002**.

Step 8 Obtain the partition distribution and node information from the command output in **Step 6** and **Step 7**, and create the JSON file for reallocation in the current directory.

To migrate data in the partition whose **Broker_ID** is **6** to the **/srv/BigData/hadoop/data1/kafka-logs** directory, the required JSON configuration file is as follows:


```
{"partitions":[{"topic": "testws","partition": 2,"replicas": [6,5],"log_dirs": ["/srv/BigData/hadoop/data1/kafka-logs","any"]}], "version":1}
```

 NOTE

- **topic** indicates the topic name, for example, **testws**.
- **partition** indicates the topic partition.
- The number in **replicas** corresponds to **Broker_ID**.
- **log_dirs** indicates the path of the disk to be migrated. In this example, **log_dirs** of the node whose **Broker_ID** is 5 is set to **any**, and that of the node whose **Broker_ID** is 6 is set to **/srv/BigData/hadoop/data1/kafka-logs**. Note that the path must correspond to the node.

Step 9 Run the following command to perform reallocation:

Security mode:

```
./kafka-reassign-partitions.sh --bootstrap-server Service IP address of Broker:21007 --command-config ../config/client.properties --zookeeper {zk_host}:{port}/kafka --reassignment-json-file Path of the JSON file compiled in Step 8 --execute
```

Normal mode:

```
./kafka-reassign-partitions.sh --bootstrap-server Service IP address of Broker:21005 --command-config ../config/client.properties --zookeeper {zk_host}:{port}/kafka --reassignment-json-file Path of the JSON file compiled in Step 8 --execute
```

If message "Successfully started reassignment of partitions" is displayed, the execution is successful.

----End

15.22 Common Issues About Kafka

15.22.1 How Do I Solve the Problem that Kafka Topics Cannot Be Deleted?

Question

How do I delete a Kafka topic if it fails to be deleted?

Answer

- Possible cause 1: The **delete.topic.enable** configuration item is not set to **true**. The deletion can be performed only when the configuration item is set to **true**.
- Possible cause 2: The **auto.create.topics.enable** configuration parameter is set to **true**, which is used by other applications and is always running in the background.

Solution:

- For cause 1: Set **delete.topic.enable** to **true** on the configuration page.
- For cause 2: Stop the application that uses the topic in the background, or set **auto.create.topics.enable** to **false** (restart the Kafka service), and then delete the topic.

16 Using KafkaManager

16.1 Introduction to KafkaManager

KafkaManager is a tool for managing Apache Kafka and provides GUI-based metric monitoring and management of Kafka clusters.

KafkaManager supports the following functions:

- Manage multiple Kafka clusters.
- Check cluster status (topics, consumers, offsets, partitions, replicas, and nodes)
- Run preferred replica election.
- Generate partition assignments with option to select brokers to use.
- Run reassignment of partitions (based on generated assignments).
- Create a topic with optional topic configurations (Multiple Kafka cluster versions are supported).
- Delete a topic (only supported on 0.8.2+ and **delete.topic.enable = true** is set in broker configuration).
- Batch generate partition assignments for multiple topics with option to select brokers to use.
- Batch run reassignment of partitions for multiple topics.
- Add partitions to an existing topic.
- Update configurations for an existing topic.
- Optionally enable JMX polling for broker-level and topic-level metrics.
- Optionally filter out consumers that do not have `ids/owner/&offsets/` directories in ZooKeeper.

16.2 Accessing the KafkaManager Web UI

You can monitor and manage Kafka clusters on the graphical KafkaManager web UI.

Prerequisites

- KafkaManager has been installed in a cluster.
- The password of user **admin** has been obtained. The password of user **admin** is specified by the user during MRS cluster creation.

Accessing the KafkaManager Web UI

Step 1 Go to the cluster details page and choose **Components > KafkaManager**.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 2 In the **KafkaManager Summary** area, click any UI link in **KafkaManager WebUI** to access the KafkaManager web UI.

You can view the following information on the KafkaManager web UI.

- Kafka cluster list
- Broker node list and metric monitoring information of Kafka clusters
- Kafka cluster replica monitoring information
- Kafka cluster consumer monitoring information

 **NOTE**

You can click the KafkaManager logo in the upper left corner on any sub-page of KafkaManager to return to the homepage of the KafkaManager web UI, where a cluster list is displayed.

----End

16.3 Managing Kafka Clusters

Kafka cluster management includes the following operations:

- [Adding a Cluster on the KafkaManager Web UI](#)
- [Updating Cluster Parameters](#)
- [Deleting a Cluster on the KafkaManager Web UI](#)


Adding a Cluster on the KafkaManager Web UI

After a Kafka cluster is created for the first time, a default Kafka cluster named **my-cluster** is created on the KafkaManager web UI. You can also add Kafka clusters that have been created on the MRS management console on the KafkaManager web UI to manage multiple Kafka clusters.

Step 1 Log in to the KafkaManager web UI.

Step 2 In the upper part of the page, choose **Cluster > Add Cluster**.

Figure 16-1 Adding a cluster



Step 3 Set the cluster parameters. For the following parameters, refer to their example values. Retain the default values for other parameters.

Table 16-1 Cluster parameters to be modified

| Parameter | Example Value | Description |
|---|---|---|
| Cluster Name | mrs-demo | Name of the cluster to be added on the KafkaManager web UI |
| Cluster Zookeeper Hosts | zk1_ip:zk1_port,
zk2_ip:zk2_port/kafka | ZooKeeper address of the cluster to be added |
| Kafka Version | 1.1.0 | Kafka version of the cluster to be added. The default value is 1.1.0 . |
| Enable JMX Polling (Set JMX_PORT env variable before starting kafka server) | Selected | - |
| Poll consumer information (Not recommended for large # of consumers) | Selected | - |
| Enable Active OffsetCache (Not recommended for large # of consumers) | Selected | - |
| Display Broker and Topic Size (only works after applying this patch) | Selected | - |

| Parameter | Example Value | Description |
|-------------------|---------------|---|
| Security Protocol | PLAINTEXT | <ul style="list-style-type: none">For a Kafka cluster with Kerberos authentication enabled, select SASL_PLAINTEXT.For a Kafka cluster with Kerberos authentication disabled, select PLAINTEXT. |

Step 4 Click **Save**.

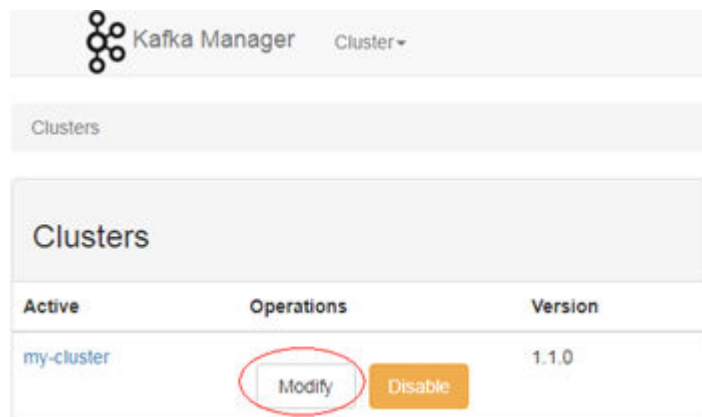
----End

Updating Cluster Parameters

Step 1 Log in to the KafkaManager web UI.

Step 2 Click **Modify** in the **Operations** column of the cluster.

Figure 16-2 Updating cluster parameters



Step 3 Go to the cluster configuration page and modify cluster parameters.

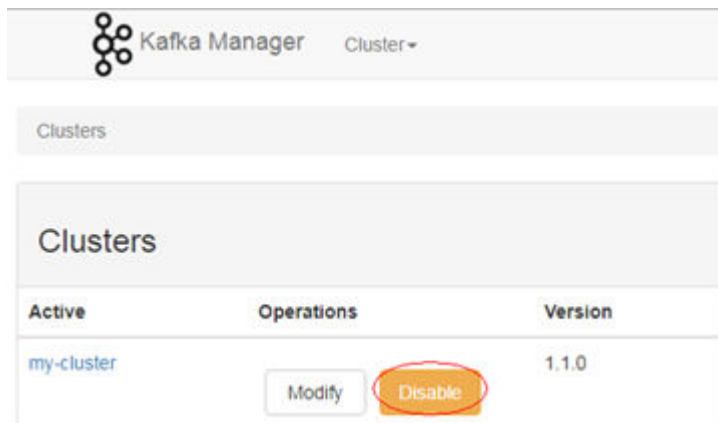
----End

Deleting a Cluster on the KafkaManager Web UI

Step 1 Log in to the KafkaManager web UI.

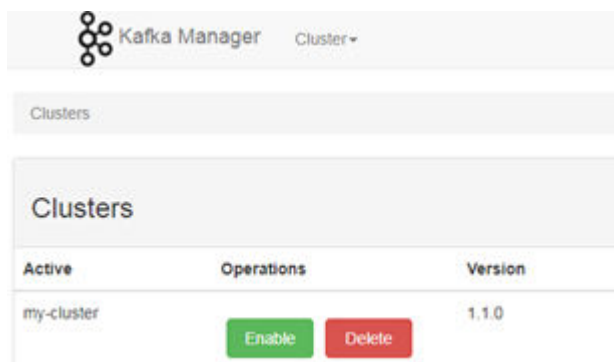
Step 2 Click **Disable** in the **Operations** column of the cluster.

Figure 16-3 Disabling a cluster



Step 3 When **Delete** or **Enable** is displayed in the **Operations** column on the cluster list page, click **Delete** to delete the cluster. You can also click **Enable** to enable the cluster.

Figure 16-4 Enabling or deleting a cluster



----End

16.4 Kafka Cluster Monitoring Management

The Kafka cluster monitoring management includes the following operations:


- [Viewing Broker Information](#)
- [Viewing Topic Information](#)
- [Viewing Consumers Information](#)
- [Modifying the Partition of a Topic Through KafkaManager](#)

Viewing Broker Information

Step 1 Log in to the KafkaManager web UI.

Step 2 On the cluster list page, click a cluster name to access the Summary page of the cluster.

Figure 16-5 Summary page of a cluster



Step 3 Click **Brokers** to access the Broker monitoring page. The page displays the Broker list and I/O statistics of the Broker nodes.

Figure 16-6 Broker monitoring page

| ← Brokers | | | | | | | Combined Metrics | | | | |
|-----------|-------------------------|-------|----------|----------|-----------|------|-----------------------------|------|-------|-------|--------|
| Id | Host | Port | JMX Port | Bytes In | Bytes Out | Size | Rate | Mean | 1 min | 5 min | 15 min |
| 1 | SSL-9093.PLAINTEXT-9092 | 21006 | 21006 | 0.00 | 0.00 | 0 B | Messages in /sec | 0.00 | 0.00 | 0.00 | 0.00 |
| 2 | SSL-9093.PLAINTEXT-9092 | 21006 | 21006 | 0.00 | 0.00 | 0 B | Bytes in /sec | 0.05 | 0.00 | 0.00 | 0.00 |
| 3 | SSL-9093.PLAINTEXT-9092 | 21006 | 21006 | 0.00 | 0.00 | 0 B | Bytes out /sec | 0.02 | 0.00 | 0.00 | 0.00 |
| | | | | | | | Bytes rejected /sec | 0.00 | 0.00 | 0.00 | 0.00 |
| | | | | | | | Failed fetch request /sec | 0.00 | 0.00 | 0.00 | 0.00 |
| | | | | | | | Failed produce request /sec | 0.00 | 0.00 | 0.00 | 0.00 |

----End

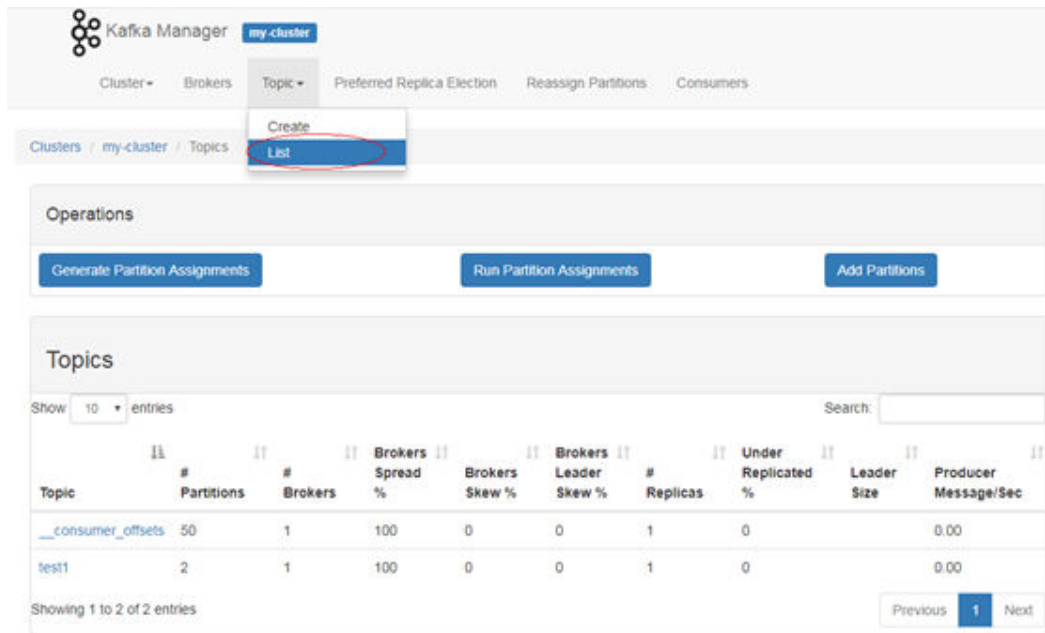
Viewing Topic Information

Step 1 Log in to the KafkaManager web UI.

Step 2 On the cluster list page, click a cluster name to access the **Summary** page of the cluster.


Step 3 Choose **Topic > List** to view the topic list of the current cluster and information about each topic.

Figure 16-7 Topic list



Step 4 Click a topic name to view details about the topic.

Figure 16-8 Topic details




----End

Viewing Consumers Information

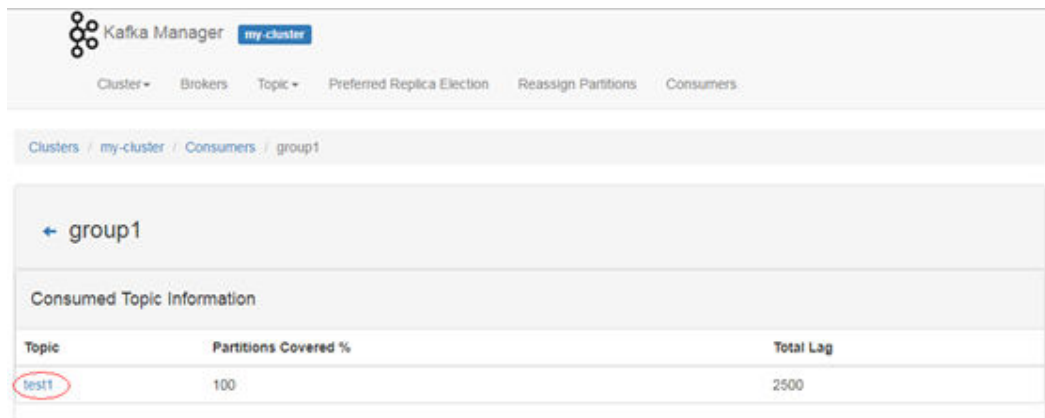
- Step 1** Log in to the KafkaManager web UI.
- Step 2** On the cluster list page, click a cluster name to access the **Summary** page of the cluster.
- Step 3** Click **Consumers** to view the consumers of the current cluster and each consumer's consumption information.

Figure 16-9 Consumers



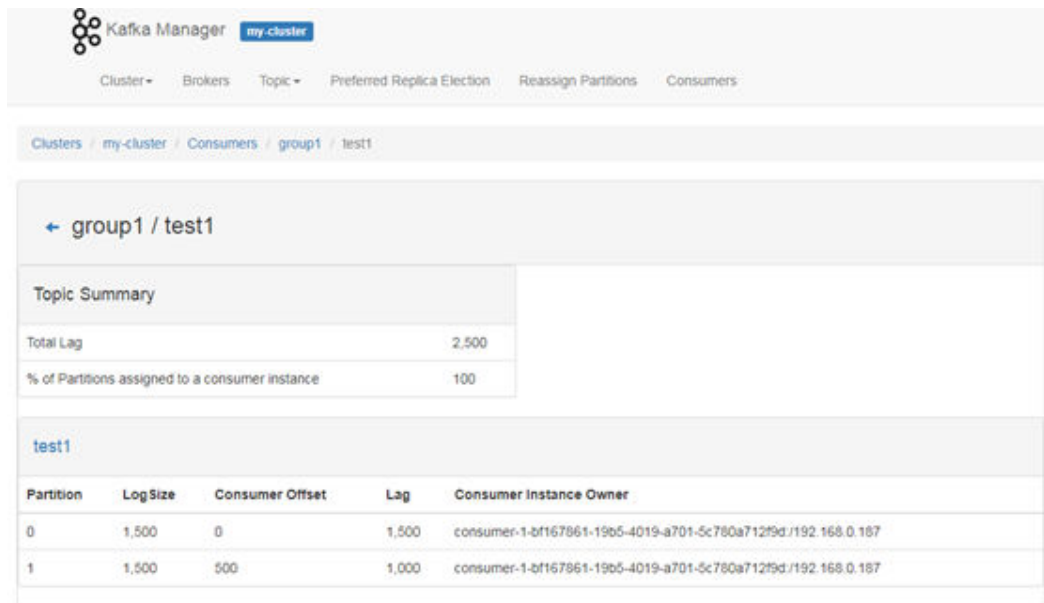
Step 4 Click a consumer name to view the list of the consumed topics.

Figure 16-10 List of topics consumed by the consumer



Step 5 Click a topic name in the topic list of the consumer to view consumption information about the topic.

Figure 16-11 Topic consumption details

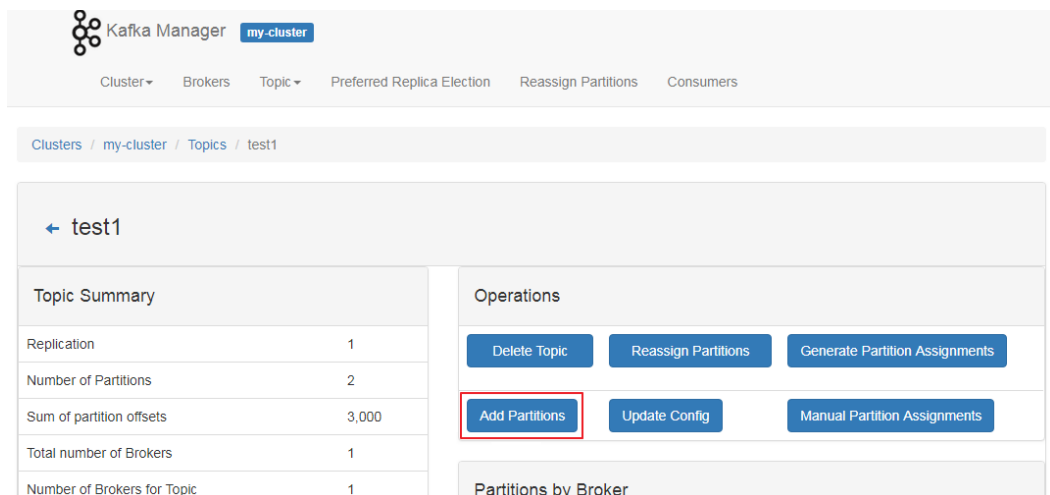


----End

Modifying the Partition of a Topic Through KafkaManager

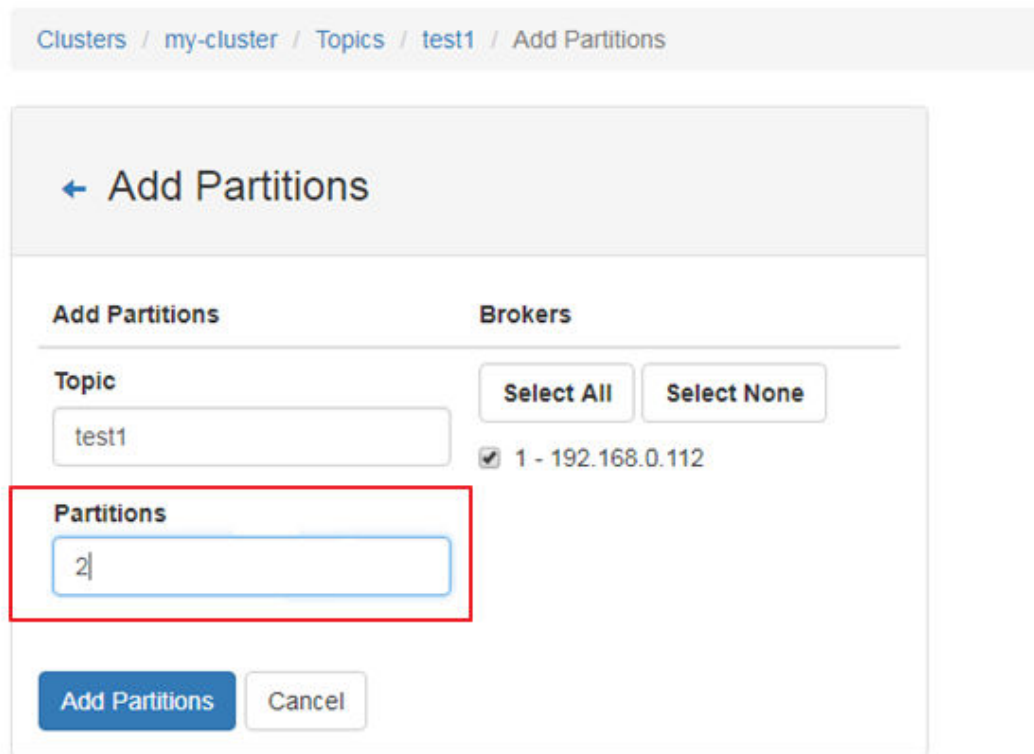
- Step 1** Log in to the KafkaManager web UI.
- Step 2** On the cluster list page, click a cluster name to access the **Summary** page of the cluster.
- Step 3** Choose **Topic > List** to access the topic list page of the current cluster.
- Step 4** Click a topic name to access the **Topic Summary** page.
- Step 5** Click **Add Partitions**. The page for adding partitions is displayed.

Figure 16-12 Adding partitions



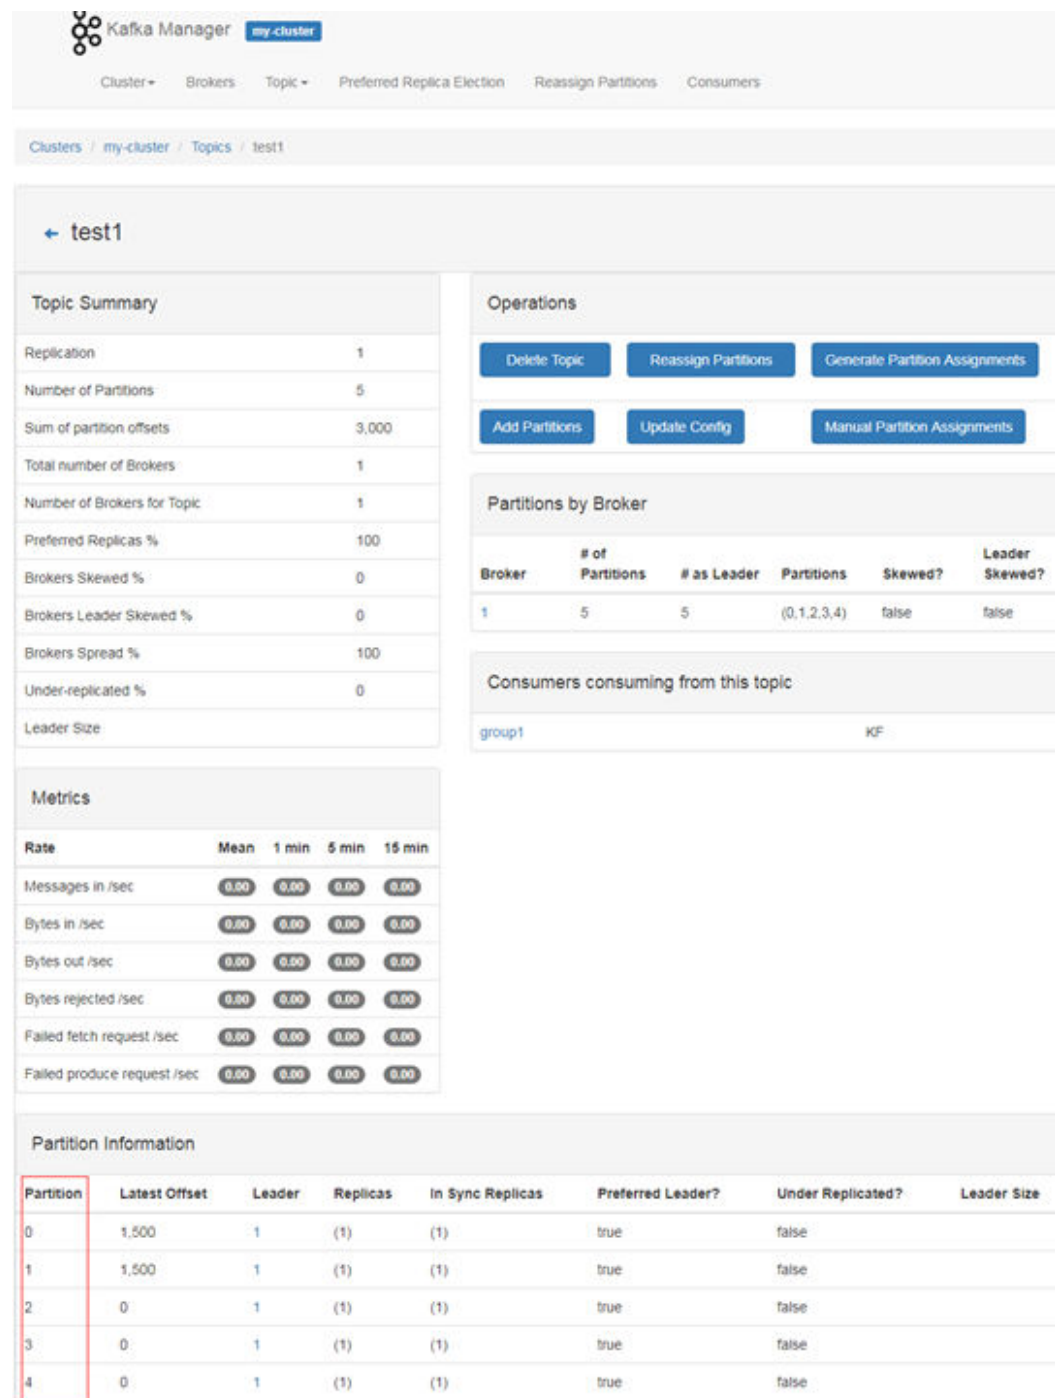
- Step 6** Confirm the topic name and modify the value of the **Partitions** parameter and click **Add Partitions** to add partitions.

Figure 16-13 Modifying the number of partitions



- Step 7** After the partitions are added successfully, click **Go to topic view** to return to the **Topic Summary** page.
- Step 8** Check the number of partitions in **Partition Information** in the lower part of the **Topic Summary** page.

Figure 16-14 Partition Information



Step 9 (Optional) If you are not satisfied with the assigned partitions, you can use the partition reassignment function to automatically reassign partitions.

1. On the **Topic Summary** page, click **Generate Partition Assignments**.
2. Select the broker instance and click **Generate Partition Assignments** to generate a partition.
3. After partition generation, click **Go to topic view** to return to the **Topic Summary** page.

4. On the **Topic Summary** page, click **Reassign Partitions** to automatically assign partitions to the broker instance of the cluster.
5. Click **Go to reassign partitions** to view details about the reassigned partitions.

Step 10 (Optional) If you are not satisfied with the automatically assigned partitions, you can manually assign the partitions.

1. On the **Topic Summary** page, click **Manual Partition Assignments** to access the page for manually assign partitions.
2. Manually assign a broker ID to each partition replica, and click **Save Partition Assignment** to save the changes.
3. Click **Go to topic view** to return to the **Topic Summary** page and view the partition details.

----End

17 Using Loader

17.1 Using Loader from Scratch

You can use Loader to import data from the SFTP server to HDFS.

This section applies to MRS clusters earlier than 3.x.

Prerequisites

- You have prepared service data.
- You have created an analysis cluster.

Procedure

Step 1 Access the Loader page.

1. Go to the cluster details page and choose **Services**.
2. Choose **Hue**. In **Hue Web UI** of **Hue Summary**, click **Hue (Active)**. The Hue web UI is displayed.
3. Choose **Data Browsers > Sqoop**.

The job management tab page is displayed by default on the Loader page.

Step 2 On the Loader page, click **Manage links**.

Step 3 Click **New link** and create **sftp-connector**. For details, see [File Server Link](#).

Step 4 Click **New link**, enter the link name, select **hdfs-connector**, and create **hdfs-connector**.

Step 5 On the Loader page, click **Manage jobs**.

Step 6 Click **New Job**.

Step 7 In **Connection**, set parameters.

1. In **Name**, enter a job name.
2. Select the source link created in [Step 3](#) and the target link created in [Step 4](#).

Step 8 In **From**, configure the job of the source link.

For details, see [ftp-connector or sftp-connector](#).

Step 9 In **To**, configure the job of the target link.

For details, see [hdfs-connector](#).

Step 10 In **Task Config**, set job running parameters.

Table 17-1 Loader job running properties

| Parameter | Description |
|--------------------------------------|--|
| Extractors | Number of Map tasks |
| Loaders | Number of Reduce tasks
This parameter is displayed only when the destination field is HBase or Hive. |
| Max. Error Records in a Single Shard | Error record threshold. If the number of error records of a single Map task exceeds the threshold, the task automatically stops and the obtained data is not returned.
NOTE
Data is read and written in batches for MYSQL and MPPDB of generic-jdbc-connector by default. Errors are recorded once at most for each batch of data. |
| Dirty Data Directory | Directory for saving dirty data. If you leave this parameter blank, dirty data will not be saved. |

Step 11 Click **Save**.

----End

17.2 How to Use Loader

This section applies to MRS clusters earlier than 3.x.

Process

The process for migrating user data with Loader is as follows:

1. Access the Loader page of the Hue web UI.
2. Manage Loader links.
3. Create a job and select a data source link and a link for saving data.
4. Run the job to complete data migration.

Loader Page

The Loader page is a graphical data migration management tool based on the open source Sqoop web UI and is hosted on the Hue web UI. Perform the following operations to access the Loader page:

1. Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).

2. Choose **Data Browsers > Sqoop**.

The job management tab page is displayed by default on the Loader page.

Loader Links

Loader links save data location information. Loader uses links to access data or save data to the specified location. Perform the following operations to access the Loader link management page:

1. Access the Loader page.
2. Click **Manage links**.

The Loader link management page is displayed.

Click **Manage jobs** to return to the job management page.

3. Click **New link** to go to the configuration page and set parameters to create a Loader link.

Loader Jobs

Loader jobs are used to manage data migration tasks. Each job consists of a source data link and a destination data link. A job reads data from the source link and saves data to the destination link to complete a data migration task.

17.3 Common Loader Parameters

This section applies to MRS 3.x or later.

Navigation Path

For details about the how to set parameters, see [Modifying Cluster Service Configuration Parameters](#).

Parameter Description

Table 17-2 Common Loader parameters

| Parameter | Description | Default Value | Value Range |
|--|--|---------------|-------------|
| mapreduce.client.submit.file.replication | Number of copies of the job files that the MapReduce task depends on in HDFS. If the number of DataNodes in the cluster is less than the value of this parameter, the number of copies is equal to the number of DataNodes. If the number of DataNodes is greater than or equal to the value of this parameter, the number of copies is the value of this parameter. | 10 | 3 to 256 |
| loader.fault.tolerance.rate | Error tolerance.
If the value is greater than 0, the error tolerance mechanism is enabled. When enabling the fault tolerance mechanism, you are advised to set the number of Map jobs to be greater than or equal to 3. It is recommended that this function be used when the job data volume is large. | 0 | 0 to 1.0 |
| loader.input.field.separator | Default input field separator. The parameter value takes effect only when input and output conversion steps are configured. The conversion steps can be left blank. If no separators are configured in job conversion steps, the default separator is used. | , | - |

| Parameter | Description | Default Value | Value Range |
|-------------------------------|--|---------------|-------------|
| loader.input.line.separator | Default input line separator. The parameter value takes effect only when input and output conversion steps are configured. The conversion steps can be left blank. If no separators are configured in job conversion steps, the default separator is used. | - | - |
| loader.output.field.separator | Default output field separator. The parameter value takes effect only when input and output conversion steps are configured. The conversion steps can be left blank. If no separators are configured in job conversion steps, the default separator is used. | , | - |
| loader.output.line.separator | Line separator of data that Loader outputs | - | - |

 NOTE

- Because it needs time to calculate the fault tolerance rate, you are recommended to use the **loader.fault.tolerance.rate** parameter when the job runtime is longer than 2 minutes to ensure user experience.
- Default separators are configured for the parameters in the preceding table for Loader. If separators are configured in the conversion steps for the jobs, the separators in the conversion steps will be used. If separators are not configured in the conversion steps, the default separators will be used.

17.4 Creating a Loader Role

Scenario

Create and configure a Loader role on FusionInsight Manager as an MRS cluster administrator. The Loader role can set Loader administrator permissions, job connections, job groups, and Loader job operation and scheduling permissions.

This section applies to MRS 3.x or later.

Prerequisites

- The MRS cluster administrator has understood service requirements.
- You have logged in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).

Procedure

Step 1 Choose **System > Permission > Role**.

Step 2 Click **Create Role** and set a role name and enter description.

Step 3 Set permissions. For details, see [Table 17-3](#).

NOTE

When setting permissions for a role, you cannot set permissions for multiple resources at the same time. If you need to set permissions for multiple resources, set them one by one.

Loader permissions:

- **Admin:** Loader administrator permission
- **Job Connector:** connection permission of Loader
- **Job Group:** permission to perform operations on Loader job groups. You can set the operation permissions of a specific job in a specified job group, including the **Edit** and **Execute** permissions of the job.
- **Job Scheduler:** permission to schedule Loader jobs

Table 17-3 Setting Loader roles

| Task | Role Authorization |
|---|--|
| Setting the Loader administrator permission | In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Loader and select Admin . |
| Setting the connection permission of Loader
(including the editing, deletion, and reference permissions of Job Connection) | <ol style="list-style-type: none"> 1. In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Loader > Job Connector. 2. In the Permission column of the specified job link, select Edit. |
| Setting the edit permission for Loader job groups
(including modifying the name of a job group, deleting a specified group, creating jobs in a specified group, importing jobs from external systems to a specified group in batches, and migrating jobs from other groups to a specified group) | <ol style="list-style-type: none"> 1. In the Permission table, choose Loader > Job Group. 2. In the Permission column of the specified job group, select Edit Group. |

| Task | Role Authorization |
|---|---|
| Setting the edit permission for all jobs in a Loader job group (including the permission to edit all existing or new jobs in the group) | <ol style="list-style-type: none"> In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Loader > Job Group. In the Permission column of the specified job group, select Edit Job. |
| Setting the execution permission for all jobs in the Loader job group (including the execution permission on all existing or new jobs in the group) | <ol style="list-style-type: none"> In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Loader > Job Group. In the Permission column of the specified job group, select Execute Job. |
| Setting the edit permission for Loader jobs (including editing, deleting, copying, and exporting jobs) | <ol style="list-style-type: none"> In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Loader > Job Group. Select a job group. In the Permission column of the specified job, select Edit. |
| Setting the execution permission for Loader jobs (including permissions to start and stop jobs and view historical records) | <ol style="list-style-type: none"> In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Loader > Job Group. Select a job group. In the Permission column of the specified job, select Execute. |
| Setting the permission for scheduling Loader jobs (including the editing, deletion, and validation permissions of Scheduler) | <ol style="list-style-type: none"> In the Configure Resource Permission table, choose <i>Name of the desired cluster</i> > Loader > Job Scheduling. In the Permission column of the specified job scheduling row, select Edit. |

 **NOTE**

- In addition to **Admin**, the preceding permissions are configured only for inventory resource information.
- Users without the preceding roles can also create tasks, groups, and connectors, but cannot perform operations on inventory resources.

Step 4 Click **OK**, and return to the **Role** page.

----End

17.5 Loader Link Configuration

This section applies to versions earlier than MRS 3.x.

Overview

Loader supports the following links. This section describes configurations of each link.

- obs-connector
- generic-jdbc-connector
- ftp-connector or sftp-connector
- hbase-connector, hdfs-connector, or hive-connector

OBS Link

An OBS link is a data exchange channel between Loader and OBS. [Table 17-4](#) describes the configuration parameters.

Table 17-4 obs-connector configuration

| Parameter | Description |
|--------------|---|
| Name | Name of a Loader connection |
| OBS Server | Enter an OBS endpoint. The common format is OBS.Region.DomainName .
Run the following command to query the endpoints of OBS:
cat /opt/Bigdata/apache-tomcat-7.0.78/webapps/web/WEB-INF/classes/cloud-obs.properties |
| Port | Specifies the port for accessing OBS data. The default value is 443 . |
| Access Key | AK for a user to access OBS |
| Security Key | SK corresponding to AK |

Relational Database Link

A relational database link is a data exchange channel between Loader and a relational database. [Table 17-5](#) describes the configuration parameters.

NOTE

Some parameters are hidden by default. They appear only after you click **Show Senior Parameter**.

Table 17-5 generic-jdbc-connector configuration

| Parameter | Description |
|---------------|---|
| Name | Name of a Loader link |
| Database Type | Data types supported by Loader links: ORACLE , MYSQL , and MPPDB |
| Host | Database access address, which can be an IP address or domain name. |
| Port | Port for accessing the database |
| Database | Name of the database saving data |
| Username | Username for accessing the database |
| Password | Password of the user Use the actual password. |

Table 17-6 Senior parameter configuration

| Parameter | Description |
|-----------------------|---|
| Fetch Size | A maximum volume of data obtained during each database access |
| Connection Properties | Drive properties exclusive to the database link supported by databases of different types, for example, autoReconnect of MYSQL . If you want to define the drive properties, click Add . |
| Identifier Enclose | Delimiter for reserving keywords in the database SQL. Delimiters defined in different databases vary. |

File Server Link

File server links include FTP and SFTP links and serve as a data exchange channel between Loader and a file server. [Table 17-7](#) describes the configuration parameters.

Table 17-7 ftp-connector or sftp-connector configuration

| Parameter | Description |
|-------------|---|
| Name | Name of a Loader link |
| Hostname/IP | Enter the file server access address, which can be a host name or IP address. |

| Parameter | Description |
|-----------|--|
| Port | Port for accessing the file server. <ul style="list-style-type: none">• Use port 21 for FTP.• Use port 22 for SFTP. |
| Username | Username for logging in to the file server |
| Password | Password of the user |

MRS Cluster Link

MRS cluster links include HBase, HDFS, and Hive links and serve as a data exchange channel between Loader and HBase, HDFS, or Hive.

When configuring an MRS cluster link, set the name, select a connector, for example, **hbase-connector**, **hdfs-connector**, or **hive-connector**, and save the settings.

17.6 Managing Loader Links (Versions Earlier Than MRS 3.x)

Scenario

You can create, view, edit, and delete links on the Loader page.

This section applies to versions earlier than MRS 3.x.

Prerequisites

You have accessed the Loader page. For details, see [Loader Page](#).

Creating a Link

Step 1 On the Loader page, click **Manage links**.

Step 2 Click **New link** and configure link parameters.

For details about the parameters, see [Loader Link Configuration](#).

Step 3 Click **Save**.

If link configurations, for example, IP address, port, and access user information, are incorrect, the link will fail to be verified and saved. In addition, VPC configurations may affect the network connectivity.

NOTE

You can click **Test** to immediately check whether the link is available.

----End

Viewing a Link

Step 1 On the Loader page, click **Manage links**.

- If Kerberos authentication is enabled for the cluster, all links created by the current user are displayed by default and other users' links cannot be displayed.
- If Kerberos authentication is disabled for the cluster, all Loader links of the cluster are displayed.

Step 2 In **Sqoop Links**, enter a link name to filter the link.

----End

Editing a Link

Step 1 On the Loader page, click **Manage links**.

Step 2 Click the link name to go to the edit page.

Step 3 Modify the link configuration parameters based on service requirements.

Step 4 Click **Test**.

If the test is successful, go to [Step 5](#). If a message displays indicating that OBS server cannot be connected, repeat [Step 3](#).

Step 5 Click **Save**.

If a Loader job has integrated into a Loader link, editing the link parameters may affect Loader running.

----End

Deleting a Link

Step 1 On the Loader page, click **Manage links**.

Step 2 Locate the row that contains the target link, and click **Delete**.

Step 3 In the dialog box, click **Yes, delete it**.

If a Loader job has integrated a Loader link, the link cannot be deleted.

----End

17.7 Managing Loader Links (MRS 3.x and Later Versions)

Scenario

You can create, view, edit, and delete links on the Loader page.

This section applies to MRS 3.x or later.

Creating a Connection

Step 1 Log in to the service page.

For MRS 3.x or later, log in to FusionInsight Manager (for details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#)) and choose **Cluster > Services**.

Step 2 Select **Loader**. On the right of **Loader WebUI**, click the link to open the Loader web UI.

Step 3 On the Loader page, click **New job**.

Step 4 Click **Add** next to **Connection** and set connection parameters.

For details about the parameters, see [Loader Connection Configuration](#).

Step 5 Click **OK**.

If connection configurations, for example, IP address, port, and access user information, are incorrect, the connection will fail to be verified and saved.

NOTE

You can click **Test** to immediately check whether the connection is available.

----End

Viewing a Connection

Step 1 On the Loader page, click **New job**.

Step 2 Click the drop-down list of **Connection** to view the connections you have created.

----End

Editing a Connection

Step 1 On the Loader page, click **New job**.

Step 2 Select the name of the connection to be edited from the **Connection** drop-down list.

Step 3 Click **Edit** next to **Connection**.

Step 4 On the dialog box displayed, modify the connection parameters based on service requirements.

Step 5 Click **Test**.

- If the test is successful, go to [Step 6](#).
- If the test fails, repeat [Step 4](#).

Step 6 Click **Save**.

If a Loader job has integrated into a Loader link, editing the link parameters may affect Loader running.

----End

Deleting a Connection

- Step 1** On the Loader page, click **New job**.
- Step 2** Select the name of the connection to be deleted from the **Connection** drop-down list.
- Step 3** Click **Delete**.
- Step 4** In the displayed dialog box, click **OK**.

If a Loader job has integrated a Loader connection, the connection cannot be deleted.

----End

Loader Connection Configuration

Loader supports the following connections:

- **generic-jdbc-connector**: For details about parameter settings, see [Table 17-8](#).
- **ftp-connector**: For details about parameter settings, see [Table 17-9](#).
- **sftp-connector**: For details about parameter settings, see [Table 17-10](#).
- **hdfs-connector**: For details about parameter settings, see [Table 17-11](#).
- **oracle-connector**: For details about parameter settings, see [Table 17-12](#).
- **mysql-fastpath-connector**: For details about parameter settings, see [Table 17-14](#).
- **oracle-partition-connector**: For details about parameter settings, see [Table 17-13](#).

Table 17-8 generic-jdbc-connector configuration

| Parameter | Description |
|-------------------|---|
| Name | Name of a Loader connection |
| Connector | Select generic-jdbc-connector . |
| JDBC Driver Class | JDBC driver classes are as follows: <ul style="list-style-type: none"> • oracle: oracle.jdbc.driver.OracleDriver • SQLServer: com.microsoft.sqlserver.jdbc.SQLServerDriver • mysql: com.mysql.jdbc.Driver • postgresql: org.postgresql.Driver |

| Parameter | Description |
|------------------------|---|
| JDBC Connection String | <p>Database access address, which can be an IP address or domain name.</p> <p>Enter the database connection string. In the following examples, 10.10.10.10 indicates the IP address and test indicates the database name.</p> <ul style="list-style-type: none"> oracle: jdbc:oracle:thin:@10.10.10.10:1521:orcl SQLServer: jdbc:sqlserver://10.10.10.10:1433;DatabaseName=test mysql: jdbc:mysql://10.10.10.10/test?&useUnicode=true&characterEncoding=GBK postgresql: jdbc:postgresql://10.10.10.10:5432/test |
| Username | Username for accessing the database |
| Password | Password of the user. Use the actual password. |

Table 17-9 ftp-connector configuration

| Parameter | Description |
|-------------------------|--|
| Name | Name of a Loader connection |
| Connector | Select ftp-connector . |
| FTP Mode | Select ACTIVE or PASSIVE . |
| FTP Protocol | <p>Select:</p> <ul style="list-style-type: none"> FTP SSL_EXPLICIT SSL_IMPLICIT TLS_EXPLICIT TLS_IMPLICIT |
| File Name Encoding Type | Encoding type of the file name or file path. |

Table 17-10 sftp-connector configuration

| Parameter | Description |
|-----------|--------------------------------|
| Name | Name of a Loader connection |
| Connector | Select sftp-connector . |

Table 17-11 hdfs-connector configuration

| Parameter | Description |
|-----------|--------------------------------|
| Name | Name of a Loader connection |
| Connector | Select hdfs-connector . |

Table 17-12 oracle-connector configuration

| Parameter | Description |
|------------------------|---|
| Name | Name of a Loader connection |
| Connector | Select oracle-connector . |
| JDBC Connection String | Connection string used for connecting to the database |
| Username | Username for accessing the database |
| Password | Password of the user. Use the actual password. |

Table 17-13 oracle-partition-connector configuration

| Parameter | Description |
|------------------------|---|
| Name | Name of a Loader connection |
| Connector | Select oracle-partition-connector . |
| JDBC Driver Class | Enter com.microsoft.jdbc.sqlserver.SQLServerDriver . |
| JDBC Connection String | Connection string used for connecting to the database |
| Username | Username for accessing the database |
| Password | Password of the user. Use the actual password. |

Table 17-14 mysql-fastpath-connector configuration

| Parameter | Description |
|-----------|-----------------------------|
| Name | Name of a Loader connection |

| Parameter | Description |
|------------------------|--|
| Connector | <p>Select mysql-fastpath-connector.</p> <p>NOTICE
When mysql-fastpath-connector is used, the mysqldump and mysqlimport commands of MySQL must be available on NodeManagers, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, see http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html. Install the MySQL client applications and tools.</p> <p>For example, you need to install the following RPM packages in the RHEL-x86 system (select the package version based on the site requirements):</p> <ul style="list-style-type: none"> mysql-community-client-5.7.23-1.el7.x86_64.rpm mysql-community-common-5.7.23-1.el7.x86_64.rpm mysql-community-devel-5.7.23-1.el7.x86_64.rpm mysql-community-embedded-5.7.23-1.el7.x86_64.rpm mysql-community-libs-5.7.23-1.el7.x86_64.rpm mysql-community-libs-compat-5.7.23-1.el7.x86_64.rpm |
| JDBC Connection String | Connection string used for connecting to the database |
| Username | Username for accessing the database |
| Password | Password of the user. Use the actual password. |

17.8 Source Link Configurations of Loader Jobs

Overview

When Loader jobs obtain data from different data sources, a link corresponding to a data source type needs to be selected and the link properties need to be configured.

This section applies to versions earlier than MRS 3.x.

obs-connector

Table 17-15 Data source link properties of **obs-connector**

| Parameter | Description |
|---------------------------|--|
| Bucket Name | OBS file system for storing source data. |
| Source Directory/
File | Actual storage form of source data. It can be either all data files in a directory or a single data file contained in the file system. |

| Parameter | Description |
|-----------------|---|
| File Format | Loader supports the following file formats of data stored in OBS: <ul style="list-style-type: none"> • CSV_FILE: Specifies a text file. When the destination link is a database link, only the text file is supported. • BINARY_FILE: Specifies binary files excluding text files. |
| Line Separator | Identifier of each line end of source data |
| Field Separator | Identifier of each field end of source data |
| Encoding Type | Text encoding type of source data. It takes effect on text files only. |
| File Split Type | The following types are supported: <ul style="list-style-type: none"> • File: The number of files is assigned to a map task by the total number of files. The calculation formula is Total number of files/Extractors. • Size: A file size is assigned to a map task by the total file size. The calculation formula is Total file size/Extractors. |

generic-jdbc-connector

Table 17-16 Data source link properties of **generic-jdbc-connector**

| Parameter | Description |
|-----------------------|---|
| Schema/
Tablespace | Name of the database storing source data. You can query and select it on the interface. |
| Table Name | Data table storing the source data. You can query and select it on the interface. |
| Partition Column | If multiple columns need to be read, use this column to split the result and obtain data. |
| Where Clause | Query statement used when accessing the database |

ftp-connector or sftp-connector

Table 17-17 Data source link properties of **ftp-connector** or **sftp-connector**

| Parameter | Description |
|---------------------------|--|
| Source Directory/
File | Actual storage form of source data. It can be either all data files in a directory or single data file contained in the file server. |

| Parameter | Description |
|-----------------|--|
| File Format | <p>Loader supports the following file formats of data stored in the file server:</p> <ul style="list-style-type: none"> • CSV_FILE: Specifies a text file. When the destination link is a database link, only the text file is supported. • BINARY_FILE: Specifies binary files excluding text files. |
| Line Separator | <p>Identifier of each line end of source data</p> <p>NOTE
If FTP or SFTP serves as a source link and File Format is set to BINARY_FILE, the value of Line Separator in the advanced properties is invalid.</p> |
| Field Separator | <p>Identifier of each field end of source data</p> <p>NOTE
If FTP or SFTP serves as a source link and File Format is set to BINARY_FILE, the value of Field Separator in the advanced properties is invalid.</p> |
| Encoding Type | Text encoding type of source data. It takes effect on text files only. |
| File Split Type | <p>The following types are supported:</p> <ul style="list-style-type: none"> • File: The number of files is assigned to a map task by the total number of files. The calculation formula is Total number of files/Extractors. • Size: A file size is assigned to a map task by the total file size. The calculation formula is Total file size/Extractors. |

hbase-connector

Table 17-18 Data source link properties of **hbase-connector**

| Parameter | Description |
|------------|---------------------------------|
| Table Name | HBase table storing source data |

hdfs-connector

Table 17-19 Data source link properties of **hdfs-connector**

| Parameter | Description |
|---------------------------|---|
| Source Directory/
File | Actual storage form of source data. It can be either all data files in a directory or single data file contained in HDFS. |

| Parameter | Description |
|-----------------|--|
| File Format | <p>Loader supports the following file formats of data stored in HDFS:</p> <ul style="list-style-type: none"> • CSV_FILE: Specifies a text file. When the destination link is a database link, only the text file is supported. • BINARY_FILE: Specifies binary files excluding text files. |
| Line Separator | <p>Identifier of each line end of source data</p> <p>NOTE
If HDFS serves as a source link and File Format is set to BINARY_FILE, the value of Line Separator in the advanced properties is invalid.</p> |
| Field Separator | <p>Identifier of each field end of source data</p> <p>NOTE
If HDFS serves as a source link and File Format is set to BINARY_FILE, the value of Field Separator in the advanced properties is invalid.</p> |
| File Split Type | <p>The following types are supported:</p> <ul style="list-style-type: none"> • File: The number of files is assigned to a map task by the total number of files. The calculation formula is Total number of files/Extractors. • Size: A file size is assigned to a map task by the total file size. The calculation formula is Total file size/Extractors. |

hive-connector

Table 17-20 Data source link properties of **hive-connector**

| Parameter | Description |
|---------------|--|
| Database Name | Name of the Hive database storing the data source. You can query and select it on the interface. |
| Table | Name of the Hive table storing the data source. You can query and select it on the interface. |

17.9 Destination Link Configurations of Loader Jobs

Overview

When Loader jobs save data to different storage locations, a destination link needs to be selected and the link properties need to be configured.

obs-connector

Table 17-21 Destination link properties of **obs-connector**

| Parameter | Description |
|------------------|--|
| Bucket Name | OBS file system for storing final data. |
| Output Directory | Directory for storing final data in the file system. A directory must be specified. |
| File Format | Loader supports the following file formats of data stored in OBS: <ul style="list-style-type: none"> • CSV_FILE: Specifies a text file. When the destination link is a database link, only the text file is supported. • BINARY_FILE: Specifies binary files excluding text files. |
| Line Separator | Identifier of each line end of final data |
| Field Separator | Identifier of each field end of final data |
| Encoding Type | Text encoding type of final data. It takes effect on text files only. |

generic-jdbc-connector

Table 17-22 Destination link properties of **generic-jdbc-connector**

| Parameter | Description |
|-------------|---|
| Schema Name | Name of the database storing final data |
| Table | Name of the table saving final data |

ftp-connector or sftp-connector

Table 17-23 Destination link properties of **ftp-connector** or **sftp-connector**

| Parameter | Description |
|------------------|--|
| Output Directory | Directory for storing final data in the file server. A directory must be specified. |
| File Format | Loader supports the following file formats of data stored in the file server: <ul style="list-style-type: none"> • CSV_FILE: Specifies a text file. When the destination link is a database link, only the text file is supported. • BINARY_FILE: Specifies binary files excluding text files. |

| Parameter | Description |
|-----------------|---|
| Line Separator | Identifier of each line end of final data
NOTE
If FTP or SFTP serves as a destination link and File Format is set to BINARY_FILE , the value of Line Separator in the advanced properties is invalid. |
| Field Separator | Identifier of each field end of final data
NOTE
If FTP or SFTP serves as a destination link and File Format is set to BINARY_FILE , the value of Field Separator in the advanced properties is invalid. |
| Encoding Type | Text encoding type of final data. It takes effect on text files only. |

hbase-connector

Table 17-24 Destination link properties of **hbase-connector**

| Parameter | Description |
|--------------------------|---|
| Table Name | Name of the HBase table saving final data. You can query and select it on the interface. |
| Method | Data can be imported to an HBase table using either BULKLOAD or PUTLIST . |
| Clear Data Before Import | Whether to clear data in the destination HBase table. Options are as follows: <ul style="list-style-type: none"> • True: Clean up data in the table. • False: Do not clean up data in the table. If you select False, an error is reported during job running if data exists in the table. |

hdfs-connector

Table 17-25 Destination link properties of **hdfs-connector**

| Parameter | Description |
|------------------|--|
| Output Directory | Directory for storing final data in HDFS. A directory must be specified. |

| Parameter | Description |
|-------------------|---|
| File Format | <p>Loader supports the following file formats of data stored in HDFS:</p> <ul style="list-style-type: none"> • CSV_FILE: Specifies a text file. When the destination link is a database link, only the text file is supported. • BINARY_FILE: Specifies binary files excluding text files. |
| Compression Codec | <p>Compression mode used when a file is saved to HDFS. The following modes are supported: NONE, DEFLATE, GZIP, BZIP2, LZ4, and SNAPPY.</p> |
| Overwrite | <p>How to process files in the output directory when files are imported to HDFS. Options are as follows:</p> <ul style="list-style-type: none"> • True: Clean up files in the directory and import new files by default. • False: Do not clean up files. If files exist in the output directory, job running fails. |
| Line Separator | <p>Identifier of each line end of final data</p> <p>NOTE
If HDFS serves as a destination link and File Format is set to BINARY_FILE, the value of Line Separator in the advanced properties is invalid.</p> |
| Field Separator | <p>Identifier of each field end of final data</p> <p>NOTE
If HDFS serves as a destination link and File Format is set to BINARY_FILE, the value of Field Separator in the advanced properties is invalid.</p> |

hive-connector

Table 17-26 Destination link properties of **hive-connector**

| Parameter | Description |
|-----------|---|
| Database | Name of the Hive database storing final data. You can query and select it on the interface. |
| Table | Name of the Hive table saving final data. You can query and select it on the interface. |

17.10 Managing Loader Jobs

Scenario

You can create, view, edit, and delete jobs on the Loader page.

This section applies to versions earlier than MRS 3.x.

Prerequisites

You have accessed the Loader page. For details, see [Loader Page](#).

Creating a Job

Step 1 On the Loader page, click **New job**.

Step 2 In **Connection**, set parameters.

1. In **Name**, enter a job name.
2. In **From link** and **To link**, select links accordingly.

After you select a link of a type, data is obtained from the specified source and saved to the destination.

NOTE

If no available link exists, click **Add a new link**.

Step 3 In **From**, configure the job of the source link.

For details, see [Source Link Configurations of Loader Jobs](#).

Step 4 In **To**, configure the job of the destination link.

For details, see [Destination Link Configurations of Loader Jobs](#).

Step 5 Check whether a database link is selected in **To link**.

Database links include:

- generic-jdbc-connector
- hbase-connector
- hive-connector

If you set **To link** to a database link, you need to configure a mapping between service data and a field in the database table.

- If you set it to a database link, go to [Step 6](#).
- If you do not set it to a database link, go to [Step 7](#).

Step 6 In **Field Mapping**, enter a field mapping. Then proceed to [Step 7](#).

Field Mapping specifies a mapping between each column of user data and a field in the database table.

Table 17-27 Field Mapping properties

| Parameter | Description |
|-------------------|--|
| Column Num | Field sequence of service data |
| Sample | First row of sample values of service data |
| Column Family | When To link is hbase-connector , you can select a column family for storing data. |
| Destination Field | Field for storing data |

| Parameter | Description |
|-----------|---|
| Type | Type of the field selected by the user |
| Row Key | When To link is hbase-connector , you need to select Destination Field as a row key. |

 **NOTE**

If the value of **From** is a connector of a file type, for example, SFTP, FTP, OBS, and HDFS files, the value of **Field Mapping** is the first row of data in the file. Ensure that the first row of data is complete. Otherwise, the Loader job will not extract columns that are not mapped.

Step 7 In **Task Config**, set job running parameters.

Table 17-28 Loader job running properties

| Parameter | Description |
|--------------------------------------|--|
| Extractors | Number of Map tasks |
| Loaders | Number of Reduce tasks
This parameter is displayed only when the destination field is HBase or Hive. |
| Max. Error Records in a Single Shard | Error record threshold. If the number of error records of a single Map task exceeds the threshold, the task automatically stops and the obtained data is not returned.
NOTE
Data is read and written in batches for MYSQL and MPPDB of generic-jdbc-connector by default. Errors are recorded once at most for each batch of data. |
| Dirty Data Directory | Specifies the directory for saving dirty data. If you leave this parameter blank, dirty data will not be saved. |

Step 8 Click **Save**.

----End

Viewing a Job

Step 1 Access the Loader page. The Loader job management page is displayed by default.

- If Kerberos authentication is enabled for the cluster, all jobs created by the current user are displayed by default and other users' jobs cannot be displayed.
- If Kerberos authentication is disabled for the cluster, all Loader jobs of the cluster are displayed.

Step 2 In **Sqoop Jobs**, enter a job name to filter the job.

Step 3 Click **Refresh** to obtain the latest job status.

----End

Editing a Job

Step 1 Access the Loader page. The Loader job management page is displayed by default.

Step 2 Click the job name to go to the edit page.

Step 3 Modify the job configuration parameters based on service requirements.

Step 4 Click **Save**.


NOTE

Basic job operations in the navigation bar on the left are **Run**, **Copy**, **Delete**, **Disable**, **History Record**, and **Show Job JSON Definition**.

----End

Deleting a Job

Step 1 Access the Loader page.

Step 2 In the row of the specified job, click .

You can also select one or more jobs and click **Delete Job** in the upper right corner of the job list.

Step 3 In the dialog box, click **Yes, delete it**.

If the state of a Loader job is **Running**, the job fails to be deleted.

----End

17.11 Preparing a Driver for MySQL Database Link

Scenario

As a component for batch data export, Loader can import and export data using a relational database.

Prerequisites

You have prepared service data.

Procedure

Procedure for MRS clusters earlier than 3.x:

Step 1 Download the MySQL JDBC driver **mysql-connector-java-5.1.21.jar** from the MySQL official website. For details about how to select the MySQL JDBC driver, see the following table.

Table 17-29 Version information

| JDBC Driver Version | MySQL Version |
|---------------------|---|
| Connector/J 5.1 | MySQL 4.1, MySQL 5.0, MySQL 5.1, and MySQL 6.0 alpha |
| Connector/J 5.0 | MySQL 4.1, MySQL 5.0 servers, and distributed transaction (XA) |
| Connector/J 3.1 | MySQL 4.1, MySQL 5.0 servers, and MySQL 5.0 except distributed transaction (XA) |
| Connector/J 3.0 | MySQL 3.x and MySQL 4.1 |

Step 2 Upload **mysql-connector-java-5.1.21.jar** to the Loader installation directory on the active and standby MRS Master nodes.

- For versions earlier than MRS 3.x, upload the package to **/opt/Bigdata/MRS_XXX/install/FusionInsight-Sqoop-1.99.7/FusionInsight-Sqoop-1.99.7/server/jdbc/**.

In the preceding path, **XXX** indicates the MRS version number. Change it based on site requirements.

Step 3 Change the owner of the **mysql-connector-java-5.1.21.jar** package to **omm:wheel**.

Step 4 Modify the **jdbc.properties** configuration file.

Change the key value of **MYSQL** to **mysql-connector-java-5.1.21.jar**, for example, **MYSQL=mysql-connector-java-5.1.21.jar**.

Step 5 Restart the Loader service.

----End

Procedure for MRS cluster 3.x and later versions:

Modify the permission on the JAR package of the relational database driver.

Step 1 Log in to the active and standby management nodes of the Loader service, obtain the driver JAR package of the relational database, and save it to the following directory on the active and standby Loader nodes: **`\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib**

 **NOTE**

The version 8.1.0.1 is used as an example. Replace it with the actual version number.

Step 2 Run the following commands as user **root** on the active and standby nodes of the Loader service to change the permission:

```
cd `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib
```

```
chown omm:wheel JAR package name
```

```
chmod 600 JARpackage name
```

Step 3 Log in to FusionInsight Manager. Choose **Cluster** and click the target cluster name. In the navigation pane on the left, choose **Services > Loader**. In the upper right corner, choose **More**, select **Restart Service**, and enter the password of the administrator to restart the Loader service.

----End

17.12 Importing Data

17.12.1 Overview

The section "Importing Data" applies to MRS 3.x or later.

Description

Loader is an ETL tool that enables MRS to exchange data and files with external data sources, such as relational databases, SFTP servers, and FTP servers. It allows data or files to be imported from relational databases or file systems to MRS.

Loader supports the following data import modes:

- Importing data from a relational database to HDFS or OBS
- Importing data from a relational database to HBase
- Importing data from a relational database to Phoenix tables
- Importing data from a relational database to Hive tables
- Importing data from an SFTP server to HDFS or OBS
- Importing data from an SFTP server to HBase
- Importing data from an SFTP server to Phoenix tables
- Importing data from an SFTP server to Hive tables
- Importing data from an FTP server to HDFS or OBS
- Importing data from an FTP server to HBase
- Importing data from an FTP server to Phoenix tables
- Importing data from an FTP server to Hive tables
- Importing data from HDFS or OBS to HBase in the same cluster

MRS needs to connect to an external data source to exchange data and files with the data source. The following connectors are used to configure connection parameters for different types of data sources:

- `generic-jdbc-connector`: relational database connector
- `ftp-connector`: FTP data source connector
- `hdfs-connector`: HDFS data source connector
- `oracle-connector`: dedicated connector for Oracle databases. `row_id` serves as partition columns. Compared with `generic-jdbc-connector`, Map jobs are more evenly distributed on `oracle-connector`, and whether indexes have been created for the partition columns does not matter.

- `mysql-fastpath-connector`: dedicated connector for MySQL databases. Data is imported and exported by using the `mysqldump` and `mysqlimport` tools of MySQL. Compared with `generic-jdbc-connector`, `mysql-fastpath-connector` delivers a faster data import and export speed.
- `sftp-connector`: SFTP data source connector
- `oracle-partition-connector`: connector that supports the Oracle partition feature, which is used to optimize the import and export of Oracle partition tables.

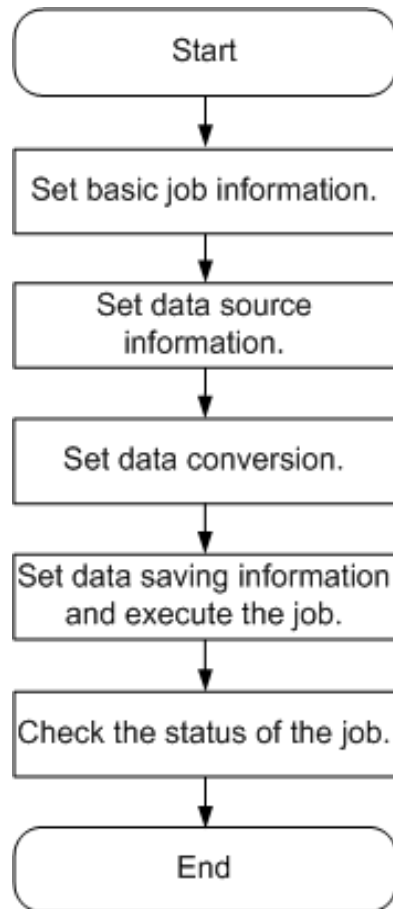
 **NOTE**

- When an FTP data source connector is used, data is not encrypted. This may result in security risks. You are advised to use an SFTP data source connector.
- You are advised to deploy the SFTP server, FTP server, database server, and Loader into separate subnets to ensure secure data import.
- For connection to relational databases, general database connectors (`generic-jdbc-connector`) or dedicated database connectors (`oracle-connector`, `oracle-partition-connector`, and `mysql-fastpath-connector`) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When `mysql-fastpath-connector` is used, the `mysqldump` and `mysqlimport` commands of MySQL must be available on NodeManager nodes, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, install the MySQL client applications and tools following the instructions at <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>.
- When `oracle-connector` is used, the connection user must be granted the select permission on the following system catalogs or views:
`dba_tab_partitions`, `dba_constraints`, `dba_tables`, `dba_segments`, `v$instance`, `dba_objects`, `v$instance`, `SYS_CONTEXT`, `dba_extents`, and `dba_tab_subpartitions`
- When `oracle-partition-connector` is used, the connection user must be granted the select permission on the following system catalogs: `dba_objects` and `dba_extents`.

Import Process

You can import data on the Loader web UI. [Figure 17-1](#) shows the data import process.

Figure 17-1 Import process



You can also use shell scripts to update and run Loader jobs. To use this method, you need to configure the installed Loader client.

17.12.2 Importing Data Using Loader

Scenario

This section describes how to import data from external data sources to MRS.

Generally, you can manually manage data import and export jobs on the Loader UI. To use shell scripts to update and run Loader jobs, you must configure the installed Loader client.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have the permission to access the HDFS or OBS directories, HBase tables, and data involved in job execution.
- You have obtained the username and password used by an external data source (SFTP server or relational database).
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.

- When using Loader to import data from SFTP, FTP, and HDFS/OBS, ensure that the input paths and input path subdirectories of the external data sources and the name of the files in these directories do not contain any of the special characters /'";,;
- If a task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.

Procedure

Step 1 Check whether data is imported from MRS to a relational database for the first time.

- If yes, go to [Step 2](#).
- If no, go to [Step 3](#).

Step 2 Modify the permission on the JAR package of the relational database driver.

1. Log in to the active and standby management nodes of the Loader service, obtain the JAR package of the relational database driver, and save it to the following directory on the active and standby Loader nodes: **\$ {BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib.**

NOTE

The version 8.1.0.1 is used as an example. Replace it with the actual version number.

2. Run the following command as user **root** on the active and standby nodes of the Loader service to modify the permission:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib
```

```
chown omm:wheel JAR package name
```


```
chmod 600 JAR package name
```

3. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **More** > **Restart Service**. Enter the password of the administrator to restart the Loader service.

Step 3 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-2 Loader web UI



- Step 4** Create a Loader data import job. Click **New Job**. Select the required job type in **1. Basic Information**, and click **Next**.
1. Set **Name** to the job name and **Type** to **Import**.
 2. Select a connection for **Connection**. By default, no connection is created. Click **Add** to create a connection, and then click **Test** to test whether the connection is available. Click **OK** when the system displays a message indicates that the test is successful.

Data sources need to be connected when MRS exchanges data and files with external data sources. **Connection** indicates the set of connection parameters for connecting to data sources.

Table 17-30 Connection configuration parameters

| Connector | Parameter | Description |
|------------------------|----------------------------|---|
| generic-jdbc-connector | JDBC Driver Class | Name of a JDBC driver class |
| | JDBC Connection String | JDBC connection string |
| | Username | Username for connecting to the database |
| | Password | Password for connecting to the database |
| | JDBC Connection Properties | JDBC connection attribute. Click Add to manually add connection attributes. <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |
| ftp-connector | FTP Server IP Address | IP address of the FTP server |
| | FTP Server Port | Port number of the FTP server |
| | FTP Username | Username for accessing the FTP server |
| | FTP Password | Password for accessing the FTP server |
| | FTP Mode | FTP access mode. Possible values are ACTIVE and PASSIVE . If this parameter is not set, FTP access is in passive mode by default. |

| Connector | Parameter | Description |
|--------------------------|-------------------------|--|
| | FTP Protocol | <p>FTP protocol.</p> <ul style="list-style-type: none"> - FTP: indicates the FTP protocol. - SSL_EXPLICIT: indicates the explicit SSL protocol. - SSL_IMPLICIT: indicates the implicit SSL protocol. - TLS_EXPLICIT: indicates the explicit TLS protocol. - TLS_IMPLICIT: indicates the implicit TLS protocol. <p>If this parameter is not set, the FTP protocol is used by default.</p> |
| | File Name Encoding Type | File name and file path encoding format supported by the FTP server. If this parameter is not set, the default format UTF-8 is used. |
| hdfs-connector | - | - |
| oracle-connector | JDBC Connection String | Connection string for a user to connect to the database |
| | Username | Username for connecting to the database |
| | Password | Password for connecting to the database |
| | Connection Properties | <p>Connection attributes. Click Add to manually add connection attributes.</p> <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |
| mysql-fastpath-connector | JDBC Connection String | JDBC connection string |
| | Username | Username for connecting to the database |
| | Password | Password for connecting to the database |
| | Connection Properties | <p>Connection attributes. Click Add to manually add connection attributes.</p> <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |
| sftp-connector | SFTP Server IP Address | IP address of the SFTP server |

| Connector | Parameter | Description |
|----------------------------|------------------------|---|
| | SFTP Server Port | Port number of the SFTP server |
| | SFTP Username | Username for accessing the SFTP server |
| | SFTP Password | Password for accessing the SFTP server |
| | SFTP Public Key | Public key of the SFTP server |
| oracle-partition-connector | JDBC Driver Class | Name of a JDBC driver class |
| | JDBC Connection String | JDBC connection string |
| | Username | Username for connecting to the database |
| | Password | Password for connecting to the database |
| | Connection Properties | Connection attributes. Click Add to manually add connection attributes. <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |

3. Set **Group** to the group to which the job belongs. By default, there is no created group. Click **Add** to create a group and click **OK**.
4. **Queue** indicates that Loader tasks are executed in a specified Yarn queue. The default value is **root.default**, which indicates that the tasks are executed in the **default** queue.
5. Set **Priority** to the priority of Loader tasks in the specified Yarn queue. The value can be **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, or **VERY_HIGH**. The default value is **NORMAL**.

Step 5 In the **2. Input Settings** area, set the data source and click **Next**.

 **NOTE**

- When creating or editing a Loader job, you can use macro definitions when configuring parameters such as the SFTP path, HDFS/OBS path, and Where condition of SQL. For details, see [Using Macro Definitions in Configuration Items](#) .
- Loader supports common field data types, such as Char, VarChar, Boolean, Binary, SmallInt, Int, BigInt, Decimal, Float, Double, Date, Time, TimeStamp, and String. The supported types may vary according to the data source. For details about the supported types, expand the field data type drop-down list of the corresponding input operator (such as Table Input) on the Loader GUI. Some database-specific fields may not be supported. For example, Loader does not support the CLOB, XMLType, and BLOB fields in Oracle.

Table 17-31 List of input configuration parameters

| Source File Type | Parameter | Description |
|---------------------------------|-----------------|--|
| sftp-connector or ftp-connector | Input Path | Input path or name of the source file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple input paths separated with semicolons (;). Ensure that the number of input paths is the same as that of SFTP servers configured for the connector. |
| | File Split Type | Indicates whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each Map in the MapReduce task for data import. FILE indicates that each Map processes one or more complete source files. The same source file cannot be allocated to different Maps. When the data is saved to the output directory, the directory structure of the input path is retained. SIZE indicates that each Map processes input files of a certain size. A source file can be split into multiple Maps. The number of files saved when data is saved to the output directory is the same as that of Maps. The file name format is import_part_xxxx , where <i>xxxx</i> is a unique random number generated by the system. |
| | Filter Type | File filtering criterion. WILDCARD indicates that a wildcard is used in filtering, and REGEX indicates that a regular expression is used in filtering. This parameter is used together with Path Filter and File Filter . The default value is WILDCARD . |
| | Path Filter | Wildcard or regular expression for filtering the directories in the input path of the source files. This parameter is used when Filter Type is set. Input Path is not used for filtering. If there are multiple filter conditions, use commas (,) to separate them. If the value is empty, the directories are not filtered. |
| | File Filter | Wildcard or regular expression for filtering the file names of the source files. This parameter is used when Filter Type is set. If there are multiple filter conditions, use commas (,) to separate them. The value cannot be left blank. |

| Source File Type | Parameter | Description |
|------------------------|---------------|---|
| | Encoding Type | Source file encoding format, for example, UTF-8. This parameter can be set only in text file import. |
| | Suffix | File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. |
| | Compression | Indicates whether to enable compressed transmission when SFTP is used to export data. true indicates that compression is enabled, and false indicates that compression is disabled. |
| hdfs-connector | Input Path | Input path of source files in HDFS |
| | Path Filter | Wildcard for filtering the directories in the input paths of the source files. Input Path is not used for filtering. If there are multiple filter conditions, use commas (,) to separate them. If the value is empty, the directories are not filtered. The regular expression filtering is not supported. |
| | File Filter | Wildcard for filtering the file names of the source files. If there are multiple filter conditions, use commas (,) to separate them. The value cannot be left blank. The regular expression filtering is not supported. |
| | Encoding Type | Source file encoding format, for example, UTF-8. This parameter can be set only in text file import. |
| | Suffix | File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. |
| generic-jdbc-connector | Schema Name | Database schema name. This parameter exists in the Table name schema. |
| | Table Name | Database table name. This parameter exists in the Table name schema. |

| Source File Type | Parameter | Description |
|------------------|---------------------------|--|
| | SQL Statement | SQL statement for the Loader to query data to be imported in Table SQL statement mode. The SQL statement requires the query condition WHERE \${CONDITIONS} . Without this condition, the SQL statement cannot be run properly, for example, select * from TABLE WHERE A>B and \${CONDITIONS} . If Table column names is set, the column specified by Table column names will replace the column queried in the SQL statement. This parameter cannot be set when Schema name or Table name is set. |
| | Table Column Names | Table columns whose content is to be imported by Loader. Use commas (,) to separate multiple fields. |
| | Partition Column Name | Database table column based on which to-be-imported data is determined. This parameter is used for partitioning in a Map job. You are advised to configure the primary key field.
NOTE <ul style="list-style-type: none"> • A partition column must have an index. If no index exists, do not specify a partition column. If a partition column without an index is specified, the database server disk I/O will be busy, the access of other services to the database will be affected, and the import will take a long period. • In multiple fields with indexes, select the field that has the most discrete value as the partition column. A partition column that is not discrete may result in load imbalance when multiple MapReduce jobs are imported. • The sorting rules of partition columns must be case-sensitive. Otherwise, data may be lost during data import. • You are not advised to select fields of the float or double type for the partition column. Otherwise, the records containing the minimum and maximum values of the partition column may fail to be imported due to precision issues. |
| | Nulls in Partition Column | Indicates whether to process records whose values are null in database table columns. If the value is true , the data whose value is null in the partition column is processed. If the value is false , the data whose value is null in the partition column is not processed. |

| Source File Type | Parameter | Description |
|--------------------------|---------------------------------------|---|
| | Whether to Specify a Partition Column | Indicates whether to specify a partition column. |
| oracle-connector | Table Name | Table name. |
| | Column Name | Column name. |
| | Query Condition | Query condition in an SQL statement |
| | Splitting Mode | Data splitting mode. The options are ROWID and PARTITION . |
| | Table Partition Name | Name of a table partition. Use commas (,) to separate the names of different partitions. |
| | Data Block Allocation Mode | Allocation method of data after being split. |
| | Read Size | Amount of data to be read each time. |
| mysql-fastpath-connector | Schema Name | Database schema name. |
| | Table Name | Database table name. |
| | Query Condition | Query condition of a specified table. |
| | Partition Column Name | <p>Database table column based on which to-be-imported data is determined. This parameter is used for partitioning in a Map job. You are advised to configure the primary key field.</p> <p>NOTE</p> <ul style="list-style-type: none"> • A partition column must have an index. If no index exists, do not specify a partition column. If a partition column without an index is specified, the database server disk I/O will be busy, the access of other services to the database will be affected, and the import will take a long period. • In multiple fields with indexes, select the field that has the most discrete value as the partition column. A partition column that is not discrete may result in load imbalance when multiple MapReduce jobs are imported. • You are not advised to select fields of the float or double type for the partition column. Otherwise, the records containing the minimum and maximum values of the partition column may fail to be imported due to precision issues. |

| Source File Type | Parameter | Description |
|----------------------------|---------------------------------------|--|
| | Nulls in Partition Column | Indicates whether to process records whose values are null in database table columns. If the value is true , the data whose value is null in the partition column is processed. If the value is false , the data whose value is null in the partition column is not processed. |
| | Whether to Specify a Partition Column | Indicates whether to specify a partition column. |
| oracle-partition-connector | Schema Name | Database schema name. |
| | Table Name | Partition table name. |
| | Query Condition | Query condition in an SQL statement. |
| | Table Column Names | Table columns whose content is to be imported by Loader. Use commas (,) to separate multiple fields. |

Step 6 In the **3. Convert** area, set the conversion operations during data transmission.

Check whether source data values in the data operation job created by the Loader can be directly used without conversion, including upper and lower case conversion, cutting, merging, and separation.

- If yes, click **Next**.
 - If no, perform [Step 6.1](#) to [Step 6.4](#).
1. No created conversion step exists by default. Drag an example conversion step on the left to the edit box to create a new conversion step.
 2. Conversion step types must be selected based on service requirements. A complete conversion process includes the following types:
 - a. Input type. Only one conversion step can be added. This parameter is mandatory if the task involves HBase or relational databases.
 - b. Conversion type, which is an intermediate conversion step. You can add one or more conversion types or do not add any conversion type.
 - c. Output type. Only one output type can be added in the last conversion step. This parameter is mandatory if the task involves HBase or relational databases.

Table 17-32 Example list

| Type | Description |
|------------|---|
| Input Type | <ul style="list-style-type: none"> <li data-bbox="762 353 1362 454">▪ CSV File Input: CSV file input step for configuring separators to generate multiple fields. <li data-bbox="762 481 1410 582">▪ Fixed-Width File Input: Text file input step for configuring the length of characters or bytes to be truncated to generate multiple fields. <li data-bbox="762 609 1430 710">▪ Table Input: relational data input step for configuring specified columns in the database as input fields. <li data-bbox="762 736 1398 837">▪ HBase Input: HBase table input step for configuring the column definition of an HBase table to a specified field. <li data-bbox="762 864 1406 965">▪ HTML Input: HTML web page data input step for obtaining the target data of the HTML web page file to the specified field. <li data-bbox="762 992 1385 1093">▪ Hive Input: Hive table input step for defining columns in a Hive table to specified fields. <li data-bbox="762 1097 1406 1198">▪ Spark Input: Spark SQL table input step for defining columns in the SparkSQL table to specified fields. Only SparkSQL can access Hive data. |

| Type | Description |
|-----------------|---|
| Conversion Type | <ul style="list-style-type: none"> <li data-bbox="762 309 1406 409">▪ Long Integer Time Conversion: Configure the conversion between a long integer value and a date. <li data-bbox="762 439 1385 501">▪ Null Value Conversion: Configure a specified value to replace the null value. <li data-bbox="762 530 1358 593">▪ Random Value Conversion: Configure new value-added fields as random data fields. <li data-bbox="762 622 1374 685">▪ Adding a Constant Field: Add a constant to directly generate a constant field. <li data-bbox="762 714 1426 815">▪ Concatenation and Conversion: Concatenate fields, connect generated fields using connection characters, and convert new fields. <li data-bbox="762 844 1422 938">▪ Separator Conversion: Configure the generated fields to be separated by separators and convert new fields. <li data-bbox="762 967 1390 1068">▪ Modulo Conversion: Configure the generated fields to be converted into new fields through modulo operation. <li data-bbox="762 1097 1422 1198">▪ Cutting Character String: Truncate a generated field based on a specified position to generate a new field. <li data-bbox="762 1227 1414 1350">▪ EL Operation Conversion: Calculate field values. Currently, the following operators are supported: md5sum, sha1sum, sha256sum, and sha512sum. <li data-bbox="762 1379 1394 1480">▪ Character String Case Conversion: Configure the generated fields to be converted to new fields through case conversion. <li data-bbox="762 1509 1318 1572">▪ Reverse String Conversion: Reverse the generated fields to generate new fields. <li data-bbox="762 1601 1394 1695">▪ Character String Space Clearing Conversion: Configure the generated fields to clear spaces and convert them to new fields. <li data-bbox="762 1724 1374 1825">▪ Row Filtering Conversion: Configure logical conditions to filter out rows that contain triggering conditions. <li data-bbox="762 1854 1406 1912">▪ Update Fields: Update the value of a specified field when certain conditions are met. |

| Type | Description |
|-------------|--|
| Output Type | <ul style="list-style-type: none"> ▪ File Output: Configure generated fields to be connected by separators and exported to a file. ▪ Table Output: Configure the mapping between output fields and specified columns in the database. ▪ HBase Output: Configure the generated fields to the columns of the HBase table. ▪ Hive Output: Configure generated fields to a column of a Hive table. ▪ Spark Output: Configure generated fields to the columns of SparkSQL tables. Only SparkSQL can access Hive data. |

The edit box allows you to perform the following tasks:

- **Rename:** Rename an example.
- **Edit:** Edit the step conversion by referring to [Step 6.3](#).
- **Delete:** Delete an example.

 **NOTE**

You can also use the shortcut key **Del** to delete the example.

3. Click **Edit** to edit the step conversion information and configure fields and data.

For details about how to set parameters in the step conversion information, see [Operator Help](#).

 **NOTE**

- When sftp-connector or ftp-connector is used to import data, the time type field in the original data must be set to a string during the data conversion so that the time can be accurate to millisecond for data import. The data that has more precise time than millisecond will not be imported.
- When generic-jdbc-connector is used to import data, it is recommended that the data length of the CHAR or VARCHAR type field be set to -1 during data conversion so that all data can be imported. This prevents the data from being truncated when the actual data length is too long.
- When generic-jdbc-connector is used to import data, the time type field in the original data must be set to a time type value during the data conversion so that the time can be accurate to second for data import. The data that has more precise time than second will not be imported.
- When data is imported to a Hive partitioned table, Hive does not scan the newly imported data by default. You need to run the following HQL statement to repair the table so that the newly imported data can be queried:

MSCK REPAIR TABLE *table_name*;

If the conversion step is incorrectly configured, the source data cannot be converted and become dirty data. The dirty data marking rules are as follows:

- In any input type step, all data becomes dirty data if the number of fields contained in the original data is less than that of configured fields, or the field values in the original data do not match the configured field type.
- In the **CSV File Input** step, **Validate input field** checks whether the input field matches the value type. If the input field and value type of a row do not match, the row is skipped and becomes dirty data.
- In the **Fixed Width File Input** step, **Fixed Length** specifies the field splitting length. If the length is greater than the length of the original field value, data splitting fails and the current row becomes dirty data.
- In the **HBase Input** step, if the HBase table name specified by **HBase Table Name** is incorrect, or no primary key column is configured for **Primary Key**, all data becomes dirty data.
- In any conversion step, rows whose conversion fails become dirty data. For example, in the **Split Conversion** step, if the number of generated fields is less than that of configured fields, or the original data cannot be converted to the String type, the current row becomes dirty data.
- In the **Filter Row Conversion** step, rows filtered by filter criteria become dirty data.
- In the **Modulo Conversion** step, if the original field value is **NULL**, the current row becomes dirty data.
- For jobs that import data to Hive/SparkSQL tables, you must configure the Hive conversion step.

4. Click **Next**.

Step 7 In the **4. Output Settings** area, set the destination location for saving data and click **Save** to save the job or click **Save and Run** to save and run the job.

Table 17-33 List of output configuration parameters

| Storage Type | Parameter | Description |
|--------------|-----------|--|
| HDFS | File Type | <p>The type of the file to be saved after data is imported to HDFS. Possible values are as follows:</p> <ul style="list-style-type: none"> • TEXT_FILE: imports a text file and stores it as a text file. • SEQUENCE_FILE: imports a text file and stores it as a sequence file. • BINARY_FILE: imports files of any format by using binary streams but not to process the files. <p>NOTE
When the file import type to TEXT_FILE or SEQUENCE_FILE, Loader automatically selects a decompression method based on the file name extension to decompress a file.</p> |

| Storage Type | Parameter | Description |
|--------------|--------------------|--|
| | Compression Format | Compression format of files imported to HDFS. Select a format from the drop-down list. If you select NONE or do not set this parameter, data is not compressed. |
| | Output Directory | Directory for storing data imported into HDFS. |
| | Operation | <p>Action during data import. When all data is to be imported from the input path to the destination path, the data is stored in a temporary directory and then copied from the temporary directory to the destination path. After the data is imported successfully, the data is deleted from the temporary directory. One of the following actions can be taken when duplicate file names exist during data transfer:</p> <ul style="list-style-type: none"> • OVERRIDE: overrides the old file. • RENAME: renames as new file. For a file without an extension, a string is added to the file name as the extension; for a file with an extension, a string is added to the extension. The string is unique. • APPEND: adds the content of the new file to the end of the old file. This action only adds content regardless of whether the file can be used. For example, a text file can be used after this operation, while a compressed file cannot. • IGNORE: reserves the old file and does not copy the new file. • ERROR: stops the task and reports an error if duplicate file names exist. Transferred files are imported successfully, while files that have duplicate names and files that are not transferred fail to be imported. |
| | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. This parameter cannot be set when Extractor Size is set. The value must be less than or equal to 3000. |

| Storage Type | Parameter | Description |
|----------------|--------------------------|---|
| | Extractor Size | Size of data processed by Maps that are started in a MapReduce task of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000 . This parameter cannot be set when Extractors is set. When a relational database connector is used, Extractor Size is unavailable. You need to set Extractors . |
| HBASE_BULKLOAD | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. |
| | Clear data before import | Indicates whether to clear data in the original table before importing data. The value true indicates that the clearing operation is performed, and the value false indicates that the clearing operation is not performed. If you do not set this parameter, the original table is not cleared by default. |
| | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. |
| | Extractor Size | HBase does not support this parameter. Please set Extractors . |
| HBASE_PUTLIST | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. |
| | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. |
| | Extractor Size | HBase does not support this parameter. Please set Extractors . |
| HIVE | Output Directory | Directory for storing data imported into Hive. |

| Storage Type | Parameter | Description |
|--------------|------------------|---|
| | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. This parameter cannot be set when Extractor Size is set. The value must be less than or equal to 3000. |
| | Extractor Size | Size of data processed by Maps that are started in a MapReduce task of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000 . This parameter cannot be set when Extractors is set. When a relational database connector is used, Extractor Size is unavailable. You need to set Extractors . |
| SPARK | Output Directory | Only SparkSQL is supported to access Hive data. You can specify the directory for storing data imported to Hive. |
| | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. This parameter cannot be set when Extractor Size is set. The value must be less than or equal to 3000. |
| | Extractor Size | Size of data processed by Maps that are started in a MapReduce task of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000 . This parameter cannot be set when Extractors is set. When a relational database connector is used, Extractor Size is unavailable. You need to set Extractors . |

Step 8 On the Loader web UI, view, start, stop, copy, delete, edit, or view historical information about created jobs.

Figure 17-3 Viewing Loader jobs

The screenshot shows a web interface for managing Loader jobs. At the top, there are buttons for 'New Job', 'Import', 'Export', 'Change Group', and 'Batch Delete'. A search bar for job names and a refresh button are also present. Below is a table with columns for Job ID, Name, Description, Start Time, Execute User, Progress, Status, and Operation. Two jobs are listed, both with a status of 'Succeeded' and 100% progress.

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|-------------|--------------------------|---------------------|--------------|----------|-----------|-----------|
| 1 | sftp2hdfs | Import from SFTP to H... | 2016-02-03 14:16:48 | admin | 100% | Succeeded | [Icons] |
| 2 | sftp2hdfs-1 | Import from SFTP to H... | 2016-02-03 11:03:54 | admin | 100% | Succeeded | [Icons] |

----End

17.12.3 Typical Scenario: Importing Data from an SFTP Server to HDFS or OBS

Scenario

Use Loader to import data from an SFTP server to HDFS or OBS.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HDFS or OBS directories and data involved in job execution.
- You have obtained the username and password of the SFTP server as well as the read permission for the source files on the SFTP server. If file name extension needs to be added after a source file is imported, the user must have the write permission of the source file.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- When using Loader to import data from the SFTP server, the input paths and input path subdirectories of the SFTP server and the name of the files in these directories do not contain any of the special characters `/'";,;`.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-4 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-5 Basic Information



1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **sftp-connector**, click **Add**, set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**. Loader allows multiple SFTP servers to be configured. Click **Add** to add the configuration information of multiple SFTP servers.

Table 17-34 Connection parameters

| Parameter | Description | Example Value |
|------------------------|--|---------------|
| Name | Name of the SFTP server connection | sftpName |
| SFTP Server IP Address | IP address of the SFTP server | 10.16.0.1 |
| SFTP Server Port | Port number of the SFTP server | 22 |
| SFTP Username | Username for accessing the SFTP server | root |
| SFTP Password | Password for accessing the SFTP server | xxxx |

| Parameter | Description | Example Value |
|-----------------|-------------------------------|---------------|
| SFTP Public Key | Public key of the SFTP server | OdDt/yn...etM |

 **NOTE**

When multiple SFTP servers are configured, the data in the specified directories of the SFTP servers is imported to the same directory in HDFS or OBS.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-35 Parameter description

| Parameter | Description | Example Value |
|-----------------|---|----------------------------|
| Input Path | <p>Input path or name of the source file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple input paths separated with semicolons (;). Ensure that the number of input paths is the same as that of SFTP servers configured for the connector.</p> <p>NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items .</p> | /opt/
tempfile;/
opt |
| File Split Type | <p>Indicates whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each Map in the MapReduce task for data import.</p> <ul style="list-style-type: none"> • FILE: indicates that the source file is split by file. That is, each Map processes one or multiple complete files, the same source file cannot be allocated to different Maps, and the source file directory structure is retained after data import. • SIZE: indicates that the source file is split by size. That is, each Map processes input files of a certain size, and a source file can be divided and processed by multiple Maps. After data is stored in the output directory, the number of saved files is the same as that of Maps. The file name format is import_part_xxxx, where <i>xxxx</i> is a unique random number generated by the system. | FILE |

| Parameter | Description | Example Value |
|---------------|---|-----------------------|
| Filter Type | <p>File filter condition. This parameter is used when Path Filter or File Filter is set.</p> <ul style="list-style-type: none"> ● WILDCARD: indicates using a wildcard. ● REGEX: indicates using a regular expression. ● If the parameter is not set, a wildcard is used by default. | WILDCARD |
| Path Filter | <p>Wildcard or regular expression for filtering the directories in the input path of the source files. This parameter is used when Filter Type is set. Input Path is not used for filtering. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. If this parameter is left empty, directories are not filtered.</p> <ul style="list-style-type: none"> ● ? matches a single character. ● * indicates multiple characters. ● Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | 1*,2*;1* |
| File Filter | <p>Wildcard or regular expression for filtering the file names of the source files. This parameter is used when Filter Type is set. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. This parameter cannot be left blank.</p> <ul style="list-style-type: none"> ● ? matches a single character. ● * indicates multiple characters. ● Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | *.txt,*.csv;
*.txt |
| Encoding Type | Source file encoding format, for example, UTF-8 and GBK. This parameter can be set only in text file import. | UTF-8 |

| Parameter | Description | Example Value |
|-------------|---|---------------|
| Suffix | File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. This parameter is valid only when the data source is a file system. You are advised to set this parameter in incremental data import.

For example, if the parameter is set to .txt and the source file is test-loader.csv , the source file name is test-loader.csv.txt after export. | .log |
| Compression | Indicates whether to enable compressed transmission when SFTP is used to export data. <ul style="list-style-type: none"> The value true indicates that compression is enabled. The value false indicates that compression is disabled. | true |


Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-36](#).

Table 17-36 Input and output parameters of the operator

| Input Type | Output Type |
|------------------|-------------|
| CSV File Input | File Output |
| HTML Input | File Output |
| Fixed File Input | File Output |

Figure 17-6 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HDFS**.

Table 17-37 Parameter description

| Parameter | Description | Example Value |
|--------------------|---|---------------|
| File Type | Type of a file after the file is imported. The options are as follows: <ul style="list-style-type: none"> • TEXT_FILE: imports a text file and stores it as a text file. • SEQUENCE_FILE: imports a text file and stores it as a sequence file. • BINARY_FILE: imports files of any format by using binary streams. | TEXT_FILE |
| Compression Format | Compression format of files imported to HDFS or OBS. Select a format from the drop-down list. If you select NONE or do not set this parameter, data is not compressed. | NONE |
| Output Directory | Directory for storing data imported into HDFS or OBS.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /user/test |

| Parameter | Description | Example Value |
|------------|---|---------------|
| Operation | <p>Action during data import. When all data is to be imported from the input path to the destination path, the data is stored in a temporary directory and then copied from the temporary directory to the destination path. After the data is imported successfully, the data is deleted from the temporary directory. One of the following actions can be taken when duplicate file names exist during data transfer:</p> <ul style="list-style-type: none"> ● OVERWRITE: overrides the old file. ● RENAME: renames as new file. For a file without an extension, a string is added to the file name as the extension; for a file with an extension, a string is added to the extension. The string is unique. ● APPEND: adds the content of the new file to the end of the old file. This action only adds content regardless of whether the file can be used. For example, a text file can be used after this operation, while a compressed file cannot. ● IGNORE: reserves the old file and does not copy the new file. ● ERROR: stops the task and reports an error if duplicate file names exist. Transferred files are imported successfully, while files that have duplicate names and files that are not transferred fail to be imported. | OVERWRITE |
| Extractors | <p>Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. This parameter cannot be set when Extractor Size is set. The value must be less than or equal to 3000. You are advised to set the parameter to the number of CPU cores on the SFTP server.</p> <p>NOTE
To improve the data import speed, ensure that the following conditions are met:</p> <ul style="list-style-type: none"> ● Each Map connection is equivalent to a client connection. Therefore, you must ensure that the maximum number of connections of the SFTP server is greater than the number of Maps. ● Ensure that the disk I/O or network bandwidth on the SFTP server does not reach the upper limit. | 20 |


| Parameter | Description | Example Value |
|----------------|---|---------------|
| Extractor Size | Size of data processed by Maps that are started in a MapReduce task of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000. This parameter cannot be set when Extractors is set. | 1000 |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-7 Viewing job details



----End

17.12.4 Typical Scenario: Importing Data from an SFTP Server to HBase

Scenario

Use Loader to import data from an SFTP server to HBase.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HBase tables or phoenix tables that are used during job execution.
- You have obtained the username and password of the SFTP server as well as the read permission for the source files on the SFTP server. If file name extension needs to be added after a source file is imported, the user must have the write permission of the source file.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- When using Loader to import data from the SFTP server, the input paths and input path subdirectories of the SFTP server and the name of the files in these directories do not contain any of the special characters `/'";,;`.

- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

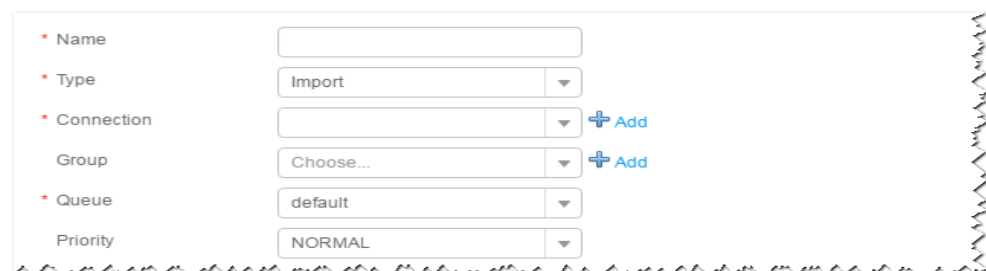
1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-8 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-9 Basic Information



1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **sftp-connector**, click **Add**, set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**. Loader

allows multiple SFTP servers to be configured. Click **Add** to add the configuration information of multiple SFTP servers.

Table 17-38 Connection parameters

| Parameter | Description | Example Value |
|------------------------|--|---------------|
| Name | Name of the SFTP server connection | sftpName |
| SFTP Server IP Address | IP address of the SFTP server | 10.16.0.1 |
| SFTP Server Port | Port number of the SFTP server | 22 |
| SFTP Username | Username for accessing the SFTP server | root |
| SFTP Password | Password for accessing the SFTP server | xxxx |
| SFTP Public Key | Public key of the SFTP server | OdDt/yn...etM |

 **NOTE**

When multiple SFTP servers are configured, the data in the specified directories of the servers is imported to HBase.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-39 Parameter description

| Parameter | Description | Example Value |
|------------|---|------------------------|
| Input Path | <p>Input path or name of the source file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple input paths separated with semicolons (;). Ensure that the number of input paths is the same as that of SFTP servers configured for the connector.</p> <p>NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items .</p> | /opt/
tempfile;/opt |

| Parameter | Description | Example Value |
|-----------------|---|---------------|
| File Split Type | <p>Indicates whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each Map in the MapReduce task for data import.</p> <ul style="list-style-type: none"> ● FILE: indicates that the source file is split by file. That is, each Map processes one or multiple complete files, the same source file cannot be allocated to different Maps, and the source file directory structure is retained after data import. ● SIZE: indicates that the source file is split by size. That is, each Map processes input files of a certain size, and a source file can be divided and processed by multiple Maps. After data is stored in the output directory, the number of saved files is the same as that of Maps. The file name format is import_part_xxxx, where xxxx is a unique random number generated by the system. | FILE |
| Filter Type | <p>File filter condition. This parameter is used when Path Filter or File Filter is set.</p> <ul style="list-style-type: none"> ● WILDCARD: indicates using a wildcard. ● REGEX: indicates using a regular expression. ● If the parameter is not set, a wildcard is used by default. | WILDCARD |
| Path Filter | <p>Wildcard or regular expression for filtering the directories in the input path of the source files. This parameter is used when Filter Type is set. Input Path is not used for filtering. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. If this parameter is left empty, directories are not filtered.</p> <ul style="list-style-type: none"> ● ? matches a single character. ● * indicates multiple characters. ● Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | 1*,2*;1* |

| Parameter | Description | Example Value |
|---------------|---|-------------------|
| File Filter | <p>Wildcard or regular expression for filtering the file names of the source files. This parameter is used when Filter Type is set. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. This parameter cannot be left blank.</p> <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | *.txt;*.csv;*.txt |
| Encoding Type | Source file encoding format, for example, UTF-8 and GBK. This parameter can be set only in text file import. | UTF-8 |
| Suffix | <p>File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. This parameter is valid only when the data source is a file system. You are advised to set this parameter in incremental data import.</p> <p>For example, if the parameter is set to .txt and the source file is test-loader.csv, the source file name is test-loader.csv.txt after export.</p> | .log |
| Compression | <p>Indicates whether to enable compressed transmission when SFTP is used to export data.</p> <ul style="list-style-type: none"> • The value true indicates that compression is enabled. • The value false indicates that compression is disabled. | true |


Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-40](#).

Table 17-40 Input and output parameters of the operator

| Input Type | Output Type |
|------------------|--------------|
| CSV File Input | HBase Output |
| HTML Input | HBase Output |
| Fixed File Input | HBase Output |

Figure 17-10 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HBASE_BULKLOAD** or **HBASE_PUTLIST** based on the actual situation.

Table 17-41 Parameter description

| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|--------------------------|---|---------------|
| HBASE_BULKLOAD | Large data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| | | Clear data before import | Indicates whether to clear data in the original table before importing data. The value true indicates that the data is cleared, and the value false indicates that the data is not cleared. If you do not set this parameter, the original table is not cleared by default. | true |

| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|----------------|---|---------------|
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. You are advised to set the parameter to the maximum number of connections on the SFTP server. | 20 |
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |
| HBASE_P UTLIST | Small data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |

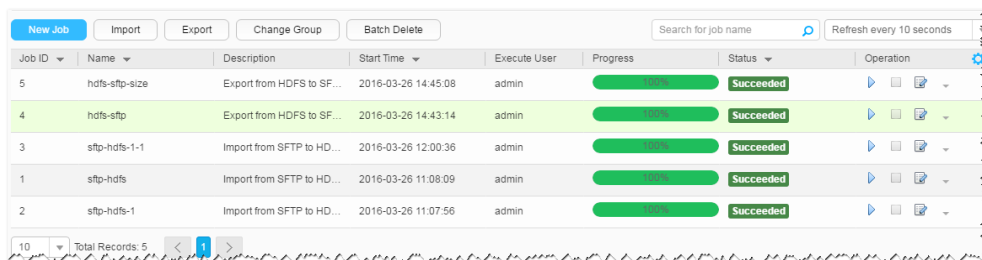
| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|--------------|---------------------|----------------|---|---------------|
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. | 20 |
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-11 Viewing job details



----End

17.12.5 Typical Scenario: Importing Data from an SFTP Server to Hive

Scenario

Use Loader to import data from an SFTP server to Hive.

Prerequisites

- You have obtained the service username and password for creating a Loader job.

- You have had the permission to access the Hive table specified in the job.
- You have obtained the username and password of the SFTP server as well as the read permission for the source files on the SFTP server. If file name extension needs to be added after a source file is imported, the user must have the write permission of the source file.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- When using Loader to import data from the SFTP server, the input paths and input path subdirectories of the SFTP server and the name of the files in these directories do not contain any of the special characters /'";.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-12 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-13 Basic Information



1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.

3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **sftp-connector**, click **Add**, set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**. Loader allows multiple SFTP servers to be configured. Click **Add** to add the configuration information of multiple SFTP servers.

Table 17-42 Connection parameters

| Parameter | Description | Example Value |
|------------------------|--|---------------|
| Name | Name of the SFTP server connection | sftpName |
| SFTP Server IP Address | IP address of the SFTP server | 10.16.0.1 |
| SFTP Server Port | Port number of the SFTP server | 22 |
| SFTP Username | Username for accessing the SFTP server | root |
| SFTP Password | Password for accessing the SFTP server | xxxx |
| SFTP Public Key | Public key of the SFTP server | OdDt/yn...etM |

 **NOTE**

When multiple SFTP servers are configured, the data in the specified directories of the servers is imported to Hive.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-43 Parameter description

| Parameter | Description | Example Value |
|-----------------|---|----------------------|
| Input Path | <p>Input path or name of the source file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple input paths separated with semicolons (;). Ensure that the number of input paths is the same as that of SFTP servers configured for the connector.</p> <p>NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items.</p> | /opt/tem
file;opt |
| File Split Type | <p>Indicates whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each Map in the MapReduce task for data import.</p> <ul style="list-style-type: none"> ● FILE: indicates that the source file is split by file. That is, each Map processes one or multiple complete files, the same source file cannot be allocated to different Maps, and the source file directory structure is retained after data import. ● SIZE: indicates that the source file is split by size. That is, each Map processes input files of a certain size, and a source file can be divided and processed by multiple Maps. After data is stored in the output directory, the number of saved files is the same as that of Maps. The file name format is import_part_xxxx, where xxxx is a unique random number generated by the system. | FILE |
| Filter Type | <p>File filter condition. This parameter is used when Path Filter or File Filter is set.</p> <ul style="list-style-type: none"> ● WILDCARD: indicates using a wildcard. ● REGEX: indicates using a regular expression. ● If the parameter is not set, a wildcard is used by default. | WIL
DC
AR
D |

| Parameter | Description | Example Value |
|---------------|---|-------------------|
| Path Filter | <p>Wildcard or regular expression for filtering the directories in the input path of the source files. This parameter is used when Filter Type is set. Input Path is not used for filtering. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. If this parameter is left empty, directories are not filtered.</p> <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | 1*,2*,1* |
| File Filter | <p>Wildcard or regular expression for filtering the file names of the source files. This parameter is used when Filter Type is set. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. This parameter cannot be left blank.</p> <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | *.txt,*.csv;*.txt |
| Encoding Type | <p>Source file encoding format, for example, UTF-8 and GBK. This parameter can be set only in text file import.</p> | UTF-8 |
| Suffix | <p>File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. This parameter is valid only when the data source is a file system. You are advised to set this parameter in incremental data import.</p> <p>For example, if the parameter is set to .txt and the source file is test-loader.csv, the source file name is test-loader.csv.txt after export.</p> | .log |

| Parameter | Description | Example Value |
|-------------|---|---------------|
| Compression | <p>Indicates whether to enable compressed transmission when SFTP is used to export data.</p> <ul style="list-style-type: none"> The value true indicates that compression is enabled. The value false indicates that compression is disabled. | true |

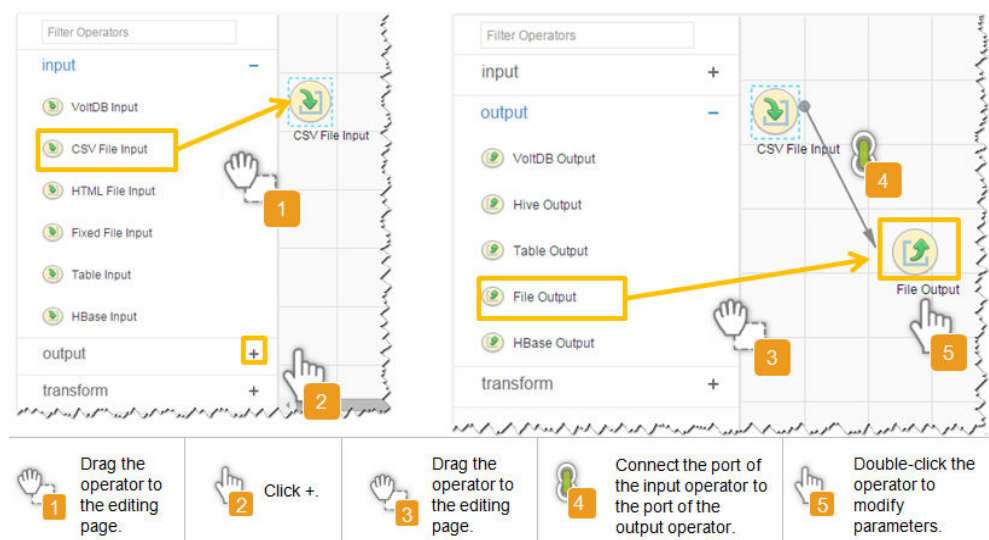
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-44](#).

Table 17-44 Input and output parameters of the operator

| Input Type | Output Type |
|------------------|-------------|
| CSV File Input | Hive Output |
| HTML Input | Hive Output |
| Fixed File Input | Hive Output |

Figure 17-14 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HIVE**.

Table 17-45 Parameter description

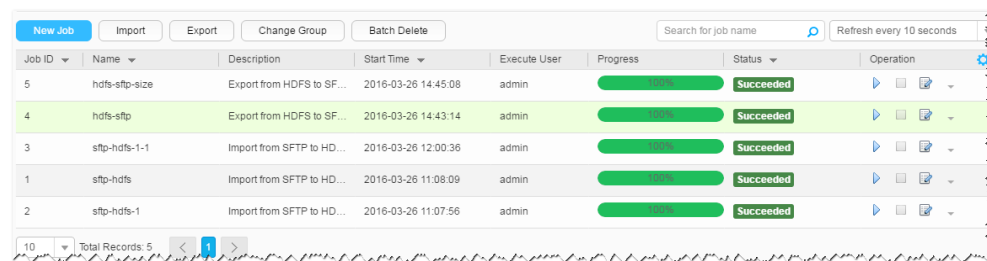
| Parameter | Description | Example Value |
|------------------|---|---------------|
| Output Directory | Directory for storing data imported into Hive.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /opt/tempfile |
| Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. You are advised to set the parameter to the maximum number of connections on the SFTP server. | 20 |
| Extractor Size | Hive does not support this parameter. Please set Extractors . | - |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-15 Viewing job details



----End

17.12.6 Typical Scenario: Importing Data from an FTP Server to HBase

Scenario

Use Loader to import data from an FTP server to HBase.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have obtained the username and password of the FTP server and the user has the read permission of the source files on the FTP server. If file name extension needs to be added after a source file is imported, the user must have the write permission of the source file.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- When using Loader to import data from the FTP server, the input paths and input path subdirectories of the FTP server and the name of the files in these directories do not contain any of the special characters `/'";,;`.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-16 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-17 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **ftp-connector**, click **Add**, set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**. Loader allows multiple FTP servers to be configured. Click **Add** to add the configuration information of multiple FTP servers.

Table 17-46 Connection parameters

| Parameter | Description | Example Value |
|-----------------------|---|---------------|
| FTP Server IP Address | IP address of the FTP server | ftpName |
| FTP Server Port | Port number of the FTP server | 22 |
| FTP Username | Username for accessing the FTP server | root |
| FTP Password | Password for accessing the FTP server | xxxx |
| FTP Mode | FTP access mode. Possible values are ACTIVE and PASSIVE . If this parameter is not set, FTP access is in passive mode by default. | PASSIVE |

| Parameter | Description | Example Value |
|-------------------------|--|---------------|
| FTP Protocol | <p>FTP protocol.</p> <ul style="list-style-type: none"> • FTP: indicates the FTP protocol. • SSL_EXPLICIT: indicates the explicit SSL protocol. • SSL_IMPLICIT: indicates the implicit SSL protocol. • TLS_EXPLICIT: indicates the explicit TLS protocol. • TLS_IMPLICIT: indicates the implicit TLS protocol. <p>If this parameter is not set, the FTP protocol is used by default.</p> | FTP |
| File Name Encoding Type | File name and file path encoding format supported by the FTP server. If this parameter is not set, the default format UTF-8 is used. | UTF-8 |

 **NOTE**

When multiple FTP servers are configured, the data in the specified directories of the servers is imported to HBase.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-47 Parameter description

| Parameter | Description | Example Value |
|------------|--|----------------------------|
| Input Path | <p>Input path or name of the source file on an FTP server. If multiple FTP server IP addresses are configured for the connector, you can set this parameter to multiple input paths separated with semicolons (;). Ensure that the number of input paths is the same as that of FTP servers configured for the connector.</p> <p>NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items .</p> | /opt/
tempfile;/o
pt |

| Parameter | Description | Example Value |
|-----------------|---|---------------|
| File Split Type | <p>Indicates whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each Map in the MapReduce task for data import.</p> <ul style="list-style-type: none"> ● FILE: indicates that the source file is split by file. That is, each Map processes one or multiple complete files, the same source file cannot be allocated to different Maps, and the source file directory structure is retained after data import. ● SIZE: indicates that the source file is split by size. That is, each Map processes input files of a certain size, and a source file can be divided and processed by multiple Maps. After data is stored in the output directory, the number of saved files is the same as that of Maps. The file name format is import_part_xxxx, where xxxx is a unique random number generated by the system. | FILE |
| Filter Type | <p>File filter condition. This parameter is used when Path Filter or File Filter is set.</p> <ul style="list-style-type: none"> ● WILDCARD: indicates using a wildcard. ● REGEX: indicates using a regular expression. ● If the parameter is not set, a wildcard is used by default. | WILDCARD |
| Path Filter | <p>Wildcard or regular expression for filtering the directories in the input path of the source files. This parameter is used when Filter Type is set. Input Path is not used for filtering. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. If this parameter is left empty, directories are not filtered.</p> <ul style="list-style-type: none"> ● ? matches a single character. ● * indicates multiple characters. ● Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | 1*,2*;1* |

| Parameter | Description | Example Value |
|---------------|---|-------------------|
| File Filter | <p>Wildcard or regular expression for filtering the file names of the source files. This parameter is used when Filter Type is set. Use semicolons (;) to separate the path filters on multiple servers and use commas (,) to separate the filter conditions of each server. This parameter cannot be left blank.</p> <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. <p>For example, when Filter type is set to WILDCARD, set the parameter to *; when Filter type is set to REGEX, set the parameter to \\.*.</p> | *.txt;*.csv;*.txt |
| Encoding Type | Source file encoding format, for example, UTF-8 and GBK. This parameter can be set only in text file import. | UTF-8 |
| Suffix | <p>File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. This parameter is valid only when the data source is a file system. You are advised to set this parameter in incremental data import.</p> <p>For example, if the parameter is set to .txt and the source file is test-loader.csv, the source file name is test-loader.csv.txt after export.</p> | .log |
| Compression | <p>Indicates whether to enable compressed transmission when FTP is used to export data.</p> <ul style="list-style-type: none"> • The value true indicates that compression is enabled. • The value false indicates that compression is disabled. | true |


Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-48](#).

Table 17-48 Input and output parameters of the operator

| Input Type | Output Type |
|------------------|--------------|
| CSV File Input | HBase Output |
| HTML Input | HBase Output |
| Fixed File Input | HBase Output |

Figure 17-18 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HBASE_BULKLOAD** or **HBASE_PUTLIST** based on the actual situation.

Table 17-49 Parameter description

| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|----------------|---|---------------|
| HBASE_BULKLOAD | Large data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |

| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|--------------------------|--|---------------|
| | | Clear data before import | Indicates whether to clear data in the original table before importing data. True indicates clearing data and False indicates not to clear data. If you do not set this parameter, the original table is not cleared by default. | true |
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. You are advised to set the parameter to the maximum number of connections on the FTP server. | 20 |
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |
| HBASE_P UTLIST | Small data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. | 20 |
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-19 Viewing job details

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|----------------|---------------------------|---------------------|--------------|----------------------------------|-----------|-----------|
| 5 | hdfs-stfp-size | Export from HDFS to SF... | 2016-03-26 14:45:08 | admin | <div style="width: 100%;"></div> | Succeeded | [Icons] |
| 4 | hdfs-stfp | Export from HDFS to SF... | 2016-03-26 14:43:14 | admin | <div style="width: 100%;"></div> | Succeeded | [Icons] |
| 3 | stfp-hdfs-1-1 | Import from SFTP to HD... | 2016-03-26 12:00:36 | admin | <div style="width: 100%;"></div> | Succeeded | [Icons] |
| 1 | stfp-hdfs | Import from SFTP to HD... | 2016-03-26 11:08:09 | admin | <div style="width: 100%;"></div> | Succeeded | [Icons] |
| 2 | stfp-hdfs-1 | Import from SFTP to HD... | 2016-03-26 11:07:56 | admin | <div style="width: 100%;"></div> | Succeeded | [Icons] |

----End

17.12.7 Typical Scenario: Importing Data from a Relational Database to HDFS or OBS

Scenario

Use Loader to import data from a relational database to HDFS or OBS.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HDFS or OBS directories and data involved in job execution.
- You have obtained the username and password of the relational database.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.
- Before the operation, perform the following steps:
 - a. Obtain the JAR package of the relational database driver and save it to the following directory on the active and standby Loader nodes: `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`.
 - b. Run the following command on the active and standby nodes as user `root` to modify the permission:


```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/
FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/
webapps/loader/WEB-INF/ext-lib
chown omm:wheel JAR package name
chmod 600 JAR package name
```
 - c. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **More** > **Restart Service**. Enter the password of the administrator to restart the Loader service.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-20 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-21 Basic Information



1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **generic-jdbc-connector** or a dedicated database connector (oracle-connector, oracle-partition-connector, or mysql-fastpath-connector), set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**.

 NOTE

- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When mysql-fastpath-connector is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManager nodes, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, install the MySQL client applications and tools following the instructions at <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>.

Table 17-50 generic-jdbc-connector connection parameters

| Parameter | Description | Example Value |
|----------------------------|---|--|
| Name | Name of a relational database connection | dbName |
| JDBC Driver Class | Name of a Java database connectivity (JDBC) driver class | oracle.jdbc.driver.OracleDriver |
| JDBC Connection String | JDBC connection string | jdbc:oracle:thin:@//10.16.0.1:1521/oradb |
| Username | Username for connecting to the database | omm |
| Password | Password for connecting to the database | xxxx |
| JDBC Connection Properties | JDBC connection attribute. Click Add to manually add the attribute. <ul style="list-style-type: none"> • Name: connection attribute name • Value: connection attribute value | <ul style="list-style-type: none"> • Name: socketTimeout • Value: 20 |

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-51 Parameter description

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Schema Name | Database schema name. This parameter exists in the Table name schema. | public |

| Parameter | Description | Example Value |
|-----------------------|--|--|
| Table Name | Database table name. This parameter exists in the Table name schema. | test |
| SQL Statement | SQL statement for Loader to query data to be imported in Table SQL statement mode. The SQL statement requires the query condition WHERE \${CONDITIONS} . Without this condition, the SQL statement cannot be run properly. An example SQL statement is as follows: select * from TABLE WHERE A>B and \${CONDITIONS} . If Table column names is set, the column specified by Table column names will replace the column queried in the SQL statement. This parameter cannot be set when Schema name or Table name is set.
NOTE
You can use macros to define SQL Where statements. For details, see Using Macro Definitions in Configuration Items . | select * from TABLE WHERE A>B and \${CONDITIONS} |
| Table Column Names | Table columns whose content is to be imported by Loader. Use commas (,) to separate multiple fields.

If the parameter is not set, all the columns are imported and the Select * order is used as the column location. | id,name |
| Partition Column Name | Database table column based on which to-be-imported data is determined. This parameter is used for partitioning in a Map job. You are advised to configure the primary key field.
NOTE <ul style="list-style-type: none"> A partition column must have an index. If no index exists, do not specify a partition column. If a partition column without an index is specified, the database server disk I/O will be busy, the access of other services to the database will be affected, and the import will take a long period. In multiple fields with indexes, select the field that has the most discrete value as the partition column. A partition column that is not discrete may result in load imbalance when multiple MapReduce jobs are imported. The sorting rules of partition columns must be case-sensitive. Otherwise, data may be lost during data import. You are not advised to select fields of the float or double type for the partition column. Otherwise, the records containing the minimum and maximum values of the partition column may fail to be imported due to precision issues. | id |

| Parameter | Description | Example Value |
|---------------------------------------|---|---------------|
| Nulls in Partition Column | Indicates whether to process records whose values are null in database table columns. <ul style="list-style-type: none"> • true: Records whose values are null are processed. • false: Records whose values are not null are processed. | true |
| Whether to Specify a Partition Column | Indicates whether to specify a partition column. | true |

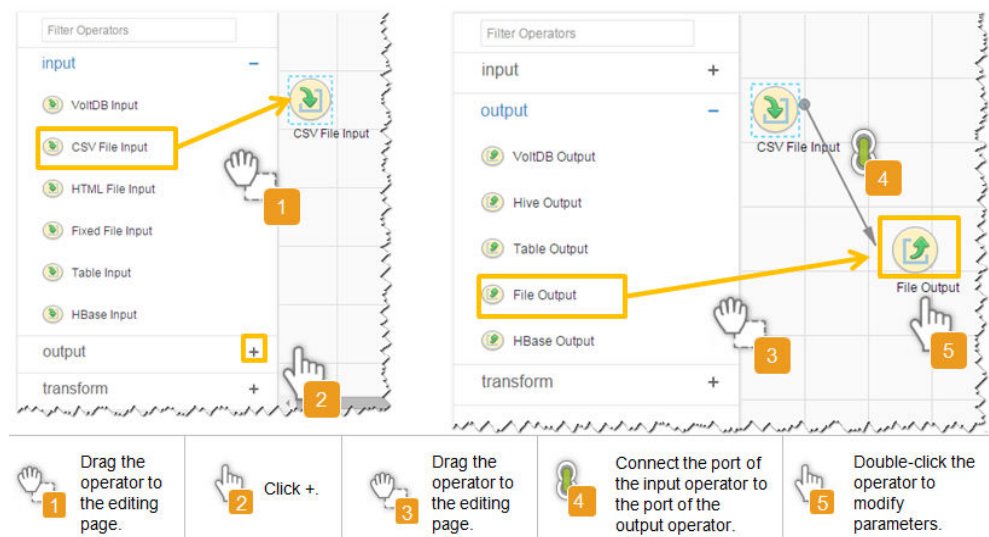
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-52](#).

Table 17-52 Input and output parameters of the operator

| Input Type | Output Type |
|-------------|-------------|
| Table input | File output |

Figure 17-22 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HDFS**.

Table 17-53 Parameter description

| Parameter | Description | Example Value |
|--------------------|---|---------------|
| File Type | Type of a file after the file is imported. The options are as follows: <ul style="list-style-type: none"> • TEXT_FILE: imports a text file and stores it as a text file. • SEQUENCE_FILE: imports a text file and stores it as a sequence file. • BINARY_FILE: imports files of any format by using binary streams. | TEXT_FILE |
| Compression Format | Compression format of files imported to HDFS or OBS. Select a format from the drop-down list. If you select NONE or leave this parameter blank, data is not compressed. | NONE |
| Output Directory | Directory for storing data imported into HDFS or OBS.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /user/test |

| Parameter | Description | Example Value |
|----------------|--|---------------|
| Operation | <p>Action during data import. When all data is to be imported from the input path to the destination path, the data is stored in a temporary directory and then copied from the temporary directory to the destination path. After the data is imported successfully, the data is deleted from the temporary directory. One of the following actions can be taken when duplicate file names exist during data transfer:</p> <ul style="list-style-type: none"> ● OVERRIDE: overrides the old file. ● RENAME: renames as new file. For a file without an extension, a string is added to the file name as the extension; for a file with an extension, a string is added to the extension. The string is unique. ● APPEND: adds the content of the new file to the end of the old file. This action only adds content regardless of whether the file can be used. For example, a text file can be used after this operation, while a compressed file cannot. ● IGNORE: reserves the old file and does not copy the new file. ● ERROR: stops the task and reports an error if duplicate file names exist. Transferred files are imported successfully, while files that have duplicate names and files that are not transferred fail to be imported. | OVERRIDE |
| Extractors | <p>Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. This parameter cannot be set when Extractor Size is set. The value must be less than or equal to 3000.</p> | - |
| Extractor Size | <p>Size of data processed by Maps that are started in a MapReduce task of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000. This parameter cannot be set when Extractors is set. When a relational database connector is used, Extractor Size is unavailable. You need to set Extractors.</p> | 1000 |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-23 Viewing job details

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|----------------|---------------------------|---------------------|--------------|----------------------------------|-----------|----------------|
| 5 | hdfs-sftp-size | Export from HDFS to SF... | 2016-03-26 14:45:08 | admin | <div style="width: 100%;"></div> | Succeeded | [Action icons] |
| 4 | hdfs-sftp | Export from HDFS to SF... | 2016-03-26 14:43:14 | admin | <div style="width: 100%;"></div> | Succeeded | [Action icons] |
| 3 | sftp-hdfs-1-1 | Import from SFTP to HD... | 2016-03-26 12:00:36 | admin | <div style="width: 100%;"></div> | Succeeded | [Action icons] |
| 1 | sftp-hdfs | Import from SFTP to HD... | 2016-03-26 11:08:09 | admin | <div style="width: 100%;"></div> | Succeeded | [Action icons] |
| 2 | sftp-hdfs-1 | Import from SFTP to HD... | 2016-03-26 11:07:56 | admin | <div style="width: 100%;"></div> | Succeeded | [Action icons] |

----End

17.12.8 Typical Scenario: Importing Data from a Relational Database to HBase

Scenario

Use Loader to import data from a relational database to HBase.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HBase tables or phoenix tables that are used during job execution.
- You have obtained the username and password of the relational database.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.
- Before the operation, perform the following steps:

- Obtain the JAR package of the relational database driver and save it to the following directory on the active and standby Loader nodes: `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`.

- Run the following command on the active and standby nodes as user `root` to modify the permission:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/
FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/
webapps/loader/WEB-INF/ext-lib
```

```
chown omm:wheel JAR package name
```

```
chmod 600 JAR package name
```

- c. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **More** > **Restart Service**. Enter the password of the administrator to restart the Loader service.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-24 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-25 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **generic-jdbc-connector** or a dedicated database connector (oracle-connector, oracle-partition-connector, or mysql-fastpath-connector), set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**.

 NOTE

- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When mysql-fastpath-connector is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManager nodes, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, install the MySQL client applications and tools following the instructions at <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>.

Table 17-54 generic-jdbc-connector connection parameters

| Parameter | Description | Example Value |
|----------------------------|---|--|
| Name | Name of a relational database connection | dbName |
| JDBC Driver Class | Name of a JDBC driver class | oracle.jdbc.driver.OracleDriver |
| JDBC Connection String | JDBC connection string | jdbc:oracle:thin:@//10.16.0.1:1521/oradb |
| Username | Username for connecting to the database | omm |
| Password | Password for connecting to the database | xxxx |
| JDBC Connection Properties | JDBC connection attribute. Click Add to manually add the attribute. <ul style="list-style-type: none"> • Name: connection attribute name • Value: connection attribute value | <ul style="list-style-type: none"> • Name: socketTimeout • Value: 20 |

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-55 Parameter description

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Schema Name | Database schema name. This parameter exists in the Table name schema. | dbo |

| Parameter | Description | Example Value |
|--------------------|---|--|
| Table Name | Database table name. This parameter exists in the Table name schema. | test |
| SQL Statement | SQL statement for Loader to query data to be imported in Table SQL statement mode. The SQL statement requires the query condition WHERE \$ {CONDITIONS} . Without this condition, the SQL statement cannot be run properly. An example SQL statement is as follows: select * from TABLE WHERE A>B and \$ {CONDITIONS} . If Table column names is set, the column specified by Table column names will replace the column queried in the SQL statement. This parameter cannot be set when Schema name or Table name is set.
NOTE
You can use macros to define SQL Where statements. For details, see Using Macro Definitions in Configuration Items . | select * from test
where \$
{CONDITIONS} |
| Table Column Names | Table columns whose content is to be imported by Loader. Use commas (,) to separate multiple fields.

If the parameter is not set, all the columns are imported and the Select * order is used as the column location. | - |

| Parameter | Description | Example Value |
|---------------------------------------|---|---------------|
| Partition Column Name | <p>Database table column based on which to-be-imported data is determined. This parameter is used for partitioning in a Map job. You are advised to configure the primary key field.</p> <p>NOTE</p> <ul style="list-style-type: none"> • A partition column must have an index. If no index exists, do not specify a partition column. If a partition column without an index is specified, the database server disk I/O will be busy, the access of other services to the database will be affected, and the import will take a long period. • In multiple fields with indexes, select the field that has the most discrete value as the partition column. A partition column that is not discrete may result in load imbalance when multiple MapReduce jobs are imported. • The sorting rules of partition columns must be case-sensitive. Otherwise, data may be lost during data import. • You are not advised to select fields of the float or double type for the partition column. Otherwise, the records containing the minimum and maximum values of the partition column may fail to be imported due to precision issues. | id |
| Nulls in Partition Column | <p>Indicates whether to process records whose values are null in database table columns.</p> <ul style="list-style-type: none"> • true: Records whose values are null are processed. • false: Records whose values are not null are processed. | true |
| Whether to Specify a Partition Column | Indicates whether to specify a partition column. | true |

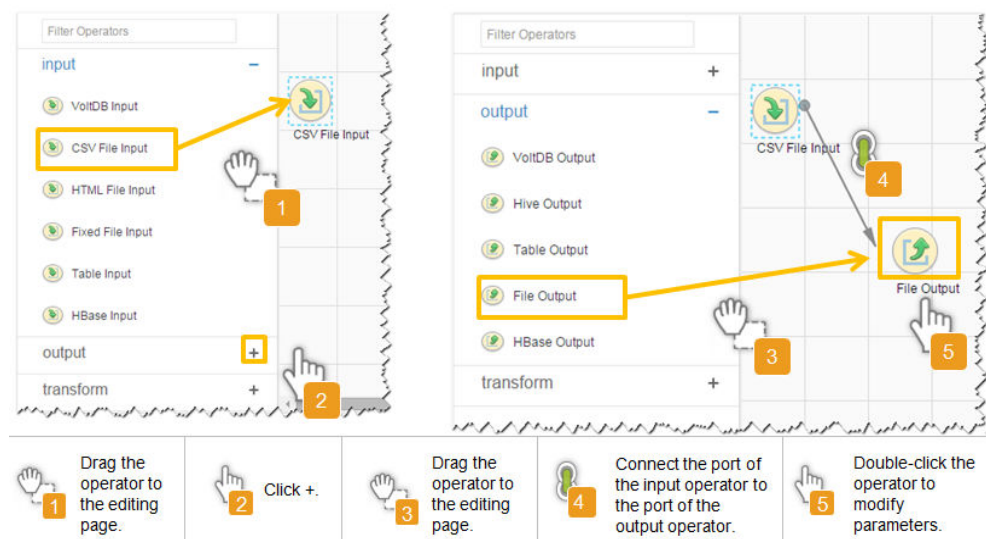
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-56](#).

Table 17-56 Input and output parameters of the operator

| Input Type | Output Type |
|-------------|--------------|
| Table input | HBase output |

Figure 17-26 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HBASE_BULKLOAD** or **HBASE_PUTLIST** based on the actual situation.

Table 17-57 Parameter description

| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|----------------|---|---------------|
| HBASE_BULKLOAD | Large data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |

| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|--------------------------|--|---------------|
| | | Clear data before import | Indicates whether to clear data in the original table before importing data. True indicates clearing data and False indicates not to clear data. If you do not set this parameter, the original table is not cleared by default. | true |
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. | 20 |
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |
| HBASE_P UTLIST | Small data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. | true |
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-27 Viewing job details

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|----------------|---------------------------|---------------------|--------------|---|-----------|---|
| 5 | hdfs-stfp-size | Export from HDFS to SF... | 2016-03-26 14:45:08 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | ▶ ⌵ ⌵ ⌵ |
| 4 | hdfs-stfp | Export from HDFS to SF... | 2016-03-26 14:43:14 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | ▶ ⌵ ⌵ ⌵ |
| 3 | stfp-hdfs-1-1 | Import from SFTP to HD... | 2016-03-26 12:00:36 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | ▶ ⌵ ⌵ ⌵ |
| 1 | stfp-hdfs | Import from SFTP to HD... | 2016-03-26 11:08:09 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | ▶ ⌵ ⌵ ⌵ |
| 2 | stfp-hdfs-1 | Import from SFTP to HD... | 2016-03-26 11:07:56 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | ▶ ⌵ ⌵ ⌵ |

----End

17.12.9 Typical Scenario: Importing Data from a Relational Database to Hive

Scenario

Use Loader to import data from a relational database to Hive.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the Hive tables that are used during job execution.
- You have obtained the username and password of the relational database.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.
- Before the operation, perform the following steps:
 - a. Obtain the JAR package of the relational database driver and save it to the following directory on the active and standby Loader nodes: `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`.
 - b. Run the following command on the active and standby nodes as user `root` to modify the permission:


```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/
FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/
webapps/loader/WEB-INF/ext-lib
chown omm:wheel JAR package name
chmod 600 JAR package name
```
 - c. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **More** > **Restart Service**. Enter the password of the administrator to restart the Loader service.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-28 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-29 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **generic-jdbc-connector** or a dedicated database connector (oracle-connector, oracle-partition-connector, or mysql-fastpath-connector), set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**.

 NOTE

- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When mysql-fastpath-connector is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManager nodes, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, install the MySQL client applications and tools following the instructions at <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>.

Table 17-58 generic-jdbc-connector connection parameters

| Parameter | Description | Example Value |
|----------------------------|---|--|
| Name | Name of a relational database connection | dbName |
| JDBC Driver Class | Name of a JDBC driver class | oracle.jdbc.driver.OracleDriver |
| JDBC Connection String | JDBC connection string | jdbc:oracle:thin:@//10.16.0.1:1521/oradb |
| Username | Username for connecting to the database | omm |
| Password | Password for connecting to the database | xxxx |
| JDBC Connection Properties | <p>JDBC connection attribute. Click Add to manually add the attribute.</p> <ul style="list-style-type: none"> • Name: connection attribute name • Value: connection attribute value <p>NOTICE
If a general connector is used to connect to the MySQL database and there a large amount of data, you need to set useCursorFetch=true in the JDBC connection string.</p> | <ul style="list-style-type: none"> • Name: socketTimeout • Value: 20 |

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-59 Parameter description

| Parameter | Description | Example Value |
|--------------------|---|--|
| Schema Name | Database schema name. This parameter exists in the Table name schema. | dbo |
| Table Name | Database table name. This parameter exists in the Table name schema. | test |
| SQL Statement | SQL statement for Loader to query data to be imported in Table SQL statement mode. The SQL statement requires the query condition WHERE \$ {CONDITIONS} . Without this condition, the SQL statement cannot be run properly. An example SQL statement is as follows: select * from TABLE WHERE A>B and \$ {CONDITIONS} . If Table column names is set, the column specified by Table column names will replace the column queried in the SQL statement. This parameter cannot be set when Schema name or Table name is set.
NOTE
You can use macros to define SQL Where statements. For details, see Using Macro Definitions in Configuration Items . | select * from test where \$ {CONDITIONS} |
| Table Column Names | Table columns whose content is to be imported by Loader. Use commas (,) to separate multiple fields.
If the parameter is not set, all the columns are imported and the Select * order is used as the column location. | - |

| Parameter | Description | Example Value |
|---------------------------------------|---|---------------|
| Partition Column Name | <p>Database table column based on which to-be-imported data is determined. This parameter is used for partitioning in a Map job. You are advised to configure the primary key field.</p> <p>NOTE</p> <ul style="list-style-type: none"> • A partition column must have an index. If no index exists, do not specify a partition column. If a partition column without an index is specified, the database server disk I/O will be busy, the access of other services to the database will be affected, and the import will take a long period. • In multiple fields with indexes, select the field that has the most discrete value as the partition column. A partition column that is not discrete may result in load imbalance when multiple MapReduce jobs are imported. • The sorting rules of partition columns must be case-sensitive. Otherwise, data may be lost during data import. • You are not advised to select fields of the float or double type for the partition column. Otherwise, the records containing the minimum and maximum values of the partition column may fail to be imported due to precision issues. | id |
| Nulls in Partition Column | <p>Indicates whether to process records whose values are null in database table columns.</p> <ul style="list-style-type: none"> • true: Records whose values are null are processed. • false: Records whose values are not null are processed. | true |
| Whether to Specify a Partition Column | Indicates whether to specify a partition column. | true |

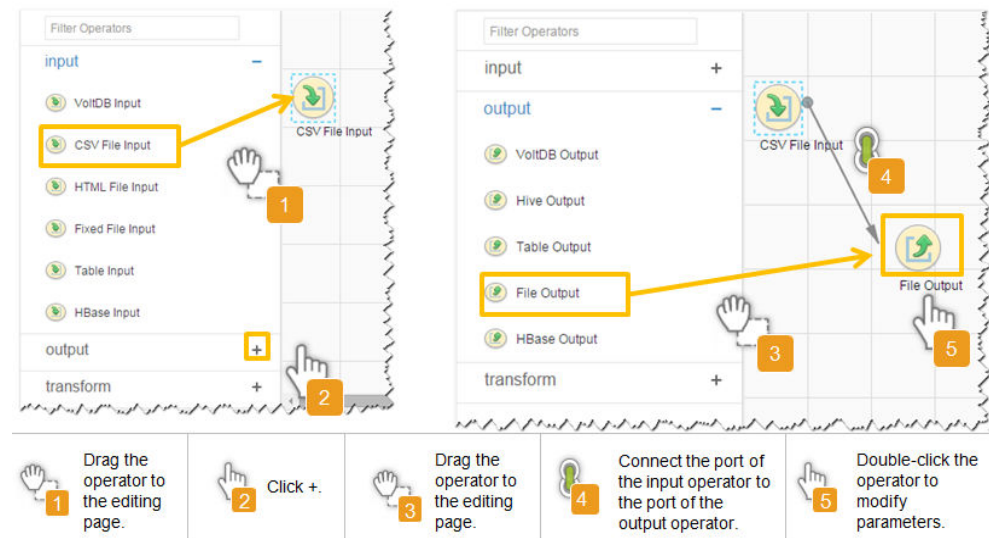
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-60](#).

Table 17-60 Input and output parameters of the operator

| Input Type | Output Type |
|-------------|-------------|
| Table input | Hive output |

Figure 17-30 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HIVE**.

Table 17-61 Parameter description

| Parameter | Description | Example Value |
|------------------|---|---------------|
| Output Directory | Directory for storing data imported into Hive.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /opt/tempfile |
| Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. You are advised to set the parameter to the maximum number of connections on the SFTP server. | 20 |
| Extractor Size | Hive does not support this parameter. Please set Extractors . | - |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-31 Viewing job details

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|---------------|---------------------------|---------------------|--------------|----------|-----------|-----------|
| 5 | hdfs-ftp-size | Export from HDFS to SF... | 2016-03-26 14:45:08 | admin | | Succeeded | |
| 4 | hdfs-ftp | Export from HDFS to SF... | 2016-03-26 14:43:14 | admin | | Succeeded | |
| 3 | ftp-hdfs-1-1 | Import from SFTP to HD... | 2016-03-26 12:00:36 | admin | | Succeeded | |
| 1 | ftp-hdfs | Import from SFTP to HD... | 2016-03-26 11:08:09 | admin | | Succeeded | |
| 2 | ftp-hdfs-1 | Import from SFTP to HD... | 2016-03-26 11:07:56 | admin | | Succeeded | |

----End

17.12.10 Typical Scenario: Importing Data from HDFS or OBS to HBase

Scenario

Use Loader to import data from HDFS or OBS to HBase.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HDFS or OBS directories and data involved in job execution.
- You have had the permission to access the HBase tables or phoenix tables that are used during job execution.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- When using Loader to import data from HDFS or OBS, the input paths and input path subdirectories of HDFS or OBS and the name of the files in these directories do not contain any of the special characters /'";,.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.

Procedure


Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).

2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-32 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-33 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **hdfs-connector**, set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-62 Parameter description

| Parameter | Description | Example Value |
|---------------|---|--------------------|
| Input Path | Input path of source files in HDFS or OBS
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /
user/
test |
| Path Filter | Wildcard for filtering the directories in the input paths of the source files. Input Path is not used for filtering. If there are multiple filter conditions, use commas (,) to separate them. If the parameter is empty, the directories are not filtered. The regular expression filtering is not supported. | * |
| File Filter | Wildcard for filtering the file names of the source files. If there are multiple filter conditions, use commas (,) to separate them. The value cannot be left blank. The regular expression filtering is not supported. | * |
| Encoding Type | Source file encoding format, for example, UTF-8. This parameter can be set only in text file import. | UTF-8 |
| Suffix | File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. | .log |

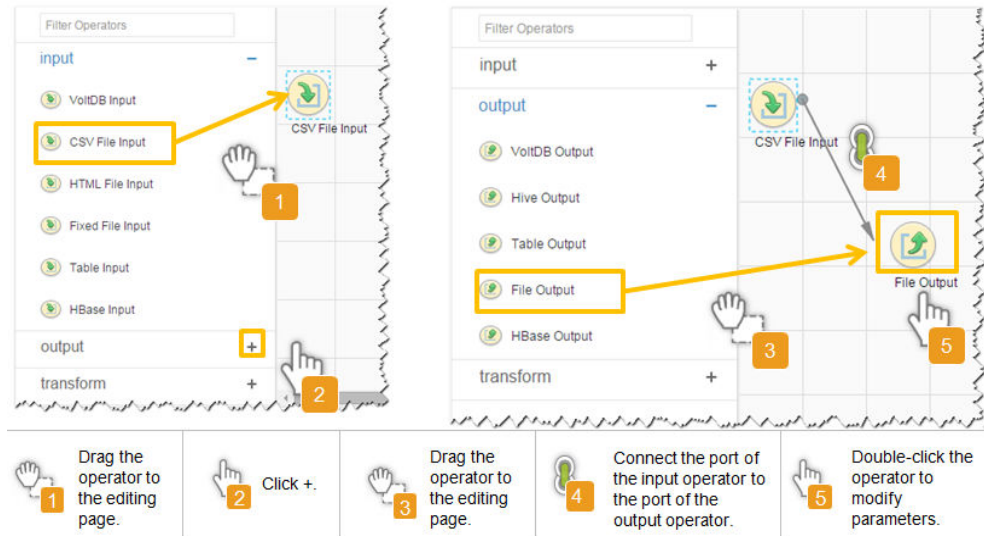
Setting Data Transformation

- Step 5** Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-63](#).

Table 17-63 Input and output parameters of the operator

| Input Type | Output Type |
|------------------|--------------|
| CSV File Input | HBase Output |
| HTML Input | HBase Output |
| Fixed File Input | HBase Output |

Figure 17-34 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **HBASE_BULKLOAD** or **HBASE_PUTLIST** based on the actual situation.

Table 17-64 Parameter description

| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|--------------------------|--|---------------|
| HBASE_BULKLOAD | Large data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| | | Clear data before import | Indicates whether to clear data in the original table before importing data. True indicates clearing data and False indicates not to clear data. If you do not set this parameter, the original table is not cleared by default. | true |
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. | 20 |

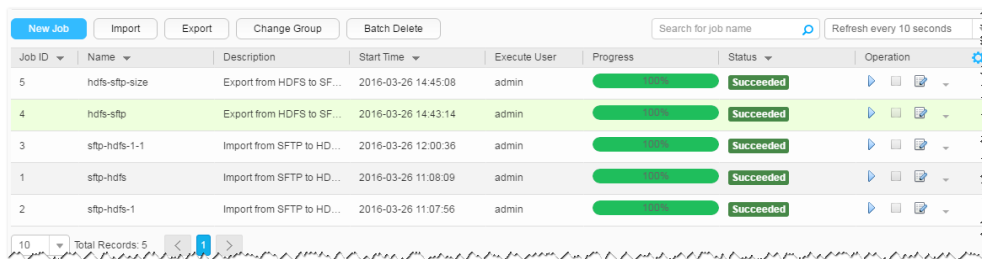
| Storage Type | Applicable Scenario | Parameter | Description | Example Value |
|----------------|---------------------|----------------|---|---------------|
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |
| HBASE_PUTLIS T | Small data volume | HBase Instance | HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| | | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. | 20 |
| | | Extractor Size | HBase does not support this parameter. Please set Extractors . | - |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-35 Viewing job details



----End

17.12.11 Typical Scenario: Importing Data from a Relational Database to ClickHouse

Scenario

This section describes how to use Loader to import data from a relational database to ClickHouse using MySQL as an example.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have created a ClickHouse table, and you have operation permissions on the table during job execution.
- You have obtained the user name and password of the MySQL database.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.
- Before the operation, perform the following steps:
 - a. Obtain the MySQL client JAR file (for example, **mysqlclient-5.8.1.jar**) from the MySQL database installation path and save it to the following directory on the active and standby Loader nodes: **`\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`**.
 - b. Obtain the **clickhouse-jdbc-*.jar** file from the ClickHouse installation directory and save it in the **`\${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`** directory on the active and standby Loader nodes.
 - c. Run the following command on the active and standby nodes as user **root** to modify the permission:

```
cd `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/  
FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/  
webapps/loader/WEB-INF/ext-lib  
chown omm:wheel JAR file name  
chmod 600 JAR file name
```
 - d. Log in to FusionInsight Manager, choose **Cluster > Services > Loader > More > Restart Service** and enter the administrator password to restart the Loader service.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster > Name of the desired cluster > Services > Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-36 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-37 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **generic-jdbc-connector** or a dedicated database connector (oracle-connector, oracle-partition-connector, or mysql-fastpath-connector), set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**.

NOTE

- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When mysql-fastpath-connector is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManager nodes, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, install the MySQL client applications and tools following the instructions at <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>.

Table 17-65 generic-jdbc-connector connection parameters

| Parameter | Description | Example |
|------------------------|--|---|
| Name | Name of a relational database connection | mysql_test |
| JDBC Driver Class | Name of a JDBC driver class | com.mysql.jdbc.Driver |
| JDBC Connection String | JDBC connection string | jdbc:mysql://10.254.144.102:3306/test?useUnicode=true&characterEncoding=UTF-8 |
| Username | Username for connecting to the database | root |
| Password | Password for connecting to the database | xxxx |

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, configure the data source information. Currently, only **Table name** is supported.

Table 17-66 Parameter description

| Parameter | Description | Example |
|-----------------------|--|---------|
| Schema name | Schema name of the specified database | public |
| Table name | Table name | test |
| Table column names | Names of the columns to be imported | id,name |
| Need partition column | Currently, only the unspecified partition mode is supported. | false |

Setting Data Transformation

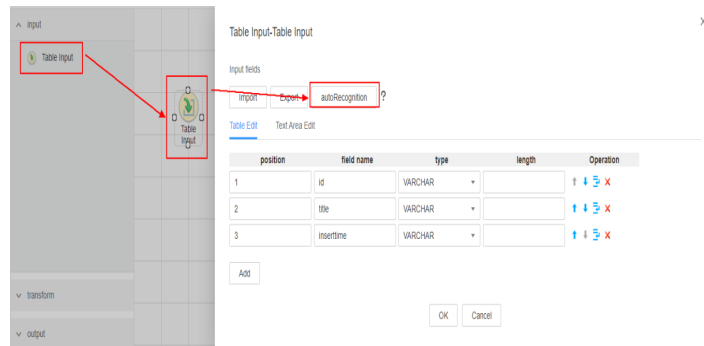
Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-67](#).

Table 17-67 Input and output parameters of the operator

| Input Type | Output Type |
|------------|-------------|
| MySQL | ClickHouse |

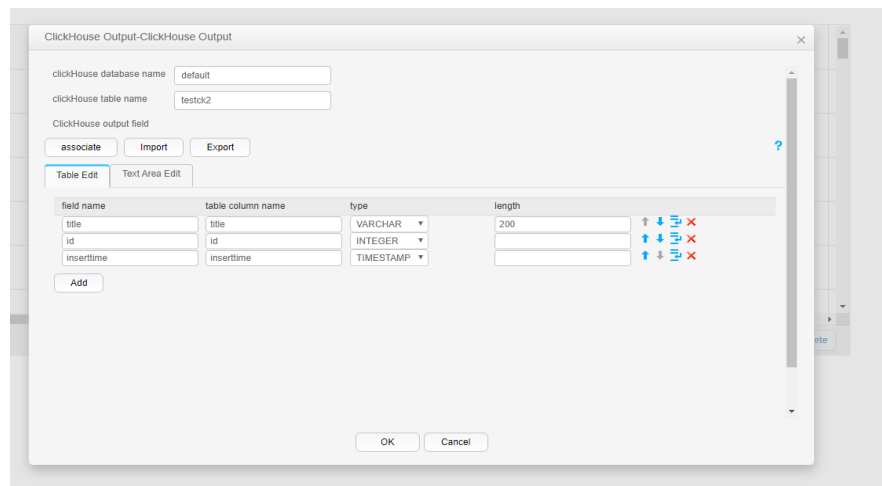
Drag **Table Input** to the grid, double-click **Table Input**, and select **autoRecognition**, as shown in **Figure 17-38**.

Figure 17-38 Operator input



Drag **ClickHouse Output** to the grid, double-click **Table Output**, and select **associate** or manually edit the table to correspond to the input table, as shown in **Figure 17-39**.

Figure 17-39 Operator output



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **CLICKHOUSE**.

Table 17-68 Parameter description

| Parameter | Description | Example |
|---------------------|----------------------------|---------|
| Storage type | Select CLICKHOUSE . | - |
| ClickHouse instance | Select ClickHouse . | - |

| Parameter | Description | Example |
|--------------------------|--|---------|
| Clear data before import | Select true or false .
NOTE
If you select true and the table to be imported is a ClickHouse distributed table, you need to manually delete the data from the local table corresponding to the ClickHouse distributed table before your import. | true |

Step 7 Click **Save and Run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-40 Viewing job details

| | | | | | | |
|---|---------------------------|--------------------------|---------------------|-----------|--|---|
| 8 | mysql_clickhouse-field... | Import from RDB to CL... | 2021-06-23 19:53:12 | admintest | ■ Succeeded | ▶ 📄 📄 ⌵ |
| 7 | mysql_clickhouse-field | Import from RDB to CL... | 2021-06-23 19:53:11 | admintest | ■ Succeeded | ▶ 📄 📄 ⌵ |
| 6 | mysql_clickhouse | Import from RDB to CL... | 2021-06-23 19:16:09 | admintest | ■ Succeeded | ▶ 📄 📄 ⌵ |

Step 9 On the ClickHouse client, check whether the data in the ClickHouse table is the same as that in the MySQL table.

----End

17.12.12 Typical Scenario: Importing Data from HDFS to ClickHouse

Scenario

Use Loader to import data from HDFS to ClickHouse.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HDFS directories and data involved in job execution.
- You have created a ClickHouse table, and you have operation permissions on the table during job execution.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- When using Loader to import data from HDFS, the input paths and input path subdirectories of HDFS and the name of the files in these directories do not contain any of the special characters `/'":;`.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-41 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-42 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Import**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **hdfs-connector**, set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set the data source information.

Table 17-69 Parameter description

| Parameter | Description | Example |
|---------------|---|--------------------|
| Input Path | Input path of source files in HDFS
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /
user/
test |
| Path Filter | Wildcard for filtering the directories in the input paths of the source files. Input Path is not used for filtering. If there are multiple filter conditions, use commas (,) to separate them. If the parameter is empty, the directories are not filtered. The regular expression filtering is not supported. | * |
| File Filter | Wildcard for filtering the file names of the source files. If there are multiple filter conditions, use commas (,) to separate them. The value cannot be left blank. The regular expression filtering is not supported. | * |
| Encoding Type | Source file encoding format, for example, UTF-8. This parameter can be set only in text file import. | UTF-8 |
| Suffix | File name extension added to a source file after the source file is imported. If this parameter is empty, no file name extension is added to the source file. | .log |

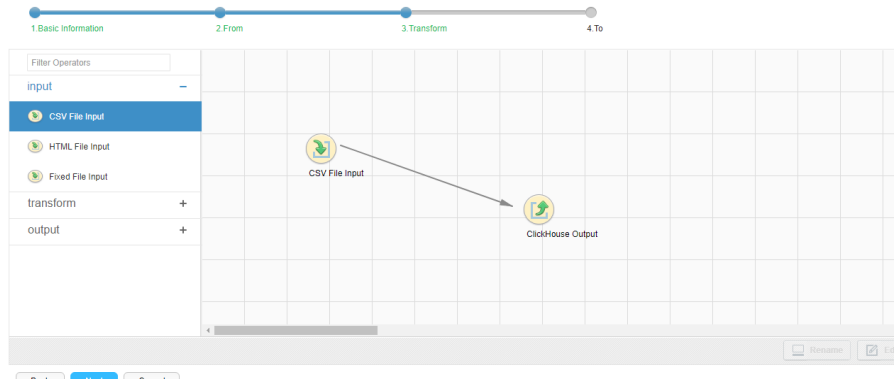
Setting Data Transformation

- Step 5** Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-70](#).

Table 17-70 Input and output parameters of the operator

| Input Type | Output Type |
|----------------|-------------------|
| CSV File Input | ClickHouse Output |

Figure 17-43 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set **Storage type** to **CLICKHOUSE** based on the actual situation.

Table 17-71 Parameter description

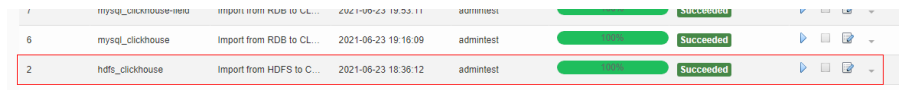
| Storage Type | Parameter | Description | Example |
|--------------|--------------------------|---|------------|
| CLICKHOUSE | ClickHouse instance | ClickHouse service instance that Loader selects from all available ClickHouse service instances in the cluster. If the selected ClickHouse service instance is not added to the cluster, the ClickHouse job cannot be run properly. | ClickHouse |
| | Clear data before import | Whether to clear data in the original table before importing data true indicates clearing data and false indicates not to clear data. If you do not set this parameter, the original table is not cleared by default.
NOTE
If you select true and the table to be imported is a ClickHouse distributed table, you need to manually delete the data from the local table corresponding to the ClickHouse distributed table before your import. | false |
| | Extractors | Number of Maps that are started at the same time in a MapReduce task of a data configuration operation. The value must be less than or equal to 3000. | 20 |
| | Extractor size | ClickHouse does not support this parameter. Please set Extractors . | - |
| | Number | Number of Map tasks. | - |

Step 7 Click **Save and Run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the Loader web UI. When **Status** is **Succeeded**, the job is complete.

Figure 17-44 Viewing job details



| ID | Name | Description | Start Time | User | Status |
|----|-----------------------|--------------------------|---------------------|-----------|-----------|
| 7 | mysql_clickhouse-1850 | Import from RDB to CL... | 2021-06-23 19:03:11 | admintest | Succeeded |
| 6 | mysql_clickhouse | Import from RDB to CL... | 2021-06-23 19:16:09 | admintest | Succeeded |
| 2 | hdfs_clickhouse | Import from HDFS to C... | 2021-06-23 18:36:12 | admintest | Succeeded |

Step 9 On the ClickHouse client, check whether the data in the ClickHouse table is the same as that imported from HDFS.

----End

17.13 Exporting Data

17.13.1 Overview

The section "Exporting Data" applies to MRS 3.x or later.

Description

Loader is an extract, transform, and load (ETL) tool for exchanging data and files between MRS and relational databases and file systems. You can use the Loader to export data or files from MRS to the relational databases or file systems.

Loader supports the following data export modes:

- Exporting data from HDFS or OBS to an SFTP server
- Exporting data from HDFS or OBS to a relational database
- Exporting data from HBase to an SFTP server
- Exporting data from HBase to a relational database
- Exporting data from Phoenix tables to an SFTP server
- Exporting data from Phoenix tables to a relational database
- Exporting data from Hive to an SFTP server
- Exporting Data from Hive to a Relational Database
- Exporting data from HBase to HDFS or OBS in the same cluster

The MRS needs to connect to the data source to exchange data and files with the external data source. The following connectors are used to configure connection parameters for different types of data sources:

- **generic-jdbc-connector**: relational database connector
- **hdfs-connector**: HDFS data source connector
- **oracle-connector**: specifies a dedicated connector for Oracle databases. **row_id** serves as partition columns. Compared with generic-jdbc-connector,

Map tasks are more evenly distributed on oracle-connector, and whether the partition columns have created indexes does not affect oracle-connector.

- **mysql-fastpath-connector**: specifies a dedicated connector for MySQL databases. Data is imported and exported by using the mysqldump and mysqlimport tools of MySQL. Compared with generic-jdbc-connector, the data import and export speed is faster.
- **sftp-connector**: SFTP data source connector
- **oracle-partition-connector**: connector that supports the Oracle partition feature, which is used to optimize the import and export of Oracle partition tables.

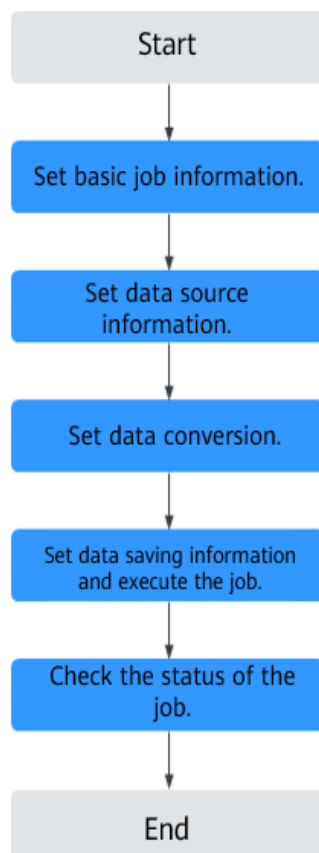
 **NOTE**

- You are advised to deploy the SFTP server, database server, and Loader into separate subnets to ensure secure data export.
- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because it is optimized for specific database types.
- When **mysql-fastpath-connector** is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManagers, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, see <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>. Install the MySQL client applications and tools.
- When oracle-connector is used, the connection user must be granted the select permission on the following system catalogs or views:
dba_tab_partitions, dba_constraints, dba_tables, dba_segments, v\$instance, dba_objects, v\$instance, dba_extents, dba_tab_partitions and dba_tab_subpartitions.
- When oracle-partition-connector is used, the connection user must be granted the select permission on the following system catalogs: dba_objects and dba_extents.

Export Process

A data export job can be executed in the Loader WebUI. [Figure 17-45](#) shows the export process.

Figure 17-45 Export process



Loader jobs can also be updated and executed using shell scripts. In this mode, the Loader client that has been installed needs to be configured.

17.13.2 Using Loader to Export Data

Scenario

This task enables you to export data from MRS to external data sources.

Generally, users can manually manage data import and export jobs on the Loader UI. To use shell scripts to update and run Loader jobs, you must configure the installed Loader client.

Prerequisites

- You have obtained the service user name and password for creating a Loader job.
- You have had the permission to access the HDFS directories, HBase tables, and data involved in job execution.
- You have obtained the user name and password used by an external data source (SFTP server or relational database).
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.

- When using Loader to export data from HDFS or OBS, the input paths and input path subdirectories of the HDFS or OBS data source and the name of the files in these directories do not contain any of the following special characters: \|"';,;
- If the job requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.

Procedure

Step 1 Check whether data is exported from Loader to a relational database for the first time.

- If yes, go to [Step 2](#).
- If no, go to [Step 3](#).

Step 2 Modify the permission on the JAR package of the RDS driver.

1. Obtain the JAR package of the relational database driver and save it to the following directory on the active and standby Loader nodes: `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`.
2. Run the following command on the active and standby nodes as user `root` to modify the permission:

```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib
```

```
chown omm:wheel JAR package name
```


```
chmod 600 JAR package name
```

3. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Service** > **Loader** > **More** > **Restart**. Enter the password of the administrator to restart the Loader service.

Step 3 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-46 Loader web UI



Step 4 Create a Loader data import job. Click **New Job**. Select the required job type in **1. Basic Information** and click **Next**.

1. Set **Name** to the job name and **Type** to **Export**.
2. Select a connection for **Connection**. By default, no connection is created. Click **Add** to create a connection, and then click **Test** to test whether the connection is available. Click **OK** when the system displays a message indicates that the test is successful.

Table 17-72 Connection configuration parameters

| Connector Type | Parameter | Description |
|--------------------------|----------------------------|--|
| generic-jdbc-connector | JDBC Driver Class | Specifies the name of a JDBC driver class. |
| | JDBC Connection String | Specifies the JDBC connection string. |
| | Username | Specifies the username for connecting to the database. |
| | Password | Specifies the password for connecting to the database. |
| | JDBC Connection Properties | Specifies JDBC connection attributes. Click Add to manually add connection attributes. <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |
| hdfs-connector | - | - |
| oracle-connector | JDBC Connection String | Specifies connection string for a user to connect to the database. |
| | Username | Specifies the username for connecting to the database. |
| | Password | Specifies the password for connecting to the database. |
| | Connection Properties | Specifies connection attributes. Click Add to manually add connection attributes. <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |
| mysql-fastpath-connector | JDBC Connection String | Specifies the JDBC connection string. |
| | Username | Specifies the username for connecting to the database. |

| Connector Type | Parameter | Description |
|----------------------------|------------------------|---|
| | Password | Specifies the password for connecting to the database. |
| | Connection Properties | Specifies connection attributes. Click Add to manually add connection attributes. <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |
| sftp-connector | SFTP Server IP | Specifies the IP address of the SFTP server. |
| | SFTP Server Port | Specifies the port number of the SFTP server. |
| | SFTP Username | Specifies the username for accessing the SFTP server. |
| | SFTP Password | Specifies the password for accessing the SFTP server. |
| | SFTP Public Key | Specifies public key of the SFTP server. |
| oracle-partition-connector | JDBC Driver Class | Specifies the name of a Java database connectivity (JDBC) driver class. |
| | JDBC Connection String | Specifies the JDBC connection string. |
| | Username | Specifies the username for connecting to the database. |
| | Password | Specifies the password for connecting to the database. |
| | Connection Properties | Specifies connection attributes. Click Add to manually add connection attributes. <ul style="list-style-type: none"> - Name: connection attribute name - Value: connection attribute value |

3. Set **Group** to the group to which the job belongs. By default, there is no created group. Click **Add** to create a group and click **OK**.
4. **Queue** indicates that Loader tasks are executed in a specified Yarn queue. The default value is **root.default**, which indicates that the tasks are executed in the **default** queue.
5. Set **Priority** to the priority of Loader tasks in the specified Yarn queue. The options can be **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, or **VERY_HIGH**. The default value is **NORMAL**.

Step 5 In the **2. Input Settings** area, set the data source and click **Next**.

 NOTE

When creating or editing a Loader job, you can use macro definitions when configuring parameters such as the SFTP path, HDFS/OBS path, and Where condition of SQL. For details, see [Using Macro Definitions in Configuration Items](#) .

Table 17-73 List of input configuration parameters

| Source File Type | Parameter | Description |
|------------------|-----------------|--|
| HDFS/OBS | Input Directory | Specifies the input path when data is exported from HDFS or OBS. |
| | Path Filter | Specifies the wildcard for filtering the directories in the input paths of the source files. Input Directory is not used in filtering. If there are multiple filter conditions, use commas (,) to separate them. If the value is empty, the directory is not filtered. The regular expression filtering is not supported. |
| | File Filter | Specifies the wildcard for filtering the file names of the source files. If there are multiple filter conditions, use commas (,) to separate them. The value cannot be left blank. The regular expression filtering is not supported. |
| | File Type | Specifies the file import type. <ul style="list-style-type: none"> • TEXT_FILE: imports a text file and stores it as a text file. • SEQUENCE_FILE: imports a text file and stores it as a sequence file. • BINARY_FILE: imports files of any format by using binary streams. |
| | File Split Type | Specifies whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each map in the MapReduce task for data export. |
| | Extractors | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. This parameter cannot be set when Extractor Size is set. The value must be less than or equal to 3000. |

| Source File Type | Parameter | Description |
|------------------|----------------|--|
| | Extractor size | Specifies the size of data processed by maps that are started in a MapReduce job of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000 . This parameter cannot be set when Extractors is set. When a relational database connector is used, Extractor size is unavailable. You need to set Extractors . |
| HBASE | HBase Instance | Specifies the HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. |
| | Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. |
| HIVE | Hive instance | Specifies the Hive service instance that Loader selects from all available Hive service instances in the cluster. If the selected Hive service instance is not added to the cluster, the Hive job cannot run properly. |
| | Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. |
| SPARK | Spark instance | Only SparkSQL can access Hive data. Specifies the SparkSQL service instance that Loader selects from all available SparkSQL service instances in the cluster. If the selected Spark service instance is not added to the cluster, the Spark job cannot be run properly. |
| | Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. |

Step 6 In the **3. Convert** area, set the conversion operations during data transmission.

Check whether source data values in the data operation job created by the Loader can be directly used without conversion, including upper and lower case conversion, cutting, merging, and separation.

- If yes, click **Next**.

- If no, perform [Step 6.1](#) to [Step 6.4](#).
- 1. No created conversion step exists by default. Drag an example conversion step on the left to the edit box to create a new conversion step.
- 2. Conversion step types must be selected based on service requirements. A complete conversion process includes the following types:
 - a. Input type. Only one conversion step can be added. This parameter is mandatory if the task involves HBase or relational databases.
 - b. Conversion type, which is an intermediate conversion step. You can add one or more conversion types or do not add any conversion type.
 - c. Output type. Only one output type can be added in the last conversion step. This parameter is mandatory if the task involves HBase or relational databases.

Table 17-74 Example list

| Type | Description |
|------------|--|
| Input Type | <ul style="list-style-type: none"> ▪ CSV File Input: CSV file input step for configuring separators to generate multiple fields. ▪ Fixed-Width File Input: Text file input step for configuring the length of characters or bytes to be truncated to generate multiple fields. ▪ Table Input: relational data input step for configuring specified columns in the database as input fields. ▪ HBase Input: HBase table input step for configuring the column definition of an HBase table to a specified field. ▪ HTML Input: HTML web page data input step for obtaining the target data of the HTML web page file to the specified field. ▪ Hive Input: Hive table input step for defining columns in a Hive table to specified fields. ▪ Spark Input: Spark SQL table input step for defining columns in the SparkSQL table to specified fields. Only Hive data can be stored and accessed. |

| Type | Description |
|-----------------|--|
| Conversion type | <ul style="list-style-type: none"> ▪ Long Integer Time Conversion: Configure the conversion between a long integer value and a date. ▪ Null Value Conversion: Configure a specified value to replace the null value. ▪ Random Value Conversion: Configure new value-added fields as random data fields. ▪ Adding a Constant Field: Add a constant to directly generate a constant field. ▪ Concatenation and Conversion: Concatenate fields, connect generated fields using connection characters, and convert new fields. ▪ Separator Conversion: Configure the generated fields to be separated by separators and convert new fields. ▪ Modulo Conversion: Configure the generated fields to be converted into new fields through modulo operation. ▪ Cutting Character String: Truncate a generated field based on a specified position to generate a new field. ▪ EL Operation Conversion: Calculate field values. Currently, the following operators are supported: md5sum, sha1sum, sha256sum, and sha512sum. ▪ Character String Case Conversion: Configure the generated fields to be converted to new fields through case conversion. ▪ Reverse String Conversion: Reverse the generated fields to generate new fields. ▪ Character String Space Clearing Conversion: Configure the generated fields to clear spaces and convert them to new fields. ▪ Row Filtering Conversion: Configure logical conditions to filter out rows that contain triggering conditions. ▪ Update Fields: Update the value of a specified field when certain conditions are met. |

| Type | Description |
|-------------|--|
| Output type | <ul style="list-style-type: none"> ▪ File Output: Configure generated fields to be connected by separators and exported to a file. ▪ Table Output: Configure the mapping between output fields and specified columns in the database. ▪ HBase Output: Configure the generated fields to the columns of the HBase table. ▪ Hive Output: Configure generated fields to a column of a Hive table. ▪ Spark Output: Configure generated fields to the columns of SparkSQL tables. Only SparkSQL can access Hive data. |

The edit box allows you to perform the following tasks:

- Re-command: Rename an example.
- Edit: Edit the step conversion by referring to [Step 6.3](#).
- Delete: Delete an example.

 **NOTE**

You can also use the shortcut key Del to delete the file.

3. Click **Edit** to edit the step conversion information and configure fields and data.

For details about how to set parameters in the step conversion information, see [Operator Help](#).

If the conversion step is incorrectly configured, the source data cannot be converted and become dirty data. The dirty data marking rules are as follows:

- In any input type step, the number of fields contained in the original data is less than the number of configured fields or the field values in the original data do not match the configured field type.
- In the **CSV File Input** step, **Validate input field** checks whether the input field matches the value type. If the input field and value type of a line do not match, the line is skipped and becomes dirty data.
- In the **Fixed Width File Input** step, **Fixed Length** specifies the field splitting length. If the length is greater than the length of the original field value, data splitting fails and the current line becomes dirty data.
- In the **HBase Input** step, if the HBase table name specified by **HBase Table Name** is incorrect, or no primary key column is configured for Primary Key, all data becomes dirty data.
- In any conversion step, lines whose conversion fails becomes dirty data. For example, in the **Split Conversion** step, the number of generated fields is less than the number of configured fields, or the original data

cannot be converted to the String type, and the current row becomes dirty data.

- In the **Filter Row Conversion** step, rows filtered by filter criteria become dirty data.
- In the **Modulo Conversion** step, if the original field value is NULL, the current row becomes dirty data.

4. Click **Next**.

Step 7 In the **4. Output Settings** area, set the destination location for saving data and click **Save** to save the job or click **Save and Run** to save and run the job.

Table 17-75 List of Output Configuration Parameters

| Data Connection Type | Parameter | Description |
|----------------------|-------------|---|
| sftp-connector | Output Path | Path or name of the export file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple paths or file names separated with semicolons (;). Ensure that the number of input paths or file names is the same as the number of SFTP servers configured for the connector. |


| Data Connection Type | Parameter | Description |
|----------------------|-------------|--|
| | Operation | <p>Specifies the action during data import. When all data is to be imported from the input path to the destination path, the data is stored in a temporary directory and then copied from the temporary directory to the destination path. After the data is imported successfully, the data is deleted from the temporary directory. One of the following actions can be taken when duplicate file names exist during data transfer:</p> <ul style="list-style-type: none"> • OVERRIDE: overrides the old file. • RENAME: renames as new file. For a file without an extension, a string is added to the file name as the extension; for a file with an extension, a string is added to the extension. The string is unique. • APPEND: adds the content of the new file to the end of the old file. This action only adds content regardless of whether the file can be used. For example, a text file can be used after this operation, while a compressed file cannot. • IGNORE: reserves the old file and does not copy the new file. • ERROR: stops the task and reports an error if duplicate file names exist. Transferred files are imported successfully, while files that have duplicate names and files that are not transferred fail to be imported. |
| | Encode type | Specifies the exported file encoding format, for example, UTF-8. This parameter can be set only in text file export. |
| | Compression | Indicates whether to enable the compressed transmission function when SFTP is used to export data. true indicates that compression is enabled, and false indicates that compression is disabled. |
| hdfs-connector | Output Path | Specifies the output directory or file name of the export file in HDFS or OBS. |

| Data Connection Type | Parameter | Description |
|----------------------------|---------------------------------|--|
| | File Format | Specifies the file export type. <ul style="list-style-type: none"> • TEXT_FILE: imports a text file and stores it as a text file. • SEQUENCE_FILE: imports a text file and stores it as a sequence file. • BINARY_FILE: imports files of any format by using binary streams. |
| | Compression codec | Specifies the compression format of files exported to HDFS or OBS. Select a format from the drop-down list. If you select NONE or do not set this parameter, data is not compressed. |
| | User-defined compression format | Name of a user-defined compression format type. |
| generic-jdbc-connector | Schema name | Specifies the database schema name. |
| | Table name | Specifies the name of a database table that is used to save the final data of the transmission. |
| | Temporary table | Specifies the name of a temporary database table that is used to save temporary data during the transmission. The fields in the table must be the same as those in the database specified by Table name . |
| oracle-partition-connector | Schema Name | Specifies the database schema name. |
| | Table Name | Specifies the name of a database table that is used to save the final data of the transmission. |
| | Temporary Table | Specifies the name of a temporary database table that is used to save temporary data during the transmission. The fields in the table must be the same as those in the database specified by Table name . |
| oracle-connector | Table Name | Destination table name to store data. |
| | Column Name | Specifies the name of the column to be written. Columns that are not specified can be set to null or the default value. |
| mysql-fastpath-connector | Schema Name | Specifies the database schema name. |

| Data Connection Type | Parameter | Description |
|----------------------|----------------------|--|
| | Table Name | Specifies the name of a database table that is used to save the final data of the transmission. |
| | Temporary Table Name | Name of the temporary table, which is used to store data. After the job is successfully executed, data is transferred to the formal table. |

Step 8 On the Loader WebUI page, you can view, start, stop, copy, delete, edit, and view historical information about created jobs.

Figure 17-47 Viewing Loader Jobs



----End

17.13.3 Typical Scenario: Exporting Data from HDFS or OBS to an SFTP Server

Scenario

This section describes how to use Loader to export data from HDFS or OBS to an SFTP server.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HDFS or OBS directories and data involved in job execution.
- You have obtained the username and password of the SFTP server and the user has the write permission of the data export directory on the SFTP server.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- When using Loader to export data from HDFS or OBS, the input paths and input path subdirectories of the HDFS or OBS data source and the name of the files in these directories do not contain any of the following special characters: \ " ; , .

- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

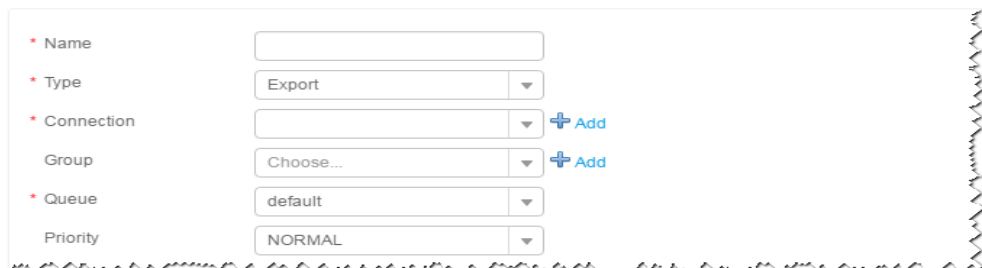
1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-48 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-49 Basic Information



1. Set **Name** to the name of the job.
2. Set **Type** to **Export**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **sftp-connector**, click **Add**, set connection parameters, and click **Test** to verify whether the connection is available. When "**Test Success**" is displayed, click **OK**. Loader

allows multiple SFTP servers to be configured. Click **Add** to add the configuration information of multiple SFTP servers.

Table 17-76 Connection parameters

| Parameter | Description | Example Value |
|------------------|--|---------------|
| Name | Specifies the name of the SFTP server connection. | sftpName |
| SFTP Server IP | Specifies the IP address of the SFTP server. | 10.16.0.1 |
| SFTP Server Port | Specifies the port number of the SFTP server. | 22 |
| SFTP Username | Specifies the user name for accessing the SFTP server. | root |
| SFTP Password | Specifies the password for accessing the SFTP server. | xxxx |
| SFTP Public Key | Specifies public key of the SFTP server. | OdDt/yn...etM |

 **NOTE**

When multiple SFTP servers are configured, the data of HDFS or OBS will be divided into multiple parts and exported to the SFTP servers randomly.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set **Source type** to **HDFS**.

Table 17-77 Data source parameters

| Parameter | Description | Example Value |
|-----------------|--|------------------------|
| Input directory | Specifies the input path when data is exported from HDFS or OBS.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /
use
r/
test |

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Path filter | <p>Specifies the wildcard for filtering the directories in the input paths of the source files. Input directory is not used in filtering. If there are multiple filter conditions, use commas (,) to separate them. If the parameter is empty, the directory is not filtered. The regular expression filtering is not supported.</p> <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. | * |
| File filter | <p>Specifies the wildcard for filtering the file names of the source files. If there are multiple filter conditions, use commas (,) to separate them. The value cannot be left blank. The regular expression filtering is not supported.</p> <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. | * |
| File Type | <p>Specifies the file import type.</p> <ul style="list-style-type: none"> • TEXT_FILE: imports a text file and stores it as a text file. • SEQUENCE_FILE: imports a text file and stores it as a sequence file. • BINARY_FILE: imports files of any format by using binary streams but not to process the files. <p>NOTE
When the file import type to TEXT_FILE or SEQUENCE_FILE, Loader automatically selects a decompression method based on the file name extension to decompress a file.</p> | TEXT_FILE |

| Parameter | Description | Example Value |
|-----------------|---|---------------|
| File Split Type | <p>Indicates whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each map in the MapReduce task for data export.</p> <ul style="list-style-type: none"> ● FILE: indicates that the source file is split by file. That is, each map processes one or multiple complete files, the same source file cannot be allocated to different maps, and the source file directory structure is retained after data import. ● SIZE: indicates that the source file is split by size. That is, each map processes input files of a certain size, and a source file can be divided and processed by multiple maps. After data is stored in the output directory, the number of saved files is the same as the number of maps. The file name format is import_part_xxxx, where xxxx is a unique random number generated by the system. | FILE |
| Extractors | <p>Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. This parameter cannot be set when Extractor Size is set. The value must be less than or equal to 3000. You are advised to set the parameter to the number of CPU cores on the SFTP server.</p> <p>NOTE
To improve the data import speed, ensure that the following conditions are met:</p> <ul style="list-style-type: none"> ● Each map connection is equivalent to a client connection. Therefore, you must ensure that the maximum number of connections of the SFTP server is greater than the number of maps. ● Ensure that the disk I/O or network bandwidth on the SFTP server does not reach the upper limit. | 20 |
| Extractor size | <p>Specifies the size of data processed by maps that are started in a MapReduce job of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000. This parameter cannot be set when Extractors is set.</p> | - |


Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-78](#).

Table 17-78 Setting the input and output parameters of the operator

| Input Type | Export Type |
|------------------------|-------------|
| CSV file input | File output |
| HTML input | File output |
| Fixed-width file input | File output |

Figure 17-50 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set the data storage mode.

Table 17-79 Parameter description

| Parameter | Description | Example Value |
|-------------|---|---------------|
| Output path | <p>Specifies the path or file name of the exported file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple paths or file names separated with semicolons (;). Ensure that the number of paths or file names is the same as the number of SFTP servers configured for the connector.</p> <p>NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items.</p> | /opt/tempfile |
| Operation | <p>Specifies the action during data import. When all data is to be imported from the input path to the destination path, the data is stored in a temporary directory and then copied from the temporary directory to the destination path. After the data is imported successfully, the data is deleted from the temporary directory. One of the following actions can be taken when duplicate file names exist during data transfer:</p> <ul style="list-style-type: none"> ● OVERWRITE: overrides the old file. ● RENAME: renames as new file. For a file without an extension, a string is added to the file name as the extension; for a file with an extension, a string is added to the extension. The string is unique. ● APPEND: adds the content of the new file to the end of the old file. This action only adds content regardless of whether the file can be used. For example, a text file can be used after this operation, while a compressed file cannot. ● IGNORE: reserves the old file and does not copy the new file. ● ERROR: stops the task and reports an error if duplicate file names exist. Transferred files are imported successfully, while files that have duplicate names and files that are not transferred fail to be imported. | OVERWRITE |
| Encode type | Specifies the exported file encoding format, for example, UTF-8. This parameter can be set only in text file export. | UTF-8 |

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Compression | <p>Indicates whether to enable the compressed transmission function when SFTP is used to export data.</p> <ul style="list-style-type: none"> The value true indicates that compression is enabled. The value false indicates that compression is disabled. | true |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the **Loader WebUI**. When **Status** is **Succeeded**, the job is complete.

Figure 17-51 Viewing a job

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|----------------|---------------------------|---------------------|--------------|----------------------------------|-----------|-----------|
| 5 | hdfs-sftp-size | Export from HDFS to SF... | 2016-03-26 14:45:08 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ □ 🔄 ⌵ |
| 4 | hdfs-sftp | Export from HDFS to SF... | 2016-03-26 14:43:14 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ □ 🔄 ⌵ |
| 3 | stfp-hdfs-1-1 | Import from SFTP to HD... | 2016-03-26 12:00:36 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ □ 🔄 ⌵ |
| 1 | stfp-hdfs | Import from SFTP to HD... | 2016-03-26 11:08:09 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ □ 🔄 ⌵ |
| 2 | stfp-hdfs-1 | Import from SFTP to HD... | 2016-03-26 11:07:56 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ □ 🔄 ⌵ |

----End

17.13.4 Typical Scenario: Exporting Data from HBase to an SFTP Server

Scenario

Use Loader to export data from HBase to an SFTP server.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HBase tables or phoenix tables that are used during job execution.
- You have obtained the username and password of the SFTP server and the user has the write permission of the data export directory on the SFTP server.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.

- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

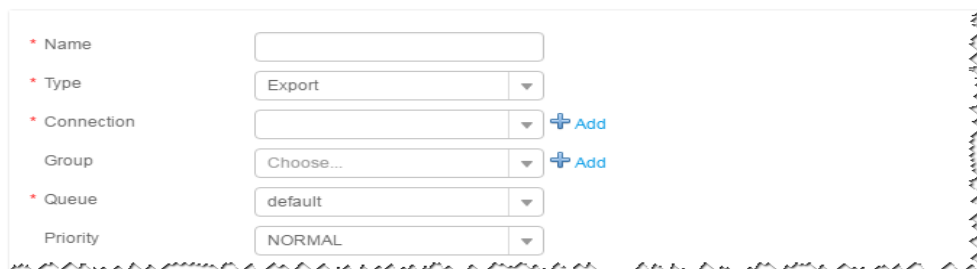
1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-52 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-53 Basic Information



1. Set **Name** to the name of the job.
2. Set **Type** to **Export**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **sftp-connector**, click **Add**, set connection parameters, and click **Test** to verify whether the connection is available. When "**Test Success**" is displayed, click **OK**. Loader allows multiple SFTP servers to be configured. Click **Add** to add the configuration information of multiple SFTP servers.

Table 17-80 Connection parameters

| Parameter | Description | Example Value |
|------------------|--|---------------|
| Name | Specifies the name of the SFTP server connection. | sftpName |
| SFTP server IP | Specifies the IP address of the SFTP server. | 10.16.0.1 |
| SFTP server port | Specifies the port number of the SFTP server. | 22 |
| SFTP username | Specifies the user name for accessing the SFTP server. | root |
| SFTP password | Specifies the password for accessing the SFTP server. | xxxx |
| SFTP public key | Specifies public key of the SFTP server. | OdDt/yn...etM |

 **NOTE**

When multiple SFTP servers are configured, the data of HBase tables or phoenix tables will be divided into multiple parts and saved to the SFTP servers randomly.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set **Source type** to **HBASE**.

Table 17-81 Data source parameters

| Parameter | Description | Example Value |
|----------------|--|---------------|
| HBase instance | Specifies the HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. You are advised to set the parameter to the maximum number of connections on the SFTP server. | 20 |

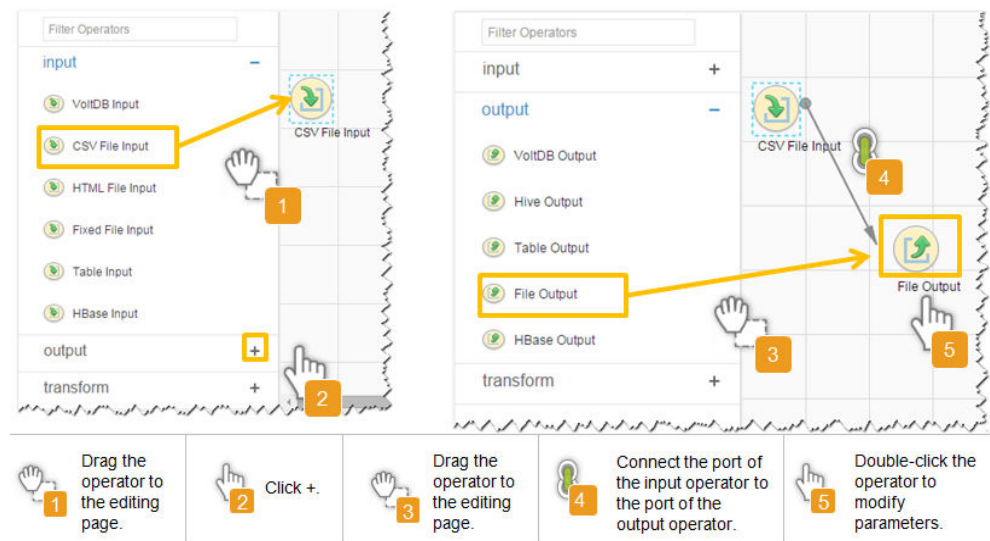
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-82](#).

Table 17-82 Setting the input and output parameters of the operator

| Input Type | Export Type |
|-------------|-------------|
| HBase input | File output |

Figure 17-54 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set the data storage mode.

Table 17-83 Parameter description

| Parameter | Description | Example Value |
|-------------|--|----------------|
| Output path | Specifies the path or file name of the exported file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple paths or file names separated with semicolons (;). Ensure that the number of paths or file names is the same as the number of SFTP servers configured for the connector.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /opt/temppfile |

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Operation | <p>Specifies the action during data import. When all data is to be imported from the input path to the destination path, the data is stored in a temporary directory and then copied from the temporary directory to the destination path. After the data is imported successfully, the data is deleted from the temporary directory. One of the following actions can be taken when duplicate file names exist during data transfer:</p> <ul style="list-style-type: none"> ● OVERRIDE: overrides the old file. ● RENAME: renames as new file. For a file without an extension, a string is added to the file name as the extension; for a file with an extension, a string is added to the extension. The string is unique. ● APPEND: adds the content of the new file to the end of the old file. This action only adds content regardless of whether the file can be used. For example, a text file can be used after this operation, while a compressed file cannot. ● IGNORE: reserves the old file and does not copy the new file. ● ERROR: stops the task and reports an error if duplicate file names exist. Transferred files are imported successfully, while files that have duplicate names and files that are not transferred fail to be imported. | OVERRIDE |
| Encode type | Specifies the exported file encoding format, for example, UTF-8. This parameter can be set only in text file export. | UTF-8 |
| Compression | <p>Indicates whether to enable the compressed transmission function when SFTP is used to export data.</p> <ul style="list-style-type: none"> ● The value true indicates that compression is enabled. ● The value false indicates that compression is disabled. | true |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the **Loader WebUI**. When **Status** is **Succeeded**, the job is complete.

Figure 17-55 Viewing a job

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|----------------|---------------------------|---------------------|--------------|----------------------------------|-----------|--|
| 5 | hdfs-sftp-size | Export from HDFS to SF... | 2016-03-26 14:45:08 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ ⌵ ⌵ |
| 4 | hdfs-sftp | Export from HDFS to SF... | 2016-03-26 14:43:14 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ ⌵ ⌵ |
| 3 | sftp-hdfs-1-1 | Import from SFTP to HD... | 2016-03-26 12:00:36 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ ⌵ ⌵ |
| 1 | sftp-hdfs | Import from SFTP to HD... | 2016-03-26 11:08:09 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ ⌵ ⌵ |
| 2 | sftp-hdfs-1 | Import from SFTP to HD... | 2016-03-26 11:07:56 | admin | <div style="width: 100%;"></div> | Succeeded | ▶ ⌵ ⌵ |

----End

17.13.5 Typical Scenario: Exporting Data from Hive to an SFTP Server

Scenario

Use Loader to export data from Hive to an SFTP server.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the Hive table specified in the job.
- You have obtained the username and password of the SFTP server and the user has the write permission of the data export directory on the SFTP server.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.

Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-56 Loader web UI

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|----------------------------|------|-------------|------------|--------------|----------|--------|-----------|
| No data available in table | | | | | | | |

Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-57 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Export**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **sftp-connector**, click **Add**, set connection parameters, and click **Test** to verify whether the connection is available. When "**Test Success**" is displayed, click **OK**. Loader allows multiple SFTP servers to be configured. Click **Add** to add the configuration information of multiple SFTP servers.

Table 17-84 Connection parameters

| Parameter | Description | Example Value |
|------------------|--|---------------|
| Name | Specifies the name of the SFTP server connection. | sftpName |
| SFTP server IP | Specifies the IP address of the SFTP server. | 10.16.0.1 |
| SFTP server port | Specifies the port number of the SFTP server. | 22 |
| SFTP username | Specifies the user name for accessing the SFTP server. | root |
| SFTP password | Specifies the password for accessing the SFTP server. | xxxx |
| SFTP public key | Specifies public key of the SFTP server. | OdDt/yn...etM |

 **NOTE**

When multiple SFTP servers are configured, the data of Hive tables will be divided into multiple parts and saved to the SFTP servers randomly.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set **Source type** to **HIVE**.

Table 17-85 Data source parameters

| Parameter | Description | Example Value |
|---------------|--|---------------|
| Hive instance | Specifies the Hive service instance that Loader selects from all available Hive service instances in the cluster. If the selected Hive service instance is not added to the cluster, the Hive job cannot run properly. | hive |
| Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. You are advised to set the parameter to the maximum number of connections on the SFTP server. | 20 |

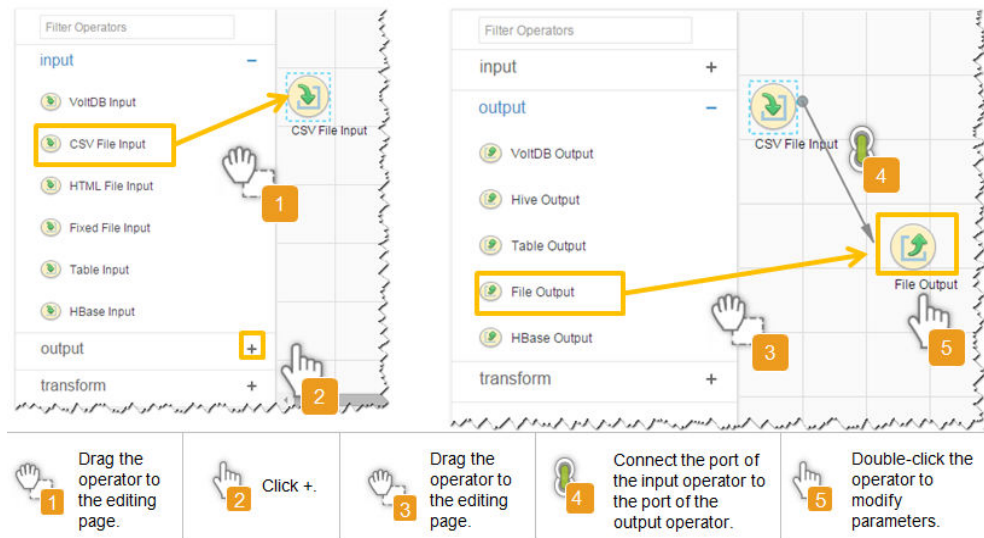
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-86](#).

Table 17-86 Setting the input and output parameters of the operator

| Input Type | Export Type |
|------------|-------------|
| Hive input | File output |

Figure 17-58 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set the data storage mode.

Table 17-87 Parameter description

| Parameter | Description | Example Value |
|-------------|--|-------------------|
| Output path | Specifies the path or file name of the exported file on an SFTP server. If multiple SFTP server IP addresses are configured for the connector, you can set this parameter to multiple paths or file names separated with semicolons (;). Ensure that the number of paths or file names is the same as the number of SFTP servers configured for the connector.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /opt/tem
pfile |

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Operation | <p>Specifies the action during data import. When all data is to be imported from the input path to the destination path, the data is stored in a temporary directory and then copied from the temporary directory to the destination path. After the data is imported successfully, the data is deleted from the temporary directory. One of the following actions can be taken when duplicate file names exist during data transfer:</p> <ul style="list-style-type: none"> • OVERRIDE: overrides the old file. • RENAME: renames as new file. For a file without an extension, a string is added to the file name as the extension; for a file with an extension, a string is added to the extension. The string is unique. • APPEND: adds the content of the new file to the end of the old file. This action only adds content regardless of whether the file can be used. For example, a text file can be used after this operation, while a compressed file cannot. • IGNORE: reserves the old file and does not copy the new file. • ERROR: stops the task and reports an error if duplicate file names exist. Transferred files are imported successfully, while files that have duplicate names and files that are not transferred fail to be imported. | OVERRIDE |
| Encode type | Specifies the exported file encoding format, for example, UTF-8. This parameter can be set only in text file export. | UTF-8 |
| Compression | <p>Indicates whether to enable the compressed transmission function when SFTP is used to export data.</p> <ul style="list-style-type: none"> • The value true indicates that compression is enabled. • The value false indicates that compression is disabled. | true |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the **Loader WebUI**. When **Status** is **Succeeded**, the job is complete.

Figure 17-59 Viewing a job

| Job ID | Name | Description | Start Time | Execute User | Progress | Status | Operation |
|--------|----------------|---------------------------|---------------------|--------------|---|-----------|----------------|
| 5 | hdfs-sftp-size | Export from HDFS to SF... | 2016-03-26 14:45:08 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | [Action icons] |
| 4 | hdfs-sftp | Export from HDFS to SF... | 2016-03-26 14:43:14 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | [Action icons] |
| 3 | sftp-hdfs-1-1 | Import from SFTP to HD... | 2016-03-26 12:00:36 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | [Action icons] |
| 1 | sftp-hdfs | Import from SFTP to HD... | 2016-03-26 11:08:09 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | [Action icons] |
| 2 | sftp-hdfs-1 | Import from SFTP to HD... | 2016-03-26 11:07:56 | admin | <div style="width: 100%; background-color: green;"></div> | Succeeded | [Action icons] |

----End

17.13.6 Typical Scenario: Exporting Data from HDFS or OBS to a Relational Database

Scenario

This section describes how to use Loader to export data from HDFS or OBS to a relational database.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HDFS or OBS directories and data involved in job execution.
- You have obtained the username and password of the relational database.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.
- Before the operation, perform the following steps:
 - a. Obtain the JAR package of the relational database driver and save it to the following directory on the active and standby Loader nodes: `$ {BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`.
 - b. Run the following command on the active and standby nodes as user root to modify the permission:


```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/
FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/
webapps/loader/WEB-INF/ext-lib
chown omm:wheel JAR package name
chmod 600 JAR package name
```
 - c. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Service** > **Loader** > **More** > **Restart**. Enter the password of the administrator to restart the Loader service.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

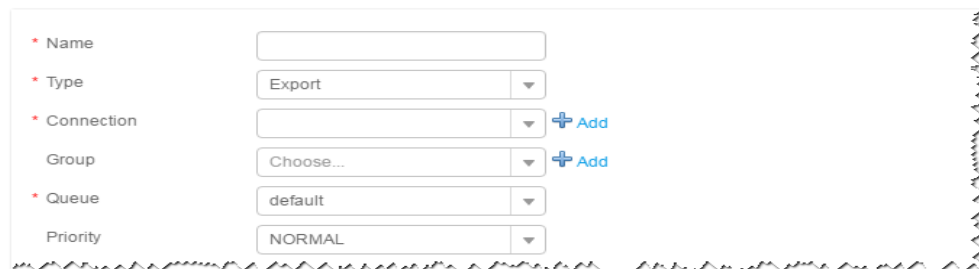
1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-60 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-61 Basic Information



1. Set **Name** to the name of the job.
2. Set **Type** to **Export**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **generic-jdbc-connector** or dedicated database connector (oracle-connector, oracle-partition-connector or mysql-fastpath-connector), set connection parameters, and click **Test** to verify whether the connection is available. When "**Test Success**" is displayed, click **OK**.

 NOTE

- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When **mysql-fastpath-connector** is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManagers, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, see <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>. Install the MySQL client applications and tools.

Table 17-88 generic-jdbc-connector connection parameters

| Parameter | Description | Example Value |
|----------------------------|---|--|
| Name | Specifies the name of a relational database connection. | dbName |
| JDBC Driver Class | Specifies the name of a Java database connectivity (JDBC) driver class. | oracle.jdbc.driver.OracleDriver |
| JDBC Connection String | Specifies the JDBC connection string. | jdbc:oracle:thin:@//10.16.0.1:1521/oradb |
| Username | Specifies the username for connecting to the database. | omm |
| Password | Specifies the password for connecting to the database. | xxxx |
| JDBC Connection Properties | JDBC connection attribute. Click Add to manually add the attribute. <ul style="list-style-type: none"> • Name: connection attribute name • Value: connection attribute value | <ul style="list-style-type: none"> • Name: socketTimeout • Value: 20 |

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set **Source type** to **HDFS**.

Table 17-89 Data source parameters

| Parameter | Description | Example Value |
|-----------------|---|--------------------|
| Input directory | Specifies the input path when data is exported from HDFS or OBS.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /
user/
test |
| Path filter | Specifies the wildcard for filtering the directories in the input paths of the source files. Input directory is not used in filtering. If there are multiple filter conditions, use commas (,) to separate them. If the parameter is empty, the directory is not filtered. The regular expression filtering is not supported. <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. | * |
| File filter | Specifies the wildcard for filtering the file names of the source files. If there are multiple filter conditions, use commas (,) to separate them. The value cannot be left blank. The regular expression filtering is not supported. <ul style="list-style-type: none"> • ? matches a single character. • * indicates multiple characters. • Adding ^ before the condition indicates negated filtering, that is, file filtering. | * |
| File Type | Specifies the file import type. <ul style="list-style-type: none"> • TEXT_FILE: imports a text file and stores it as a text file. • SEQUENCE_FILE: imports a text file and stores it as a sequence file. • BINARY_FILE: imports files of any format by using binary streams but not to process the files. NOTE
When the file import type to TEXT_FILE or SEQUENCE_FILE , Loader automatically selects a decompression method based on the file name extension to decompress a file. | TEXT_FILE |

| Parameter | Description | Example Value |
|-----------------|---|---------------|
| File split type | <p>Indicates whether to split source files by file name or size. The files obtained after the splitting are used as the input files of each map in the MapReduce task for data export.</p> <ul style="list-style-type: none"> ● FILE: indicates that the source file is split by file. That is, each map processes one or multiple complete files, the same source file cannot be allocated to different maps, and the source file directory structure is retained after data import. ● SIZE: indicates that the source file is split by size. That is, each map processes input files of a certain size, and a source file can be divided and processed by multiple maps. After data is stored in the output directory, the number of saved files is the same as the number of maps. The file name format is import_part_xxxx, where xxxx is a unique random number generated by the system. | FILE |
| Extractors | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. This parameter cannot be set when Extractor size is set. The value must be less than or equal to 3000. | 20 |
| Extractor size | Specifies the size of data processed by maps that are started in a MapReduce job of a data configuration operation. The unit is MB. The value must be greater than or equal to 100. The recommended value is 1000. This parameter cannot be set when Extractors is set. When a relational database connector is used, Extractor size is unavailable. You need to set Extractors . | - |

Setting Data Transformation


Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-90](#).

Table 17-90 Setting the input and output parameters of the operator

| Input Type | Export Type |
|----------------|--------------|
| CSV file input | Table output |

| Input Type | Export Type |
|------------------------|--------------|
| HTML Input | Table output |
| Fixed-width file input | Table output |

Figure 17-62 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set the data storage mode.

Table 17-91 Parameter description

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Schema name | Specifies the database schema name. | dbo |
| Table Name | Specifies the name of a database table that is used to save the final data of the transmission.
NOTE
Table names can be defined using macros. For details, see Using Macro Definitions in Configuration Items . | test |

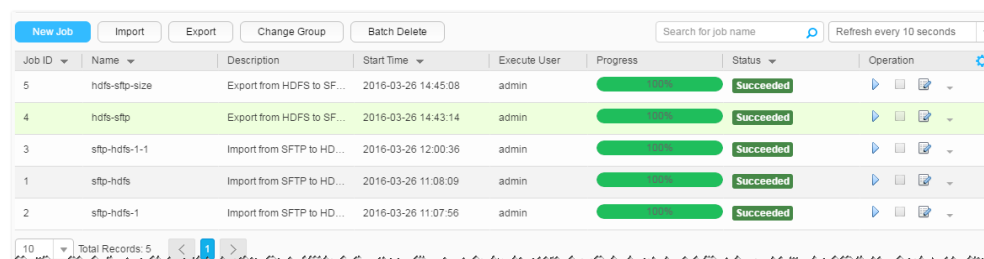
| Parameter | Description | Example Value |
|-----------------|--|---------------|
| Temporary table | <p>Specifies the name of a temporary database table that is used to save temporary data during the transmission. The fields in the table must be the same as those in the database specified by Table name.</p> <p>NOTE
A temporary table is used to prevent dirty data from being generated in the destination table when data is exported to the database. Data is migrated from the temporary table to the destination table only after all data is successfully written to the temporary table. Using temporary tables increases the job execution time.</p> | tmp_test |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the **Loader WebUI**. When **Status** is **Succeeded**, the job is complete.

Figure 17-63 Viewing a job



----End

17.13.7 Typical Scenario: Exporting Data from HBase to a Relational Database

Scenario

Use Loader to export data from HBase to a relational database.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the HBase tables or phoenix tables that are used during job execution.
- You have obtained the username and password of the relational database.

- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.
- Before the operation, perform the following steps:
 - a. Obtain the JAR package of the relational database driver and save it to the following directory on the active and standby Loader nodes: `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`.
 - b. Run the following command on the active and standby nodes as user root to modify the permission:


```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/
FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/
webapps/loader/WEB-INF/ext-lib
chown omm:wheel JAR package name
chmod 600 JAR package name
```
 - c. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Service** > **Loader** > **More** > **Restart**. Enter the password of the administrator to restart the Loader service.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-64 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-65 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Export**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **generic-jdbc-connector** or dedicated database connector (oracle-connector, oracle-partition-connector or mysql-fastpath-connector), set connection parameters, and click **Test** to verify whether the connection is available. When "**Test Success**" is displayed, click **OK**.

NOTE

- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When **mysql-fastpath-connector** is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManagers, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, see <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>. Install the MySQL client applications and tools.

Table 17-92 generic-jdbc-connector connection parameters

| Parameter | Description | Example Value |
|-------------------|---|---------------------------------|
| Name | Specifies the name of a relational database connection. | dbName |
| JDBC Driver Class | Specifies the name of a Java database connectivity (JDBC) driver class. | oracle.jdbc.driver.OracleDriver |

| Parameter | Description | Example Value |
|----------------------------|---|--|
| JDBC Connection String | Specifies the JDBC connection string. | jdbc:oracle:thin:@//10.16.0.1:1521/oradb |
| Username | Specifies the username for connecting to the database. | omm |
| Password | Specifies the password for connecting to the database. | xxxx |
| JDBC Connection Properties | JDBC connection attribute. Click Add to manually add the attribute. <ul style="list-style-type: none"> Name: connection attribute name Value: connection attribute value | <ul style="list-style-type: none"> Name: socketTimeout Value: 20 |

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set **Source type** to **HBASE**.

Table 17-93 Data source parameters

| Parameter | Description | Example Value |
|----------------|---|---------------|
| HBase instance | Specifies the HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. | 20 |

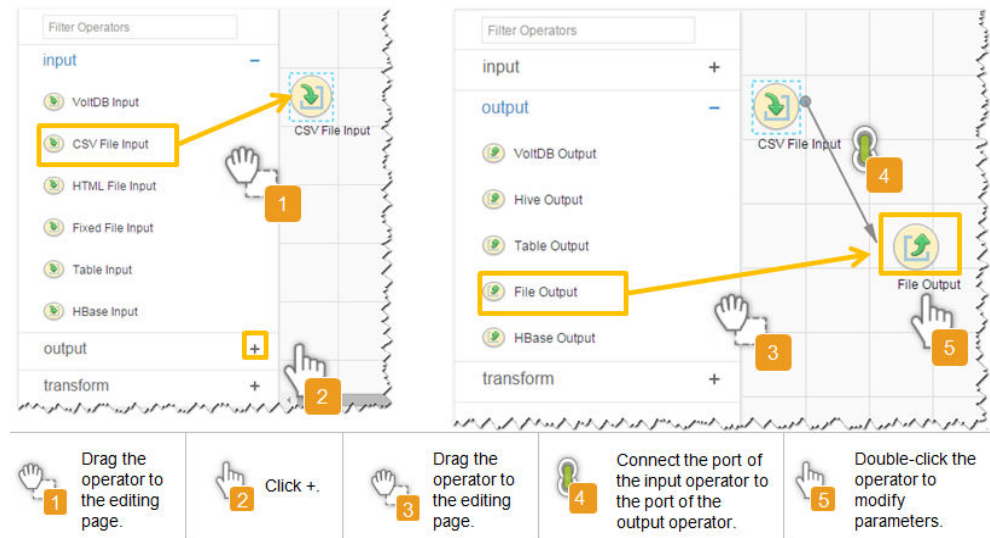
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-94](#).

Table 17-94 Setting the input and output parameters of the operator

| Input Type | Export Type |
|-------------|--------------|
| HBase input | Table output |

Figure 17-66 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set the data storage mode.

Table 17-95 Parameter description

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Schema name | Specifies the database schema name. | dbo |
| Table name | Specifies the name of a database table that is used to save the final data of the transmission.
NOTE
Table names can be defined using macros. For details, see Using Macro Definitions in Configuration Items . | test |


| Parameter | Description | Example Value |
|-----------------|--|---------------|
| Temporary table | <p>Specifies the name of a temporary database table that is used to save temporary data during the transmission. The fields in the table must be the same as those in the database specified by Table name.</p> <p>NOTE
A temporary table is used to prevent dirty data from being generated in the destination table when data is exported to the database. Data is migrated from the temporary table to the destination table only after all data is successfully written to the temporary table. Using temporary tables increases the job execution time.</p> | tmp_test |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the **Loader WebUI**. When **Status** is **Succeeded**, the job is complete.

Figure 17-67 Viewing a job



----End

17.13.8 Typical Scenario: Exporting Data from Hive to a Relational Database

Scenario

Use Loader to export data from Hive to a relational database.

Prerequisites

- You have obtained the service username and password for creating a Loader job.
- You have had the permission to access the Hive tables that are used during job execution.
- You have obtained the username and password of the relational database.

- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.
- Before the operation, perform the following steps:
 - a. Obtain the JAR package of the relational database driver and save it to the following directory on the active and standby Loader nodes: `${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/webapps/loader/WEB-INF/ext-lib`.
 - b. Run the following command on the active and standby nodes as user root to modify the permission:


```
cd ${BIGDATA_HOME}/FusionInsight_Porter_8.1.0.1/install/
FusionInsight-Sqoop-1.99.3/FusionInsight-Sqoop-1.99.3/server/
webapps/loader/WEB-INF/ext-lib
chown omm:wheel JAR package name
chmod 600 JAR package name
```
 - c. Log in to FusionInsight Manager. Choose **Cluster** > *Name of the desired cluster* > **Service** > **Loader** > **More** > **Restart**. Enter the password of the administrator to restart the Loader service.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-68 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-69 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Export**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **generic-jdbc-connector** or dedicated database connector (oracle-connector, oracle-partition-connector or mysql-fastpath-connector), set connection parameters, and click **Test** to verify whether the connection is available. When "**Test Success**" is displayed, click **OK**.

NOTE

- For connection to relational databases, general database connectors (generic-jdbc-connector) or dedicated database connectors (oracle-connector, oracle-partition-connector, and mysql-fastpath-connector) are available. However, compared with general database connectors, dedicated database connectors perform better in data import and export because they are optimized for specific database types.
- When **mysql-fastpath-connector** is used, the **mysqldump** and **mysqlimport** commands of MySQL must be available on NodeManagers, and the MySQL client version to which the two commands belong must be compatible with the MySQL server version. If the two commands are unavailable or the versions are incompatible, see <http://dev.mysql.com/doc/refman/5.7/en/linux-installation-rpm.html>. Install the MySQL client applications and tools.

Table 17-96 generic-jdbc-connector connection parameters

| Parameter | Description | Example Value |
|-------------------|---|---------------------------------|
| Name | Specifies the name of a relational database connection. | dbName |
| JDBC Driver Class | Specifies the name of a Java database connectivity (JDBC) driver class. | oracle.jdbc.driver.OracleDriver |

| Parameter | Description | Example Value |
|----------------------------|---|--|
| JDBC Connection String | Specifies the JDBC connection string. | jdbc:oracle:thin:@//10.16.0.1:1521/oradb |
| Username | Specifies the username for connecting to the database. | omm |
| Password | Specifies the password for connecting to the database. | xxxx |
| JDBC Connection Properties | JDBC connection attribute. Click Add to manually add the attribute. <ul style="list-style-type: none"> Name: connection attribute name Value: connection attribute value | <ul style="list-style-type: none"> Name: socketTimeout Value: 20 |

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set **Source type** to **HIVE**.

Table 17-97 Data source parameters

| Parameter | Description | Example Value |
|---------------|--|---------------|
| Hive instance | Specifies the Hive service instance that Loader selects from all available Hive service instances in the cluster. If the selected Hive service instance is not added to the cluster, the Hive job cannot run properly. | hive |
| Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. | 20 |

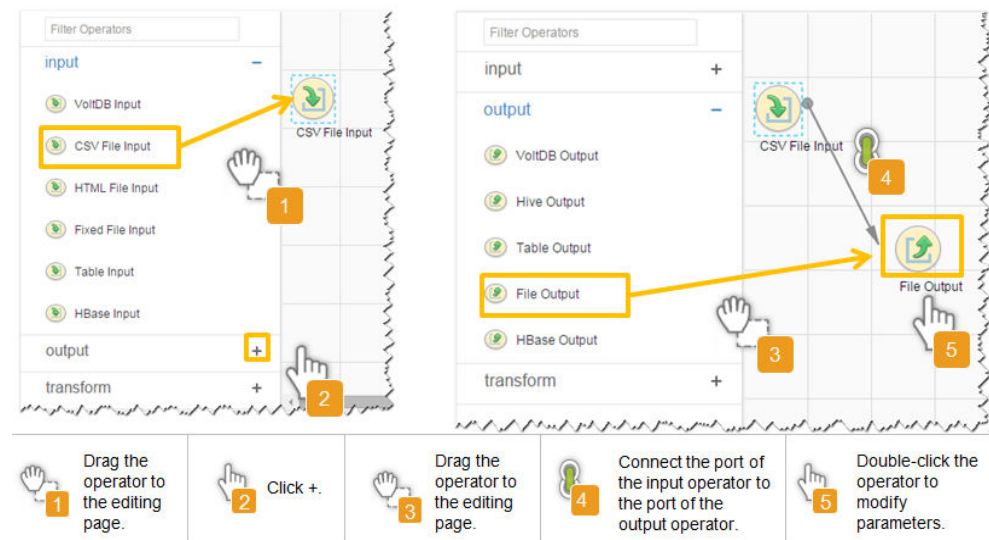
Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-98](#).

Table 17-98 Setting the input and output parameters of the operator

| Input Type | Export Type |
|------------|--------------|
| Hive input | Table output |

Figure 17-70 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set the data storage mode.

Table 17-99 Parameter description

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Schema name | Specifies the database schema name. | dbo |
| Table name | Specifies the name of a database table that is used to save the final data of the transmission.
NOTE
Table names can be defined using macros. For details, see Using Macro Definitions in Configuration Items . | test |

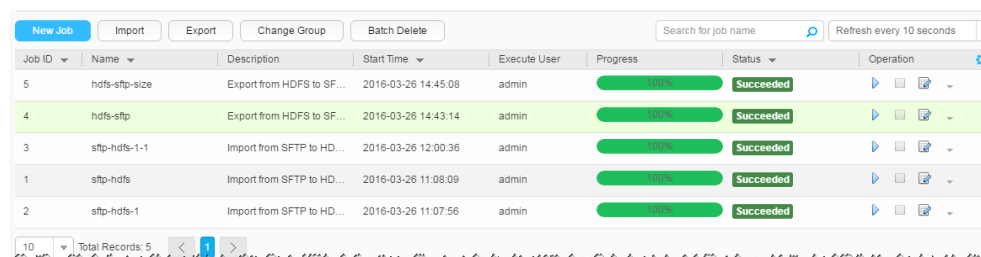
| Parameter | Description | Example Value |
|-----------------|--|---------------|
| Temporary table | <p>Specifies the name of a temporary database table that is used to save temporary data during the transmission. The fields in the table must be the same as those in the database specified by Table name.</p> <p>NOTE
A temporary table is used to prevent dirty data from being generated in the destination table when data is exported to the database. Data is migrated from the temporary table to the destination table only after all data is successfully written to the temporary table. Using temporary tables increases the job execution time.</p> | tmp_test |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the **Loader WebUI**. When **Status** is **Succeeded**, the job is complete.

Figure 17-71 Viewing a job



----End

17.13.9 Typical Scenario: Importing Data from HBase to HDFS or OBS

Scenario

This section describes how to use Loader to export data from HBase to HDFS or OBS.

Prerequisites

- You have obtained the service user name and password for creating a Loader job.
- You have had the permission to access the HDFS or OBS directories and data involved in job execution.

- You have had the permission to access the HBase tables or phoenix tables that are used during job execution.
- No disk space alarm is reported, and the available disk space is sufficient for importing and exporting data.
- If a configured task requires the Yarn queue function, the user must be authorized with related Yarn queue permission.
- The user who configures a task must obtain execution permission on the task and obtain usage permission on the related connection of the task.


Procedure

Setting Basic Job Information

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-72 Loader web UI



Step 2 Click **New Job** to go to the **Basic Information** page and set basic job information.

Figure 17-73 Basic Information

1. Set **Name** to the name of the job.
2. Set **Type** to **Export**.
3. Set **Group** to the group to which the job belongs. No group is created by default. You need to click **Add** to create a group and click **OK** to save the created group.
4. Set **Queue** to the Yarn queue that executes the job. The default value is **root.default**.
5. Set **Priority** to the priority of the Yarn queue that executes the job. The default value is **NORMAL**. The options are **VERY_LOW**, **LOW**, **NORMAL**, **HIGH**, and **VERY_HIGH**.

Step 3 In the **Connection** area, click **Add** to create a connection, set **Connector** to **hdfs-connector**, set connection parameters, and click **Test** to verify whether the connection is available. When "Test Success" is displayed, click **OK**.

Setting Data Source Information

Step 4 Click **Next**. On the displayed **From** page, set **Source type** to **HBASE**.

Table 17-100 Parameter description

| Parameter | Description | Example |
|----------------|---|---------|
| HBase instance | Specifies the HBase service instance that Loader selects from all available HBase service instances in the cluster. If the selected HBase service instance is not added to the cluster, the HBase job cannot be run properly. | HBase |
| Quantity | Specifies the number of maps that are started at the same time in a MapReduce job of a data configuration operation. The value must be less than or equal to 3000. | 20 |


Setting Data Transformation

Step 5 Click **Next**. On the displayed **Transform** page, set the transformation operations in the data transformation process. For details about how to select operators and set parameters, see [Operator Help](#) and [Table 17-101](#).

Table 17-101 Setting the input and output parameters of the operator

| Input Type | Export Type |
|-------------|-------------|
| HBase input | File output |

Figure 17-74 Operator operation procedure



Setting Data Storage Information and Executing the Job

Step 6 Click **Next**. On the displayed **To** page, set the data storage mode.

Table 17-102 Parameter description

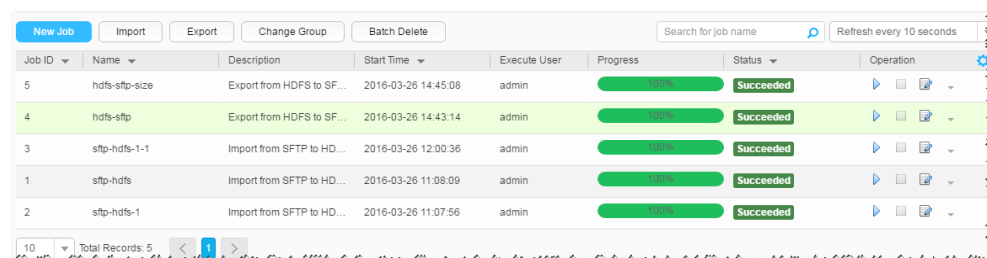
| Parameter | Description | Example |
|-------------------|---|------------|
| Output path | Specifies the output directory or file name of the export file in the HDFS or OBS.
NOTE
You can use macros to define path parameters. For details, see Using Macro Definitions in Configuration Items . | /user/test |
| File Format | Specifies the file export type. <ul style="list-style-type: none"> TEXT_FILE: imports a text file and stores it as a text file. SEQUENCE_FILE: imports a text file and stores it as a sequence file. BINARY_FILE: imports files of any format by using binary streams. | TEXT_FILE |
| Compression codec | Specifies the compression format of files exported to HDFS or OBS. Select a format from the drop-down list. If you select NONE or do not set this parameter, data is not compressed. | NONE |

Step 7 Click **Save and run** to save and run the job.

Checking the Job Execution Result

Step 8 Go to the **Loader WebUI**. When **Status** is **Succeeded**, the job is complete.

Figure 17-75 Viewing a job



----End

17.14 Managing Jobs

17.14.1 Migrating Loader Jobs in Batches

Scenario

Loader allows jobs to be migrated in batches from a group (source group) to another group (target group).

This section applies to MRS 3.x or later.

Prerequisites


- The source group and target group exist.
- The current user has the **Group Edit** permission for the source group and target group.
- The current user has the **Jobs Edit** permission for the source group or the **Edit** permission for the jobs to be migrated.

Procedure

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-76 Loader web UI



Step 2 Click **Change Group**. The Job Migration page is displayed.

Step 3 In **Source group**, select the group to which the jobs to be migrated belong; in **target group**, select the group to which the jobs are to be migrated.

Step 4 Set **Select Change Type** to a migration type.

- **All**: migrates all the jobs in the source group to the target group.
- **Specify Job**: migrates the specified jobs in the source group to the target group. Select **Specify Job**. In the job list, select the jobs to be migrated.

Step 5 Click **OK** to start job migration. In the displayed dialog box, if the progress bar is 100%, the job migration is complete.

----End

17.14.2 Deleting Loader Jobs in Batches

Scenario

Loader allows existing jobs to be deleted in batches.

This section applies to MRS 3.x or later.

Prerequisites


The current user has the **Edit** permission for the jobs to be deleted or the **Jobs Edit** permission for the group to which the jobs belong.

Procedure

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-77 Loader web UI



Step 2 Click **Batch Delete**. The Batch Delete page is displayed.

Step 3 Set **Batch Delete** to a job deletion type.

- **ALL**: deletes all jobs.
- **Specify Job**: deletes specified jobs. Select **Specify Job**. In the job list, select the jobs to be deleted.

Step 4 Click **OK** to start the job deletion. In the displayed dialog box, if the progress bar is 100%, the job deletion is complete.

----End

17.14.3 Importing Loader Jobs in Batches

Scenario

Loader allows all jobs of a configuration file to be imported in batches.

This section applies to MRS 3.x or later.

Prerequisites

The current user has the **Jobs Edit** permission of the group to which the jobs to be imported belong.

NOTE


If the group to which the jobs to be imported belong does not exist, the group is automatically created first. The current user is the creator of the group and has the **Jobs Edit** permission of the group.

Procedure

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-78 Loader web UI



Step 2 Click **Import**. The Export Job page is displayed.

Step 3 On the **Import** page, specify the path of the configuration file whose jobs are to be imported.

Step 4 Click **Upload** to start the job import. In the displayed dialog box, if the progress bar is 100%, the job import is complete.

----End

17.14.4 Exporting Loader Jobs in Batches

Scenario

Loader allows existing jobs to be exported in batches.

This section applies to MRS 3.x or later.

Prerequisites


The current user has the **Edit** permission for the jobs to be exported or the **Jobs Edit** permission of the group to which the jobs belong.

Procedure

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-79 Loader web UI



Step 2 Click **Export**. The job export page is displayed.

Step 3 Set **Batch Delete** to a job export type.

- **ALL**: exports all jobs.
- **Specify Job**: exports specified jobs. Select **Specify Job**. In the job list, select the jobs to be exported.
- **Specify Group**: exports all the jobs in a specified group. Select **Specify Group**. In the group list, select the group whose jobs are to be exported.

Export Password: exports the connector password. If this parameter is selected, the password is exported as an encrypted string.

Step 4 Click **OK** to start the job export. In the displayed dialog box, if the progress bar is 100%, the job import is complete.

----End

17.14.5 Viewing Historical Job Information

Scenario

Query the execution status and execution duration of a Loader job during routine maintenance. You can perform the following operations on the job:

- **Dirty Data**: Query data that fails to be processed or data that is filtered out during job execution, and check which source data does not meet transformation or cleaning rules.
- **Logs**: Query log information about job execution in MapReduce.

This section applies to MRS 3.x or later.

Prerequisites


You have obtained the username and password for logging in to the Loader WebUI.

Procedure

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-80 Loader web UI



Step 2 Query historical records of a Loader job.

1. Locate the row that contains the job to be viewed.
2. Click **More** and select **View History** to view the job execution history.

Figure 17-81 Viewing historical records




Table 17-103 Parameters

| Name | Description |
|--------------------|--|
| Rows/Files Read | Indicates the number of rows (files) read from the input source. |
| Rows/Files Written | Number of rows (files) written to the output source. |
| Rows/Files Skipped | <ul style="list-style-type: none"> – Indicates the number of bad rows (files) recorded during transformation. The input format is incorrect, so transformation cannot be performed. – Number of rows that are skipped after filtering conditions are configured during conversion. |

----End

17.15 Operator Help

17.15.1 Overview

The "Operator Help" section applies to MRS 3.x and later versions.

Conversion Process

Loader reads data at the source end, uses an input operator to convert data into fields by certain rules, use a conversion operator to clean or convert the fields, and finally use an output operator to process the fields and export the output result to the target end.

- A job for performing data conversion can have only one input operator and one output operator.
- Data that does not meet conversion rules will become dirty data and be skipped.

NOTE

- When importing data from a relational database to HDFS or OBS, you do not need to configure data conversion. Data is separated by commas (,) and saved to HDFS or OBS.
- When exporting data from HDFS or OBS to a relational database, you do not need to configure data conversion. Data is separated by commas (,) and saved to the relational database.

Operator Description

Loader operators have three types:

- **Input Operators**
First step of data conversion. This type of operator converts data into fields. Only one input operator can be used in each conversion. The input operator is mandatory in HBase or Hive data import and export.
- **Conversion Operators**
Intermediate conversion step of data conversion. This type of operator is optional. The conversion operators can be used together in any combination. Conversion operators can process only fields. Therefore, an input operator must be used first to convert data into fields.
- **Output Operators**
Last step of data conversion. Only one output operator can be used in each conversion for exporting processed fields. The output operator is mandatory in HBase or Hive data import and export.

Table 17-104 List of operator types

| Node Type | Description |
|-----------|---|
| Input: | <ul style="list-style-type: none">• CSV file input: Each line in the file is converted into multiple input fields based on the specified delimiter.• Fixed-width file input: Each line of the file is converted into multiple input fields based on the characters or bytes with configurable length.• Table input: converts specified columns in a relational database table into input fields of the same quantity.• HBase input: converts specified columns in an HBase table into input fields of the same quantity.• HTML input: converts elements in an HTML file into input fields.• Hive input: converts specified columns in a Hive table into input fields of the same quantity. |

| Node Type | Description |
|-----------|--|
| Convert | <ul style="list-style-type: none"> ● Long integer to time conversion: Implements the conversion between long integer values and date types. ● Null value conversion: Replaces a null value with a specified value. ● Constant field adding: Generates a constant field. ● Random value conversion: generates a random number field. ● Concatenation and conversion: concatenates existing fields to generate new fields. ● Delimiter conversion: separates existing fields with specified separators to generate new fields. ● Modulo conversion: performs modulo operation on an existing field to generate a new field. ● Character string cutting: cuts existing string fields by the specified start position and end position to generate new fields. ● EL operation conversion: specifies a calculator to calculate field values. Currently, the following operators are supported: md5sum, sha1sum, sha256sum, and sha512sum. ● Character string case conversion: converts the upper and lower cases of existing fields to generate new fields. ● Character string reverse conversion: reverses existing character string fields to generate new fields. ● Character string space clearing and conversion: clears the spaces on the left and right of the existing character string fields to generate new fields. ● Row filtering conversion: filters rows that contain triggering conditions by configuring logic conditions. ● Domain update: updates fields values when certain conditions are met. |
| Output: | <ul style="list-style-type: none"> ● Hive Output: exports existing fields to a Hive table. ● Table output: exports existing fields to a relational database table. ● File output: uses delimiters to concatenate existing fields and exports new fields to a file. ● HBase Output: exports existing fields to an HBase table. |

Field Description

Fields in the job configuration are data items defined by Loader based on service requirements to match user data. Fields have specific types and the fields types must be consistent with the actual user data types.

17.15.2 Input Operators

17.15.2.1 CSV File Input

Overview

The **CSV File Input** operator imports all files that can be opened by using a text editor.

Input and Output

- Input: test files
- Output: fields

Parameter Description

Table 17-105 Operator parameter description

| Parameter | Description | Type | Mandatory | Default Value |
|-------------------|--|---------|-----------|---------------|
| Delimiter | Delimiter in a CSV file for separating data lines. | string | Yes | , |
| Line Delimiter | Line delimiter, which can be any string specified by users based on the actual situation. The OS line delimiter is used by default. | string | No | \n |
| Filename as field | User-defined field whose value is the name of the file that stores the current data. | string | No | None |
| Absolute path | Indicates whether the file name used as the value of Filename as field contains an absolute path. Selecting the option button indicates that the file name contains an absolute path; deselecting the option button indicates that the file name does not contain a path. | boolean | No | Deselect |

| Parameter | Description | Type | Mandatory | Default Value |
|----------------------|--|------|-----------|---------------|
| Validate input field | Checks whether the input field matches the value type. If the value is NO , no check is performed. If the value is YES , whether the input field matches the value type is checked. If the input fields do not match the value type, the line is skipped. | enum | Yes | YES |
| Input fields | <p>Information about input fields:</p> <ul style="list-style-type: none"> • position: Position of the field after data lines in the source file are separated by delimiters. The position sequence starts from 1. • field name: Field name. • type: Field type. • date format: If the field type is DATE, TIME, or TIMESTAMP, you must specify a time format. If the field type is set to other values, the time format is invalid. An example time format is yyyyMMdd HH:mm:ss. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. | map | Yes | None |

Data Processing Rule

- Each data line is separated into multiple fields by using delimiters and the fields are used by the subsequent conversion operator.
- If the field value does not match the actual type, the data in the line will become dirty data.

- If the number of input field columns is equal to the number of field columns actually included in the original data, the data in the line will become dirty data.

Example

The following figure shows the source file.

```
2016,year  
year,2016
```

Configure the **CSV File Input** operator, set **Delimiter** to a comma (,), and generate fields A and B.

Delimiter: ,

Line Delimiter:

Filename as field:

Absolute path:

Validate input field: YES

Input fields

Import Export

Table Edit Text Area Edit

| position | field name | type | date format | length | |
|----------|------------|---------|-------------|--------|---------|
| 1 | A | VARCHAR | | | ↑ ↓ ↺ ✕ |
| 2 | B | VARCHAR | | | ↑ ↓ ↺ ✕ |

Add

Fields A and B are generated, as shown in the following figure.

```
2016,year  
year,2016
```

17.15.2.2 Fixed File Input

Overview

The **Fixed File Input** operator converts each line in a file into multiple fields by character or byte of a configurable length.

Input and Output

- Input: text file
- Output: fields

Parameter Description

Table 17-106 Operator parameter description

| Parameter | Description | Type | Mandatory | Default Value |
|-------------------|--|--------|-----------|---------------|
| Line Delimiter | Line delimiter, which can be any string specified by users based on the actual situation. The OS line delimiter is used by default. | string | No | \n |
| Fixed length unit | Length unit. The options are char and byte . | enum | Yes | char |
| Input fields | <p>Information about input fields:</p> <ul style="list-style-type: none"> • fixed length: Field length. The ending of the first field is the starting of the second field, the ending of the second field is the starting of the third field, and so on. • field name: Names of input fields. • type: Field type. • date format: If the field type is DATE, TIME, or TIMESTAMP, you must specify a time format. If the field type is set to other values, the time format is invalid. An example time format is yyyyMMdd HH:mm:ss. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. | map | Yes | None |

Data Processing Rule

- The source file is split based on the input field length to generate fields.
- If the field value does not match the actual type, the data in the line will become dirty data.
- If the field split length is greater than the length of the original field value, the data split fails and the line becomes dirty data.

Example

The following figure shows the source file.

```
fusionInsightbigdataprodu
```

Configure the **Fixed File Input** operator to generate fields A, B, and C.

| fixed length | filed name | type | date format | length |
|--------------|------------|---------|-------------|--------|
| 13 | A | VARCHAR | | |
| 7 | B | VARCHAR | | |
| 7 | C | VARCHAR | | |

The three fields are generated, as shown in the following figure.

```
fusionInsight,bigdata,product
```

17.15.2.3 Table Input

Overview

Table Input operator converts specified columns in a relational database table into input fields of the same quantity.

Input and Output

- Input: table columns
- Output: fields

Parameter Description

Table 17-107 Operator parameters description

| Parameter | Description | Type | Mandatory | Default Value |
|--------------|--|------|-----------|---------------|
| Input fields | <p>Information about relational database input fields:</p> <ul style="list-style-type: none"> • position: position of input fields • field name: input field name • type: field type • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. | map | Yes | None |

Data Processing Rule

- Fields are generated in a specified order. Table columns to be converted are specified by **From** in step 2 of job configuration. If **Table column names** is set, the value is the table columns to be converted; if **Table column names** is not set, the table columns to be converted are all table columns in the table by default or the columns specified by the query conditions set by **Table SQL statement**.
- The number of input fields cannot be greater than number of specified columns; otherwise, all data becomes dirty data.
- If the field value does not match the actual type, the data in the line will become dirty data.

Example

Use SQL Server 2014 as an example. Run the following command to create a **test** table:

```
create table test (id int, name text, value text);
```

Insert three data lines to the test table:

```
insert into test values (1,'zhangshan','zhang');
```

```
insert into test values (2,'lisi','li');
```

insert into test values (3,'wangwu','wang');

Query the table:

| | id | name | value |
|---|----|-----------|-------|
| 1 | 1 | zhangshan | zhang |
| 2 | 2 | lisi | li |
| 3 | 3 | wangwu | wang |

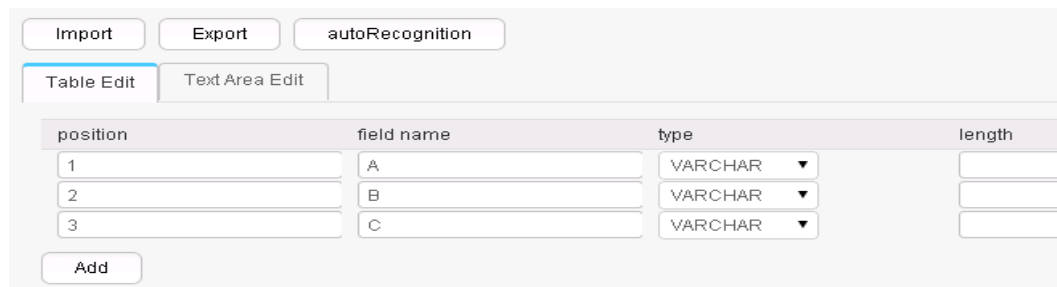
Configure the **Table Input** operator to generate the following fields:

After the data connector is set, click **Automatic Identification**. The system will automatically read fields in the database and select required fields for adding. You only need to optimize or modify the fields manually based on service scenarios.

NOTE

- This operation will overwrite existing data in the table.
- After you click **autoRecognition**, manually check the field types automatically identified by the system to ensure that they are consistent with the actual ones in the table.

For example, the system automatically identifies the **date** type in the Oracle database as the **timestamp** type. If you do not manually change the type, an error will be reported when data is queried in the Hive table.



| position | field name | type | length |
|----------|------------|---------|--------|
| 1 | A | VARCHAR | |
| 2 | B | VARCHAR | |
| 3 | C | VARCHAR | |

Configure the output operator to output data to HDFS or OBS. The result is as follows:

```
1,zhangshan,zhang
2,lisi,li
3,wangwu,wang
```

17.15.2.4 HBase Input

Overview

The **HBase Input** operator converts specified columns in an HBase table into input fields of the same quantity.

Input and Output

- Input: HBase table columns
- Output: fields

Parameter Description

Table 17-108 Operator parameter description

| Parameter | Description | Type | Mandatory | Default Value |
|--------------------|---|--------|-----------|---------------|
| Hbase Table Type | HBase table type. The options include normal (common HBase table) and phoenix . | enum | Yes | normal |
| HBase table name | HBase table name. Only one HBase table is supported. | string | Yes | None |
| HBase input fields | <p>HBase input information:</p> <ul style="list-style-type: none"> • family name: HBase column family name. • column name: HBase column name. • field name: Names of input fields. • type: Field type. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. • is rowkey: Indicates whether a column is a primary key column. A common HBase table can have only one primary key, while a phoenix table can have multiple primary keys. If multiple primary keys are configured, they are combined according to the configuration sequence. At least one primary key column must be configured. | map | Yes | None |

Data Processing Rule

- If the HBase table name does not exist, the job fails to be submitted.
- If the configured column names are inconsistent with the HBase table column names, the data cannot be read and the number of imported data records is 0.
- If the number of input field columns is greater than the number of field columns actually included in the original data, all data becomes dirty data.
- If the field value does not match the actual type, the data in the line will become dirty data.

Example

Use the data export from HBase to sqlserver2014 as an example.

In sqlserver2014, run the following statement to create an empty data test_1 for storing HBase data:

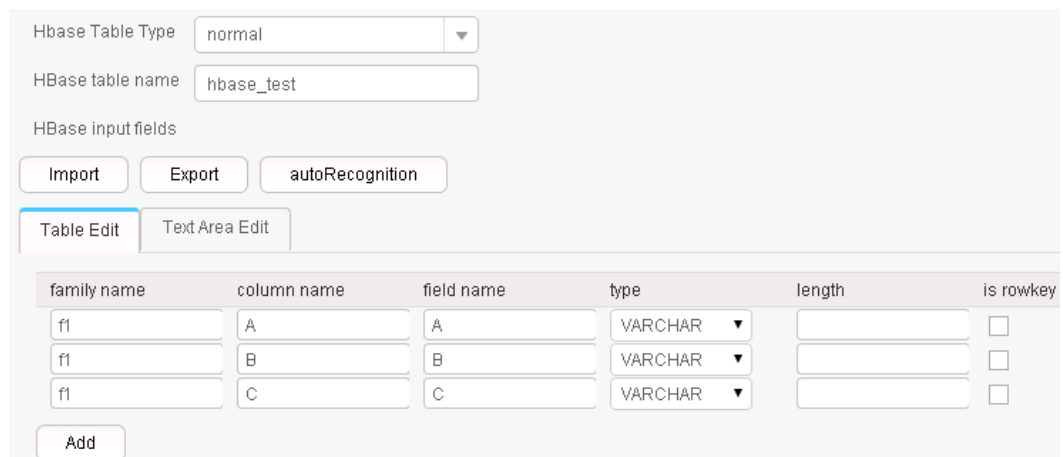
```
create table test_1 (id int, name text, value text);
```

Configure the **HBase Input** operator to generate fields A, B, and C.

After the database connection is set up, click **autoRecognition**. The system will automatically read fields in the database and select required fields for adding. You only need to optimize or modify the fields manually based on service scenarios.

NOTE

Performing this operation will overwrite existing data in the table.



| family name | column name | field name | type | length | is rowkey |
|-------------|-------------|------------|---------|--------|--------------------------|
| fl | A | A | VARCHAR | | <input type="checkbox"/> |
| fl | B | B | VARCHAR | | <input type="checkbox"/> |
| fl | C | C | VARCHAR | | <input type="checkbox"/> |

Use the **Table Out** operator to export A, B, and C to the test_1 table.

```
select * from test_1;
```

| | id | name | value |
|---|----|-----------|-------|
| 1 | 1 | zhangshan | zhang |
| 2 | 2 | lisi | li |
| 3 | 3 | wangwu | wang |

17.15.2.5 HTML Input

Overview

HTML Input operator imports a regular HTML file and converts elements in the HTML file into input fields.

Input and Output

Input: HTML file

Output: multiple fields

Parameter Description

Table 17-109 Operator parameters description

| Parameter | Description | Type | Mandatory | Default Value |
|----------------------|--|---------|-----------|---------------|
| parent tag | Upper-layer HTML tag of all fields for limiting the search scope. | string | Yes | None |
| Filename as field | User-defined field whose value is the name of the file that stores the current data. | string | No | None |
| Absolute file name | Whether the file name used as the value of Filename as field contains an absolute path. Selecting the option button indicates that the file name contains an absolute path; deselecting the option button indicates that the file name does not contain a path. | boolean | No | No |
| Validate input field | Whether to check the type matching between the input field and the value. If the value is NO , the field is not checked. If the value is YES , the field will be checked. If the input field does not match the value type, the line is skipped. | enum | Yes | YES |

| Parameter | Description | Type | Mandatory | Default Value |
|--------------|---|------|-----------|---------------|
| Input fields | <p>Information about input fields:</p> <ul style="list-style-type: none"> • position: Position of the field. The position sequence starts from 1. • field name: field name • field tag: field tag • keyword: A keyword can be configured to match the content of the tag. Wildcards are supported. For example, if the tag content is name, you can configure the keyword *name*. • type: field type • date format: If the field type is DATE, TIME, or TIMESTAMP, you need to specify the time format. If the field type is neither of them, the time format is invalid. The example time format is yyyyMMdd HH:mm:ss. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. | map | Yes | None |

Data Processing Rule

- **parent tag** is configured first to limit the search scope. The value of **parent tag** must exist; otherwise, the obtained content is empty.
- **Input fields** are configured so that the sub-tags can be used to precisely locate the tags of fields. If the tags are the same, keywords will be used for precise matching.
- The keyword is used to match the content of the field. The configuration method is similar to that of the **File filter** field in the **From** settings. The wildcard (*) is supported. The following three tags are provided to assist in locating the field:
 - a. **#PART**: indicates the values matched by wildcard *. If there are multiple *, you can specify an order from left to right and obtain content that matches the sequence number *. For example, **#PART1** indicates to

- obtain the value that matches the first * and **#PART8** indicates to obtain the value that matches the eighth *).
- b. **#NEXT**: indicates that you can obtain the value next to the value that matches the tag.
- c. **#ALL**: indicates that you can obtain all the values that match the tag.
- If the tag is configured incorrectly, the obtained value is empty, but no error is reported.

Example

The following figure shows the source file.

```
<html>
<body>
<table>
<tr>
<td>name:zhangshan</td>
<td>department:FusionInght</td>
<td>age:25</td>
</tr>
</table>
</body>
</html>
```

Configure the **HTML Input** operator to generate fields A, B, and C.

position	field name	field tag	keyword	type	date format	length
1	A	td	name:*PART1	VARCHAR		
2	B	td	department:*PAR	VARCHAR		
3	C	td	age:*PART1	VARCHAR		

Three fields are generated, as shown in the following figure.

```
zhangshan,FusionInght,25
```

17.15.2.6 Hive input

Overview

The **Hive Input** operator converts specified columns in an HBase table into input fields of the same quantity.

Input and Output

- Input: Hive table columns
- Output: fields

Parameters

Table 17-110 Operator parameters description

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
Hive database	Name of a Hive database	String	No	default
Hive table name	Name of the Hive table configured Only one Hive table is supported.	String	Yes	None
Partition filter	Configures the partition filter can export data of specific partitions. The parameter is null by default and data of the whole table can be exported. For example, to export data of a table whose partition field's locale value is CN or US , the input is as follows: locale = "CN" or locale = "US"	String	No	-
Hive input field	Configures the input information of Hive <ul style="list-style-type: none"> • column name: Hive column name. • field name: Input field name. • type: Field type. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. 	map	Yes	-

Data Processing Rule

- If the Hive table name does not exist, the job fails to be submitted.
- If the configured column names are inconsistent with the Hive table column names, the data cannot be read and the number of imported data records is 0.
- If the field value does not match the actual type, the data in the line will become dirty data.

Example

Use the data export from Hive to SQL Server 2014 as an example.

In SQL Server 2014, run the following statement to create an empty table **test_1** for storing Hive data. Run the following statement:

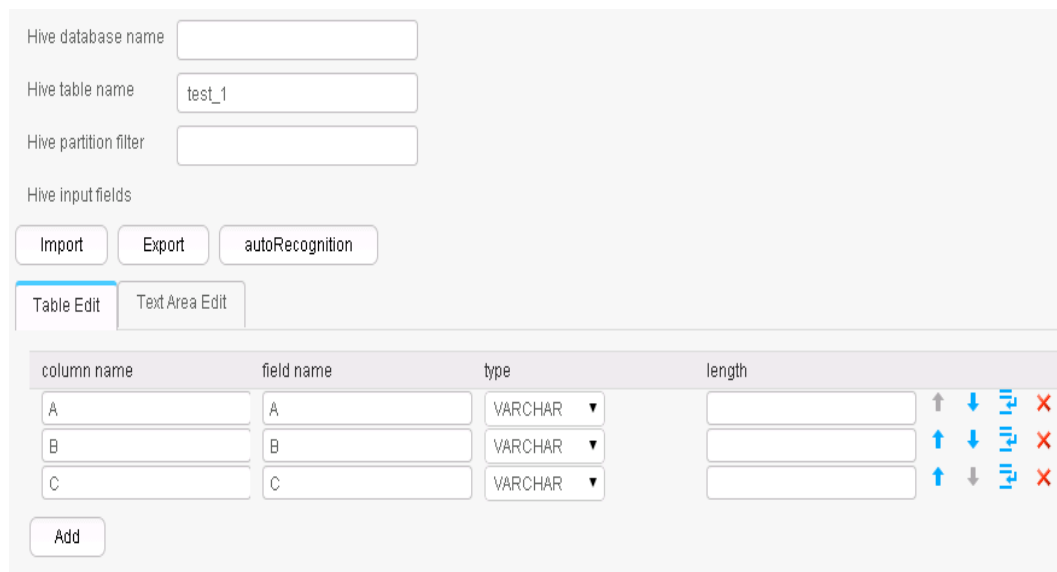
```
create table test_1 (id int, name text, value text);
```

Configure the **Hive Input** operator to generate fields A, B, and C.

After the data connector is set, click **Automatic Identification**. The system will automatically read fields in the database and select required fields for adding. You only need to optimize or modify the fields manually based on service scenarios.

NOTE

Performing this operation will overwrite existing data in the table.



column name	field name	type	length	
A	A	VARCHAR		↑ ↓ ↕ ×
B	B	VARCHAR		↑ ↓ ↕ ×
C	C	VARCHAR		↑ ↓ ↕ ×

Use the **Table Out** operator to export A, B, and C to the **test_1** table.

```
select * from test_1;
```

	id	name	value
1	1	zhangshan	zhang
2	2	lisi	li
3	3	wangwu	wang

17.15.2.7 Spark Input

Overview

The **Spark Input** operator converts specified columns in an SparkSQL table into input fields of the same quantity.

Input and Output

- Input: SparkSQL table column
- Output: fields

Parameters

Table 17-111 Operator parameters description

Parameter	Description	Type	Mandatory	Default Value
Spark database	Name of a Spark SQL database	String	No	default
Spark table name	Configures the SparkSQL table name. Only one SparkSQL table is supported.	String	Yes	None
Partition filter	Configures the partition filter can export data of specific partitions. The parameter is null by default and data of the whole table can be exported. For example, to export data of a table whose partition field's locale value is CN or US , the input is as follows: locale = "CN" or locale = "US"	String	No	-

Parameter	Description	Type	Mandatory	Default Value
Input fields of Spark	<p>Configures the input information of SparkSQL</p> <ul style="list-style-type: none"> column name: SparkSQL column name. field name: Input field name. type: Field type. length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. 	map	Yes	-

Data Processing Rule

- If the SparkSQL table name does not exist, the job fails to be submitted.
- If the configured column names are inconsistent with the SparkSQL table column names, the data cannot be read and the number of imported data records is 0.
- If the field value does not match the actual type, the data in the line will become dirty data.

Example

Use the data export from Spark to SQL Server 2014 as an example.

In SQL Server 2014, run the following statement to create an empty table **test_1** for storing SparkSQL data. Run the following statement:

```
create table test_1 (id int, name text, value text);
```

Configure the **Spark Input** operator to generate fields A, B, and C.

After the data connector is set, click **Automatic Identification**. The system will automatically read fields in the database and select required fields for adding. You only need to optimize or modify the fields manually based on service scenarios.

NOTE

Performing this operation will overwrite existing data in the table.

column name	field name	type	length
A	A	VARCHAR	
B	B	VARCHAR	
C	C	VARCHAR	

Use the **Table Out** operator to export A, B, and C to the **test_1** table.

```
select * from test_1;
```

	id	name	value
1	1	zhangshan	zhang
2	2	lisi	li
3	3	wangwu	wang

17.15.3 Conversion Operators

17.15.3.1 Long Date Conversion

Overview

The **Long Date Conversion** operator performs long integer and date conversion.

Input and Output

- Input: fields to be converted
- Output: new fields

Parameter Description

Table 17-112 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
convert type	Types of long integer and date conversion: <ul style="list-style-type: none"> • long to date: converts long integers to date. • long to time: converts long integers to time. • long to timestamp: converts long integers to timestamp. • date to long: converts date to long integers. • time to long: converts time to long integers. • timestamp to long: converts timestamp to long integers. 	enum	Yes	long to date
input field name	Name of input fields to be converted. Set this parameter to the names of fields generated in the previous conversion step.	string	Yes	None
output field name	Names of output fields.	string	Yes	None
field unit	Unit of a long integer field. According to convert type , the value is an input field or generated field. The options are second and millisecond .	enum	Yes	second
output field type	Output field type. The options are BIGINT , DATE , TIME , and TIMESTAMP .	enum	Yes	BIGINT
date format	Time field format, for example, yyyyMMdd HH:mm:ss .	string	No	None

Data Processing Rule

- If the original data includes null values, no conversion is performed.
- If the number of input field columns is greater than the number of field columns actually included in the original data, all data becomes dirty data.
- If a type conversion error occurs, the current data is saved as dirty data.

Example

Use the **CSV File Input** operator to generate fields A and B.

The following figure shows the source file.

```
1453431755874,2016-01-22 10:40:00
```

Configure the **Long Date Conversion** operator to generate four new fields C, D, E, and F. Their types are DATE, TIME, TIMESTAMP, and BIGINT, respectively.

The following figure shows the output of the conversion.

```
1453431755874,2016-01-22,2016-01-22,11:02:35,20160122 11:02:35,1453430400000
```

17.15.3.2 Null Value Conversion

Overview

The **null value conversion** operator replaces null values with specified values.

Input and Output

- Input: fields with null values
- Output: original fields with new values

Parameter Description

Table 17-113 Operator parameters description

Parameter	Description	Node Type	Mandatory	Default Value
Input field name	Names of fields that may have null values. Set this parameter to the names of existing fields.	string	Yes	None
Replace by this value	Specified values for replacing null values.	string	Yes	None

Data Processing Rule

When field values are empty, specified values are added.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
,value1  
key2,value2  
key3,
```

Configure the **null value conversion** operator, as shown in the following figure.

The screenshot shows the configuration interface for the null value conversion operator. At the top, there are 'Import' and 'Export' buttons. Below them are two tabs: 'Table Edit' (which is selected) and 'Text Area Edit'. The main area contains a table with two columns: 'input field name' and 'Replace by this value'. Under 'input field name', there are two input fields containing 'A' and 'B'. Under 'Replace by this value', there are two input fields containing 'newKey' and 'newValue'. At the bottom left, there is an 'Add' button.

After replacement, the values of fields A and B are as follows:

```
newKey,value1  
key2,value2  
key3,newValue
```

17.15.3.3 Constant Field Addition

Overview

The **Add Constants** operator generates constant fields.

Input and Output

- Input: none
- Output: constant fields

Parameter Description

Table 17-114 Operator parameters description

Parameter	Description	Type	Mandatory	Default Value
Constant fields	<p>Information about constant fields:</p> <ul style="list-style-type: none"> • output field name: Names of the configured fields. • type: field type • date format: If the field type is DATE, TIME, or TIMESTAMP, you need to specify the time format. If the field type is neither of them, the time format is invalid. The example time format is yyyyMMdd HH:mm:ss. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. • constant value: Constant value of the correct type. 	map	Yes	None

Data Processing Rule

This operator generates constant fields of the specified type.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
,value1
key2,value2
key3,
```

Configure the **Add Constants** operator to add fields C and D.

After adding the constants, fields A, B, C, and D are generated, as shown in the following figure.

```
,value1,constantsvalue1,2016
key2,value2,constantsvalue1,2016
key3,,constantsvalue1,2016
```

17.15.3.4 Random Value Conversion

Overview

Generate Random operator configures new values as random value fields.

Input and Output

- Input: none
- Output: random value fields

Parameter Description

Table 17-115 Operator parameters description

Parameter	Description	Type	Mandatory	Default Value
output field name	Names of generated random value fields.	string	Yes	None
length	Field length.	map	Yes	None
type	Field type. The options are VARCHAR , INTEGER , and BIGINT .	enum	Yes	VARCHAR

Data Processing Rule

The operator generates random value fields of specified type.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
,value1  
key2,value2  
key3,
```

Configure the random value conversion operator to generate fields C, D, and E.

output field name	type
C	VARCHAR ▼
D	INTEGER ▼
E	BIGINT ▼

Five fields are generated.

```
,value1,2druceak69ril,769974975,8452014577467885098  
key2,value2,7oq2dku93q9cg,1631427868,867914116689501757  
key3,,2jg5e7b1m17kq,654806209,2477823020516316030
```

The random value fields generated each time are different.

17.15.3.5 Concat Fields

Overview

The **Concat Fields** operator concatenates existing fields by using delimiters to generate new fields.

Input and Output

- Input: fields to be concatenated
- Output: new fields

Parameter Description

Table 17-116 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Output field name	Name of a field generated after concatenation.	string	Yes	None
Delimiter	Concatenation character. The value can be blank.	string	No	Empty string
Fields to be merged	Names of fields to be concatenated. field name must be set to the names of fields generated in the previous conversion step. Multiple field names can be added.	map	Yes	None

Data Processing Rule

- Use delimiters to concatenate the fields specified by **Fields to be merged** in order and assign the output to **Output field name**.
- If the value of a field is null, the value is changed to an empty string and then concatenated with other field values.

Example

Use the **CSV File Input** operator to generate fields A, B, and C.

The following figure shows the source file.

```
happy,new,year
welcome,to,2016
```

Configure the **Concat Fields** operator, set **Delimiter** to blank space, and generate field D.

After concatenation, fields A, B, C, and D are generated, as shown in the following figure.

```
happy,new,year,happy new year
welcome,to,2016,welcome to 2016
```

17.15.3.6 Extract Fields

Overview

The **Extract Fields** separates an existing field by using delimiters to generate new fields.

Input and Output

- Input: field to be separated
- Output: new fields

Parameter Description

Table 17-117 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Input field name	Name of a field to be separated. Set this parameter to the name of a field generated in the previous conversion step.	string	Yes	None
Delimiter	Delimiter.	string	Yes	None

Parameter	Description	Type	Mandatory	Default Value
Fields extracted	Fields generated after field separation. Multiple fields can be generated after field separation. <ul style="list-style-type: none"> position: Position of fields generated after field separation. output field name: Names of output fields. 	map	Yes	None

Data Processing Rule

- The value of the input field is separated by specified delimiters and the segments are assigned to the new fields.
- If the number of field columns after separation is greater than the actual number allowed by the original data, the line will become dirty data.

Example

Use the **CSV File Input** operator to generate field A.

The following figure shows the source file.

```
happy new year
welcome to 2016
```

Configure the **Extract Fields** operator, set **Delimiter** to blank space, and generate three fields B, C, and D.

Input field name:

Delimiter:

Fields extracted

position	output field name
<input type="text" value="1"/>	<input type="text" value="B"/>
<input type="text" value="2"/>	<input type="text" value="C"/>
<input type="text" value="3"/>	<input type="text" value="D"/>

After conversion, fields A, B, C, and D are generated, as shown in the following figure.

```
happy new year,happy,new,year
welcome to 2016,welcome,to,2016
```

17.15.3.7 Modulo Integer

Overview

The **Modulo Integer** operator performs modulo operations on integer fields to generate new fields.

Input and Output

- Input: integer fields
- Output: new fields

Parameter Description

Table 17-118 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Modulo fields	Modulo operation information: <ul style="list-style-type: none"> • input field name: Names of input fields. Set this parameter to the names of fields generated in the previous conversion step. • output field name: Names of output fields. • modulus: Values used for a modulo operation. 	map	Yes	None

Data Processing Rule

- The operator generates new fields and the values are those after the modulo operation.
- The field values must be integers; otherwise, the current line becomes dirty data.

Example

Use the **CSV File Input** operator to generate fields A and B.

The following figure shows the source file.


```
10,12  
2015,2016
```

Configure the **Modulo Integer** operator to generate two new fields C and D.

input field name	output field name	modulus
A	C	3
B	D	3

After the modulo operation, fields A, B, C, and D are generated, as shown in the following figure.

```
10,12,1,0  
2015,2016,2,0
```

17.15.3.8 String Cut

Overview

The **String Cut** operator cuts existing fields to generate new fields.

Input and Output

- Input: fields to be cut
- Output: new fields

Parameter Description

Table 17-119 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Fields to be cut	<p>Information about a cut field:</p> <ul style="list-style-type: none"> • input field name: Names of input fields. Set this parameter to the names of fields generated in the previous conversion step. • output field name: Names of output fields. • start position: Cutting start position, starting from sequence 1. • end position: Cutting end position. If the length of the cut string cannot be determined, you can set this parameter to -1, indicating the end of a string to be cut. • output field type: Type of output fields. • output field length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When output field type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When output field type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. 	map	Yes	None

Data Processing Rule

- **start position** and **end position** are used to cut the original fields and generate new fields.
- If **end position** is set to -1, the end of a string is to be cut. In other cases, the value of **end position** must be greater than the value of **start position**.
- If the value of **start position** or **end position** is greater than the length of the input field, the line will become dirty data.

Example

Use the **CSV File Input** operator to generate fields A and B.

The following figure shows the source file.

```
number,product  
FusionInsight,Bigdata
```

After configuring the **String Cut** operator, fields C and D are generated.

input field name	output field name	start position	end position	output field type	output field length
A	C	1	3	VARCHAR	
B	D	1	4	VARCHAR	

After cutting, the following fields are generated.

```
number,product,num,prod  
FusionInsight,Bigdata,Fus,Bigd
```

17.15.3.9 EL Operation

Overview

The **EL Operation** operator calculates field values and generates new fields. The algorithms that are currently supported include md5sum, sha1sum, sha256sum, and sha512sum.

Input and Output

- Input: fields to be converted
- Output: fields generated after the EL expression conversion

Parameter Description

Table 17-120 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Field generated by el operation	<p>EL expression configuration:</p> <ul style="list-style-type: none"> • name: Name of the expression output result. • el expression: Expression. The format is <i>expression name(input field name,value indicating whether to use lower case letters to indicate the output result)</i>, for example, md5sum(fieldname,true). <ul style="list-style-type: none"> - md5sum: generates md5 values. - sha1sum: generates sha1 values. - sha256sum: generates sha256 values. - sha512sum: generates sha512 values. • type: Type of the expression output result. VARCHAR is recommended. • date format: Format of the expression output result. • length: Length of the expression output result. 	map	Yes	None

Data Processing Rule

- The operator calculates fields values and generates new fields.
- The type of the new fields can only be VARCHAR.

Example

Use the **CSV File Input** operator to generate fields A and B.

The following figure shows the source file.

```
2016,year
year,2016
```

Configure the **EL Operation** operator to generate fields C, D, E, and F.

name	el expression	type	date format	length
C	md5sum(A,false)	VARCHAR		
D	sha1sum(A,true)	VARCHAR		
E	sha256sum(B,false)	VARCHAR		
F	sha512sum(B,true)	VARCHAR		

Six fields are generated, as shown in the following figure.

```
2016,year,95192C98732387165BF8E396C0F2DAD2,ab39c54239118a4b086b878b7878100f769dd1
97,4CB4EA25583C25647247AE96FC90225D99AD7A6FABC3E2C2FD13C502E323CD9E,779edfe0463b2
596e7a83e4c59083e19242e8c51eace8e2ec57704643be5e15ba80f79af227cf3ea2e2362b4081377
96a1d82cb0535652b99844bb9a62019563
year,2016,84CDC76CABF41BD7C961F6AB12F117D8,4ff0b1538469338a0073e2cdaab6a517801b6a
b4,DA6E2F539726FABD1F8CD7C9469A22B36769137975B28ABC65FE2DC29E659B77,da0ae9104086a
1c58f89f82766ac55a02c8ab44277ce39f959ec0e73391bef651c6f9793657396ce47fbd846068465
ccbf3056764424bed9be7789bd1101ace7
```

17.15.3.10 String Operations

Overview

The **String Operations** operator converts the upper and lower cases of existing fields to generate new fields.

Input and Output

- Input: fields whose case is to be converted
- Output: new fields after conversion

Parameter Description

Table 17-121 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Fields to be processed	<p>Information about fields for string case conversion:</p> <ul style="list-style-type: none"> input field name: Names of input fields. Set this parameter to the names of fields generated in the previous conversion step. output field name: Names of output fields. lower/upper: Indicates whether data is to be converted into uppercase letters or lowercase letters. 	map	Yes	None

Data Processing Rule

- Case conversion is performed for strings.
- If the input data is null, no case conversion is performed.

Example

Use the **CSV File Input** operator to generate fields A and B.

The following figure shows the source file.

```
number,product
FusionInsight,Bigdata
```

After configuring the **String Operations** operator, fields C and D are generated.

input field name	output field name	lower/upper
A	C	Upper
B	D	Lower

After conversion, four fields are generated, as shown in the following figure.

```
number,product,NUMBER,product
FusionInsight,Bigdata,FUSIONINSIGHT,bigdata
```

17.15.3.11 String Reverse

Overview

The **String Reverse** operator reverses existing fields to generate new fields.

Input and Output

- Input: fields to be reversed
- Output: new fields

Parameter Description

Table 17-122 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Fields to be reversed	<p>Information about fields for string reversal conversion:</p> <ul style="list-style-type: none"> • input field name: Names of input fields. Set this parameter to the names of fields generated in the previous conversion step. • output field name: Names of output fields. • type: Field type. • out field length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. 	map	Yes	None

Data Processing Rule

- Value reversal conversion is performed for fields.
- If the input data is null, no reversal conversion is performed.
- If the number of input field columns is greater than the number of field columns actually included in the original data, all data becomes dirty data.

Example

Use the **CSV File Input** operator to generate fields A and B.

The following figure shows the source file.

```
number,product  
FusionInsight,Bigdata
```

Configure the **String Reverse** operator to generate two new fields C and D.

input field name	output field name	type	output field length
A	C	VARCHAR	
B	D	VARCHAR	

After conversion, four fields are generated, as shown in the following figure.

```
number,product,rebmun,tcudorp  
FusionInsight,Bigdata,thgisInoisuF,atadgiB
```

17.15.3.12 String Trim

Overview

The **String Trim** operator clears spaces contained in existing fields to generate new fields.

Input and Output

- Input: fields whose spaces are to be cleared
- Output: new fields

Parameter Description

Table 17-123 Operator parameter description

Parameter	Description	Type	Mandatory	Default Value
Fields to be trimmed	<p>Information about fields for clearing spaces contained in strings:</p> <ul style="list-style-type: none"> • input field name: Names of input fields. Set this parameter to the names of fields generated in the previous conversion step. • output field name: Names of output fields. • trim type: Space clearing mode (clearing starting spaces, ending spaces, or starting and ending spaces). 	map	Yes	None

Data Processing Rule

- Clearing spaces at both ends of a value supports clearing spaces at the left end, at the right end, and at both ends.
- If the input data is null, no conversion is performed.
- If the number of input field columns is greater than the number of field columns actually included in the original data, all data becomes dirty data.

Example

Use the **CSV File Input** operator to generate fields A, B, and C.

The following figure shows the source file.

```
welcome ,to , 2016
happy ,new , year
```

Configure the **String Trim** operator to generate three new fields D, E, and F.

The screenshot shows the configuration interface for the String Trim operator. At the top, there are buttons for 'Import' and 'Export'. Below them are 'Table Edit' and 'Text Area Edit' tabs. The main configuration area is a table with three columns: 'input field name', 'output field name', and 'trim type'. The table contains three rows of configuration:

input field name	output field name	trim type
A	D	both
B	E	right
C	F	left

At the bottom of the table, there is an 'Add' button.

Six fields are generated, as shown in the following figure.

```
welcome ,to , 2016,welcome,to,2016
happy ,new , year,happy,new,year
```

17.15.3.13 Filter Rows

Overview

This **Filter Rows** operator filters rows that contain triggering conditions by configuring logic conditions.

Input and Output

- Input: fields used to create filter conditions
- Output: none

Parameter Description

Table 17-124 Operator parameters description

Parameter	Description	Type	Mandatory	Default Value
Condition logic connector	Condition logic connector. The options include AND and OR .	enum	Yes	AND
Conditions	Filter condition information: <ul style="list-style-type: none"> • input field name: Names of the input fields. Set this parameter to the names of the fields generated in the previous conversion step. • operator: Operator • comparative value. You can directly enter the value of a field referenced in the #{Existing field name} format. 	map	Yes	None

Data Processing Rule

- When the condition logic is **AND**, if no filtering condition is added, all data becomes dirty data; if the original data meets all the added filtering conditions, the current line becomes dirty data.
- When the condition logic is **OR**, if no filter condition is added, all data becomes dirty data; if the original data meets any of the added filter conditions, the current line becomes dirty data.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
test, product
FusionInsight,Bigdata
```

Configure the **Filter Rows** operator to filter out lines that contain **test**.

input field name	operator	comparative value
A	==	test

After the conversion, enter the original fields. The result is as follows:

```
FusionInsight,Bigdata
```

17.15.3.14 Update Fields Operator

Overview

The **Update Fields** operator updates fields values when certain conditions are met.

The types supported at present include **BIGINT**, **DECIMAL**, **DOUBLE**, **FLOAT**, **INTEGER**, **SMALLINT**, and **VARCHAR**. When the type is **VARCHAR** and the operator is **+**, strings will be added to the end of field values. The operator **-** is not supported. For other types, **+** and **-** indicate addition and subtraction of values. For all types, **=** indicates new value assignment.

Input and Output

Input: field

Output: input field

Parameter Description

Table 17-125 Operator parameters description

Parameter	Description	Type	Mandatory	Default Value
update field name	Fields to be updated	string	Yes	None
update operator	Operator, which can be +, -, or =.	enum	Yes	+
update value	Values to be updated	The type is the same as the field type.	No	None
Condition logic connector	Condition logic connector. The options include AND and OR .	enum	Yes	AND
Conditions	Filter condition information: <ul style="list-style-type: none"> • input field name: Names of the input fields. Set this parameter to the names of the fields generated in the previous conversion step. • operator: Operator • comparative value. You can directly enter the value of a field referenced in the #{Existing field name} format. 	map	Yes	None

Data Processing Rule

- The operator checks whether conditions are met. If yes, the operator updates the field values. If no, the operator does not update the field values.
- If the field values are digits, the updated values are digits.
- If the fields are of the string type, the operator - cannot be used.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
test, product
FusionInsight,Bigdata
```

Configure the **Update Fields** operator to update a value by adding **good** to the end of the value if the value is **test**.

The screenshot shows the configuration for the Update Fields operator. The main configuration area includes:

- update field name: A
- update operator: +
- update value: good
- Conditions logic connector: AND

Below this, there are 'Import' and 'Export' buttons, and a 'Conditions' section with 'Table Edit' and 'Text Area Edit' tabs. The 'Table Edit' tab is active, showing a table with the following configuration:

input field name	operator	comparative value
A	==	test

An 'Add' button is located below the table.

The following figure shows the output result.

```
testgood ,product  
FusionInsight,Bigdata
```

17.15.4 Output Operators

17.15.4.1 Hive output

Overview

The **Hive Output** operator exports existing fields to specified columns of a Hive table.

Input and Output

- Input: fields to be exported
- Output: Hive table

Parameters

Table 17-126 Operator parameters description

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
Hive file storage format	<p>Hive configuration file storage format. CSV, ORC, and RC are supported at present.</p> <p>NOTE</p> <ul style="list-style-type: none"> Parquet is a column-based storage format. In this format, the output field names of Loader be the same as the field names in Hive tables. For Hive of versions later than 1.2.0, a field name, instead of field number, is used to parse ORC files. Therefore, the output field names of Loader must be the same as those in Hive tables. 	enum	Yes	CSV
Hive file compression format	Hive table file compression format. Select a format from the drop-down list. If you select NONE or do not set this parameter, data is not compressed.	enum	Yes	NONE
Hive ORC file version	Version of the ORC file (when the storage format of the Hive table file is ORC).	enum	Yes	0.12
Output delimiter	Delimiter.	string	Yes	None

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
Output fields	<p>Information about output fields:</p> <ul style="list-style-type: none"> • position: Position of output fields. • field name: Names of output fields. • type: Field type. If type is set to DATE, TIME, or TIMESTAMP, you must specify a time format. If type is set to other values, the time format is invalid. An example time format is yyyyMMdd HH:mm:ss. • decimal format: scale and precision of the decimal. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. • partition key: indicates whether a column is a partition column. You can specify zero or multiple partition columns. If multiple primary keys are configured, they are combined according to the configuration sequence. 	map	Yes	None

Data Processing Rule

- The field values are exported to the Hive table.
- If one or more columns are specified as partition columns, the **Partition Handlers** feature is displayed on the **To** page in Step 4 of the job configuration. **Partition Handlers** specifies the number of handlers for processing data partitioning.
- If no column is designated as partition columns, input data does not need to be partitioned, and **Partition Handlers** is hidden by default.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
2016,year
year,2016
```

Configure the **Hive Output** operator to export a_str and b_str to the Hive table.

After the execution is complete, view the table data.

```
0: jdbc:hive2://10.52.0.97:21066/> select * from hive_test;
+-----+-----+
| hive_test.a_str | hive_test.b_str |
+-----+-----+
| 2016            | year            |
| year           | 2016           |
+-----+-----+
2 rows selected (1.6 seconds)
```

17.15.4.2 Spark Output

Overview

The **Spark Output** operator exports existing fields to specified columns of a Spark SQL table.

Input and Output

- Input: fields to be exported
- Output: SparkSQL table

Parameter Description

Table 17-127 Operator parameters description

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
Spark file storage format	<p>SparkSQL configuration file storage format. CSV, ORC, RC and PARQUET are supported at present.</p> <p>NOTE</p> <ul style="list-style-type: none"> PARQUET is a column-based storage format. In this format, the output field names of Loader be the same as the field names in the SparkSQL table. For Hive of versions later than 1.2.0, a field name, instead of field number, is used to parse ORC files. Therefore, the output field names of Loader must be the same as those in the SparkSQL table. 	enum	Yes	CSV
Spark file compression format	<p>SparkSQL table file compression format. Select a format from the drop-down list. If you select NONE or do not set this parameter, data is not compressed.</p>	enum	Yes	NONE
Spark ORC file version	<p>Version of the ORC file (when the storage format of the SparkSQL table file is ORC).</p>	enum	Yes	0.12
Output delimiter	<p>Delimiter.</p>	string	Yes	None

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
Output fields	<p>Information about output fields:</p> <ul style="list-style-type: none"> • position: Position of output fields. • field name: Names of output fields. • type: Field type. If type is set to DATE, TIME, or TIMESTAMP, you must specify a time format. If type is set to other values, the time format is invalid. An example time format is yyyyMMdd HH:mm:ss. • decimal format: scale and precision of the decimal. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. • partition key: indicates whether a column is a partition column. You can specify zero or multiple partition columns. If multiple primary keys are configured, they are combined according to the configuration sequence. 	map	Yes	None

Data Processing Rule

- The field values are exported to the SparkSQL table.
- If one or more columns are specified as partition columns, the **Partition Handlers** feature is displayed on the **To** page in Step 4 of the job configuration. **Partition Handlers** specifies the number of handlers for processing data partitioning.
- If no column is designated as partition columns, input data does not need to be partitioned, and **Partition Handlers** is hidden by default.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
2016, year
year, 2016
```

Configure the **Spark Output** operator to export A and B to the SparkSQL table.

17.15.4.3 Table Output

Overview

The **Table Output** operator exports output fields to specified columns in a relational database table.

Input and Output

- Input: fields to be exported
- Output: relational database table

Parameters

Table 17-128 Operator parameters description

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
Output delimiter	Delimiter. NOTE This configuration applies only to the MySQL dedicated connector. If the data column content contains the default delimiter, you need to set a user-defined delimiter. Otherwise, data disorder may occur.	stri ng	No	,

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
Line delimiter	Line delimiter, which can be any string specified by users based on the actual situation. Any character string is supported. The OS line delimiter is used by default. NOTE This configuration applies only to the MySQL dedicated connector. If the data column content contains the default delimiter, you need to set a user-defined delimiter. Otherwise, data disorder may occur.	string	No	\n
Output fields	Information about relational database output fields: <ul style="list-style-type: none"> • field name: Names of output fields. • table column name: Names of database table columns. • type: Field type. The value must be consistent with the field type configured in the database. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. 	map	Yes	None

Data Processing Rule

The field values are exported to the table.

Example

Use the data export from HBase to sqlserver2014 as an example.

In sqlserver2014, run the following statement to create an empty data test_1 for storing HBase data. Run the following statement:

```
create table test_1 (id int, name text, value text);
```

Use the HBase Input operator to generated three fields A, B, and C.

Use the **Table Output** operator to export A, B, and C to the test_1 table.

The command output is as follows:

	id	name	value
1	1	zhangshan	zhang
2	2	lisi	li
3	3	wangwu	wang

17.15.4.4 File Output

Overview

The **File Output** operator uses delimiters to concatenate existing fields and exports new fields to a file.

Input and Output

- Input: fields to be exported
- Output: files

Parameter Description

Table 17-129 Operator parameters description

Parameter	Description	Type	Mandatory	Default Value
Output delimiter	Set a delimiter.	string	Yes	None

Parameter	Description	Type	Mandatory	Default Value
Line breaker	Line delimiter, which can be any string specified by users based on the actual situation. Any character string is supported. The OS line delimiter is used by default.	string	No	\n
Output fields	<p>Information about output fields:</p> <ul style="list-style-type: none"> • position: Position of output fields. • field name: Names of output fields. • type: Field type. If type is set to DATE, TIME, or TIMESTAMP, you must specify a time format. If type is set to other values, the time format is invalid. The example time format is yyyyMMdd HH:mm:ss. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. 	map	No	None

Data Processing Rule

The field is exported to a file.

Example

Use the **CSV File Input** operator to generate two fields A and B.

The following figure shows the source file.

```
number,product
FusionInsight,Bigdata
```

Configure the **File Output** operator, set **Output delimiter** to a comma (,), and export A and B to a file, as shown in the following figure.

The following figure shows the result.

```
number,product
FusionInsight,Bigdata
```

17.15.4.5 HBase Output

Overview

The **HBase Output** operator exports existing fields to specified columns of an HBase Outputtable.

Input and Output

- Input: fields to be exported
- Output: HBase table

Parameters

Table 17-130 Operator parameters description

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
HBase table type	HBase table type. The options include normal (common HBase table) and phoenix.	enu m	Yes	norm al

Parameter	Description	No de Type	Man dator y	Defa ult Valu e
NULL value processing mode	Null value processing mode. Selecting the option button indicates to convert null values to empty strings and save them. Deselecting the option button indicates the data is not saved.	boo lea n	No	The optio n butto n is not select ed.
HBase output fields	<p>HBase output information:</p> <ul style="list-style-type: none"> • field name: Names of output fields. • table name: HBase table name. • column family name: HBase column family name. If it is not configured during HBase/Phoenix table creation, the default value is '0'. • column name: HBase column name. • type: Field type. If type is set to DATE, TIME, or TIMESTAMP, you must specify a time format. If type is set to other values, the time format is invalid. An example time format is yyyyMMdd HH:mm:ss. • length: Field value length. If the actual field value is excessively long, the value is cut based on the configured length. When type is set to CHAR, spaces are added to the field value for supplement if the actual field value length is less than the configured length. When type is set to VARCHAR, no space is added to the field value for supplement if the actual field value length is less than the configured length. • Primary Key: Indicates whether a column is a primary key column. A common HBase table can have only one primary key, while a phoenix table can have multiple primary keys. If multiple primary keys are configured, they are combined according to the configuration sequence. At least one primary key column must be configured. 	ma p	Yes	None

Data Processing Rule

- The field values are exported to the HBase table.
- When the original data contains NULL values, if the **NULL value processing mode** is selected, the NULL values are converted to empty strings and saved. If the **NULL value processing mode** button is not selected, the data is not saved.

Example

Using table input as an example, after the fields are generated, the HBase Output operator exports them to the related HBase table and stores the data in the test table, as shown in the following figure.

	id	name	value
1	1	zhangshan	zhang
2	2	lisi	li
3	3	wangwu	wang

Create an HBase table.

```
create 'hbase_test','f1','f2';
```

Configure the **HBase Output** operator, as shown in the following figure.

After the job execution is complete, view the data in the hbase_test table.

```
hbase(main):001:0> scan 'hbase_test'
ROW
1
1
2
2
3
3
3 row(s) in 0.2720 seconds

COLUMN+CELL
column=f1:B, timestamp=1455855645760, value=zhangshan
column=f1:C, timestamp=1455855645760, value=zhang
column=f1:B, timestamp=1455855645760, value=lisi
column=f1:C, timestamp=1455855645760, value=li
column=f1:B, timestamp=1455855645760, value=wangwu
column=f1:C, timestamp=1455855645760, value=wang
```

17.15.4.6 ClickHouse Output

Overview

The **ClickHouse Output** operator exports existing fields to specified columns of a ClickHouse table.

Input and Output

- Input: fields to be exported
- Output: ClickHouse table

Parameters

Table 17-131 Operator parameters

Parameter	Description	Type	Mandatory	Default Value
ClickHouse database name	Database where the ClickHouse table is located.	string	Yes	default
ClickHouse table name	Name of the ClickHouse table to which data is written.	string	Yes	None

Data Processing Rule

The field values are exported to the ClickHouse table.

Example

Use the **CSV File Input** operator to generate 12 fields.

The following figure shows the source file.

```
1, 'b', 'abcd', '2021-06-15', '12:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
2, 'abc', 'abcd', '2021-06-15', '12:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
3, 'ab', 'abcd', '2021-06-15', '12:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
4, 'abcdef', 'abcd', '2021-06-15', '12:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
5, 'a', 'abcd', '2021-06-15', '12:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
6, 'bg', 'cde', '2020-06-15', '13:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
7, 'f', 'cde', '2020-06-15', '13:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
8, 'h', 'cde', '2020-06-15', '13:00:06', '2021-06-15 12:00:06', 1, 12, 6.8, 18.6, 12.8, true
```

Run the following statements to create a ClickHouse table:

```
CREATE TABLE IF NOT EXISTS testck4 ON CLUSTER default_cluster(
```

```
  a Int32,
```

```
  b VARCHAR(100) NOT NULL,
```

```
  c char(100),
```

```
  d DateTime,
```

```
  e DateTime,
```

```
  f DateTime,
```


```
  g smallint,
```

```


h bigint,
l Float32,
j Float64,
k decimal(10,2),
m boolean
)
ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/testck4',
'{replica}')
PARTITION BY toYYYYMM(d)ORDER BY a;

```

Configure the **ClickHouse Output** operator, as shown in the following figure.



After the job execution is complete, view the data in the **testck4** table.



17.15.5 Associating, Editing, Importing, or Exporting the Field Configuration of an Operator

Scenario

This section describes how to associate, import, or export the field configuration information of an operator when creating or editing a Loader job.

- Associating the field configuration of an operator
Associate the field configuration information of an input operator with an output operator.
- Editing the field configuration of an operator
Edit the field configuration information of an operator.
- Importing the field configuration of an operator
Import the field configuration information to an operator by using an operator export file or operator template file.
- Exporting the field configuration of an operator
Export the field configuration information of an operator to a JSON file and save the file to a local directory.

Prerequisites

You have obtained the username and password for logging in to the Loader web UI.


Procedure

- **Associating Field Configuration of an Operator**

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-82 Loader web UI



Step 2 Edit an existing job or create a new job. The **Transform** page is displayed.

Step 3 Double-click a specified input operator (such as **CSV File Input**) to go to the edit page. Add the configuration information to the parameter table of the input field.

Step 4 Double-click a specified output operator (such as **File Output**) to go to the edit page, click **associate**, and select the required field information in the displayed **associate** dialog box.

NOTE

- The field name already exists in the field table of the output operator and is not displayed in the **associate** window.
- You can also select the required field from the **field name** list. The corresponding configuration information is displayed in the parameter table of the output field.

Step 5 Click **OK**. The selected field is displayed in the parameter table of the output field.


----End

• **Editing the Field Configuration of an Operator**

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-83 Loader web UI

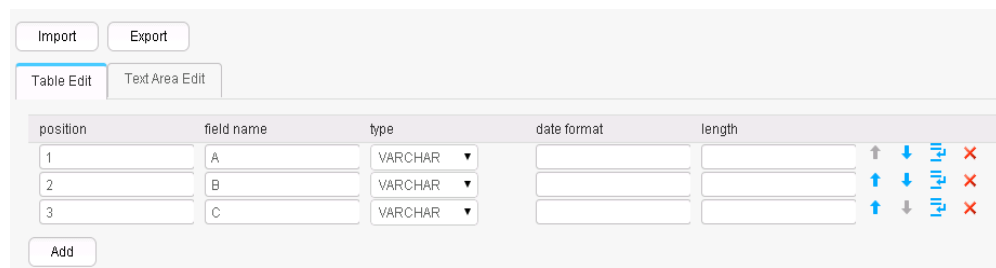


Step 2 Edit an existing job or create a new job. The **Transform** page is displayed.

Step 3 Double-click a specified operator (such as **CSV File Input**) to go to the edit page. On the **Table Edit** tab page of the input field, click **Add** and enter the field information based on the parameter requirements of the operator.

Step 4 You can move (up or down), insert a row under, and delete a field by clicking buttons corresponding to the field.

Click **Text Area Edit** to edit the field list in text format. Use commas (,) to separate field attributes.



Step 5 Click **OK**.


----End

• **Importing the Field Configuration of an Operator**

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-84 Loader web UI




Step 2 Edit an existing job or create a new job. The **Transform** page is displayed.

Step 3 Double-click a specified operator to go to the editing page and add related configuration information to the parameter table of the input or output field. Click **Import**.

Step 4 Select an import type.

- **Export File**
Field configuration information is imported by using the JSON file exported by the operator.
- **Specified Template**
Field configuration information is imported by using the TXT file compiled based on the operator template.

Step 5 Click  and select the upload file path.

Step 6 Click **Upload**. The field configuration information is imported to the operator.


----End

- **Exporting the Field Configuration of an Operator**

Step 1 Access the Loader web UI.

1. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
2. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.
3. Click **LoaderServer(Node name, Active)**. The Loader web UI is displayed.

Figure 17-85 Loader web UI



Step 2 Edit an existing job or create a new job. The **Transform** page is displayed.

Step 3 Double-click a specified operator to go to the editing page, add related configuration information to the parameter table of the input or output field, and click **Export**.

Step 4 Select an export type.

- All
All field information is exported as a JSON file and saved to a local directory.
- Specified Field Name
Fields selected in the field list are exported as a JSON file and saved to a local directory.

Step 5 Click **OK**.

----End

17.15.6 Using Macro Definitions in Configuration Items

When creating or editing Loader jobs, users can use macro definitions during parameter configuration. Then the parameters can be automatically changed to corresponding macro values when a job is implemented.

NOTE

- The macro definitions take effect in the job only.
- Macro definitions can be imported and exported together with an import or export job. If a job uses macro definitions, the exported job includes the macro definitions. Macro definitions are imported by default when a job is imported.
- For details about the format of the first parameter in the **dateformattime** macro, see **java.text.SimpleDateFormat.java**. The restrictions of the target system must be followed. For example, HDFS or OBS directories do not support special characters.

Macro Definitions of Loader

At present, Loader supports the following time macro definitions by default:

Table 17-132 Common macro definitions of Loader

Name	Result After the Replacement	Description
@{dateformat("yyyy-MM-dd")}@	2016-05-17	Indicates the current date.
@{dateformat("yyyy-MM-dd HH:mm:ss")}@	2016-05-17 16:50:00	Indicates current date and time
@{timestamp()}@	1463476137557	Indicates milliseconds since 1970.

Name	Result After the Replacement	Description
<code>@{dateformat("yyyy-MM-dd HH:mm:ss",-7,DAYS)}@</code>	2016-05-10 16:50:00	Indicates the latest seven days (the present time minus seven days). The second parameter supports addition and subtraction. The third parameter is a time unit for calculation. According to definitions in the java.util.concurrent.TimeUnit.java , time units include DAYS, HOURS, MINUTES, and SECONDS.

In the following scenarios, parameters can be configured by using macro definitions.

- Specifying a data directory that is named by the current date
The parameter is set to `/user/data/inputdate_@{dateformat("yyyy-MM-dd")}@`.
- Querying data in the latest seven days by using SQL
`select * from table where time between '@{dateformat("yyyy-MM-dd HH:mm:ss",-7,DAYS)}@' and '@{dateformat("yyyy-MM-dd HH:mm:ss")}@'`
- Specifying a table that is named by the current date
The parameter is set to `table_@{dateformat("yyyy-MM-dd")}@parmvalue`.

17.15.7 Operator Data Processing Rules

In Loader data import and export tasks, each operator defines different processing rules for null values and empty strings in raw data. Dirty data cannot be imported or exported.

The following table describes the operator data processing rules for each conversion procedure.

Table 17-133 Data processing rules

Procedure	Description
CSV file input	<ul style="list-style-type: none"> • If a delimiter appears twice consecutively in the original data, an empty string field is generated. • It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. • If a type conversion error occurs, the current data is saved as dirty data.
Fixed file input	<ul style="list-style-type: none"> • If the original data includes null values, no conversion is performed. • It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. • If the configured field conversion type is different from the actual type of the original data, all data becomes dirty data. For example, convert the string type to the numeric type. • If the configured field split length is greater than the length of the original field value, the data split fails and the current line becomes dirty data.
Table input	<ul style="list-style-type: none"> • If the original data includes null values, no conversion is performed. • It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. • If the configured field conversion type is different from the actual type of the original data, all data becomes dirty data. For example, convert the string type to the numeric type.

Procedure	Description
HBase input	<ul style="list-style-type: none"> ● If the original data includes null values, no conversion is performed. ● If the HBase table name is incorrect, all data becomes dirty data. ● If the primary key column is not configured in Is rowkey, all data becomes dirty data. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. ● If the configured field conversion type is different from the actual type of the original data, all data becomes dirty data. For example, convert the string type to the numeric type.
Long integer time conversion	<ul style="list-style-type: none"> ● If the original data includes null values, no conversion is performed. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. ● If a type conversion error occurs, the current data is saved as dirty data.
Null value conversion	<ul style="list-style-type: none"> ● If the original data contains null values, data is converted to a specified value. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data.
Random value conversion	Processing of null value and empty string is not involved, and dirty data is not generated.
Constant field addition	Processing of null value and empty string is not involved, and dirty data is not generated.
Concat fields	<ul style="list-style-type: none"> ● If the original data contains null values, data is converted to empty string. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data.
Extracts fields	<ul style="list-style-type: none"> ● If the original data contains null values, the current line becomes dirty data. ● If the number of field columns after separation is greater than the actual number allowed by the original data, the line will become dirty data.

Procedure	Description
Modulo integer	<ul style="list-style-type: none"> ● If the original data contains null values, the current line becomes dirty data. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. ● If data type conversion fails, the current line becomes dirty data.
String cut	<ul style="list-style-type: none"> ● If the input data is null, no conversion is performed. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. ● If the start or end position of the string to be truncated is greater than the length of the input field, the current line becomes dirty data.
EL operation	<ul style="list-style-type: none"> ● If the input data is null, no conversion is performed. ● Enter the value of one or more fields and output the calculation result. ● When the input type is incompatible with the operator, the current row is dirty data.
String case conversion	<ul style="list-style-type: none"> ● If the input data is null, no conversion is performed. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data.
String reverse	<ul style="list-style-type: none"> ● If the input data is null, no conversion is performed. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data.
String trim	<ul style="list-style-type: none"> ● If the input data is null, no conversion is performed. ● It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data.

Procedure	Description
Filter rows	<ul style="list-style-type: none"> When the condition logic is AND, if no filter condition is added, all data becomes dirty data; if the original data meets all the added filter conditions, the current line becomes dirty data. When the condition logic is OR, if no filter condition is added, all data becomes dirty data; if the original data meets all the added filter conditions, the current line becomes dirty data.
File output	<ul style="list-style-type: none"> If the input data is null, no conversion is performed.
Table output	<ul style="list-style-type: none"> It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. If data type conversion fails, the current line becomes dirty data.
HBase output	<ul style="list-style-type: none"> If the original data contains null values and Store null column is set to true, data is converted to empty string and saved. If Store null column is set to false, data will not be saved. It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. If data type conversion fails, the current line becomes dirty data.
Hive output	<ul style="list-style-type: none"> If one or more columns are designated as partition columns, the Partition Handlers feature is displayed on the To page. Partition Handlers specifies the number of handlers for processing data partitioning. If no column is designated as partition columns, input data does not need to be partitioned, and Partition Handlers is hidden by default. It can be configured that all data becomes dirty data when the number of input field columns is greater than the number of field columns actually included in the original data. If data type conversion fails, the current line becomes dirty data.

17.16 Client Tools

17.16.1 Running a Loader Job by Using Commands

Scenario

Generally, users can manually manage data import and export jobs on the Loader UI. If you need to update and run Loader jobs by executing the shell script, you must configure the installed Loader client.

NOTE

Loader is incompatible with the client of an earlier version. If you reinstall the cluster or the Loader service, download and install the client again, and then use the client.

This section applies to MRS 3.x or later.

Prerequisites

- The Loader client has been installed. During the installation of the Loader client using a non-root user, if another user wants to use the client, the user needs to be authorized by the user who installs the client or a user with more rights (the Loader client installation directory needs to be granted with right 755). Please pay attention to the security problems after the authorization.
- The user for accessing the Loader service has been created. If the user is a machine-machine user, the keytab file must be downloaded..

Procedure

Step 1 Configure the Loader shell client.

1. Log in to the node where the client is located as the user who installs the client.
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 go to the Loader client installation directory, for example, **/opt/client/Loader**:

```
cd /opt/client/Loader
```

4. Run the following command to configure environment variables:

```
source/opt/client/bigdata_env
```

5. If the cluster is in security mode, run the following command to authenticate the user. In normal mode, user authentication is not required.

```
kinit Component service user
```

6. Run the following command to modify the tool authorization configuration file **login-info.xml**, save the file, and exit. For the parameters in the configuration file, see [Table 17-134](#).

```
vi loader-tools-1.99.3/loader-tool/job-config/login-info.xml
```

Table 17-134 Parameters of **login-info.xml**

Parameter	Description
hadoop.config.path	Storage directory of the core-site.xml , hdfs-site.xml , and krb5.conf configuration files of the MRS cluster. These three files are stored in the Loader Client installation directory/Loader/loader-tools-1.99.3/loader-tool/hadoop-config/ directory by default.
authentication.type	Authentication type of the Loader service. Set this parameter based on MRS cluster authentication mode. <ul style="list-style-type: none"> - kerberos indicates the security mode. - simple indicates the normal mode.
user.keytab	Whether to use the keytab file for authentication. The options are true , and false .
authentication.user	User for login when the normal mode or password authentication is used. In the keytab login mode, this parameter does not need to be set.
authentication.password	Encrypted password of the user for accessing the Loader service if the keytab file authentication is not used in the security mode. NOTE Run the following command to encrypt the password as the user who installs the client. When the encryption tool runs for the first time, a random dynamic key is automatically generated and stored in .loader-tools.key . The encryption tool uses this dynamic key to encrypt passwords every time. After .loader-tools.key is deleted, a new random key will be generated and stored in .loader-tools.key when the encryption tool runs. sh Loader client installation directory/Loader/loader-tools-1.99.3/encrypt_tool password

Parameter	Description
authentication.principal	Machine-Machine username for accessing the Loader service when the keytab file authentication is used in the security mode.
authentication.keytab	Absolute keytab file directory of the Machine-Machine user for accessing the Loader service when the keytab file authentication is used in the security mode.
zookeeper.quorum	IP address and port for accessing ZooKeeper. The value format is IP1:port,IP2:port,IP3:port . The default port number is 2181.
sqoop.server.list	Floating IP address and port for accessing Loader. The value format is floatip:port . The default port number is 21351 .

Step 2 Use the Loader shell client.

1. Run the following command to go to the Loader shell client directory. For example, if the Loader client installation directory is **/opt/client/Loader**, run the following command:
cd /opt/client/Loader/loader-tools-1.99.3/shell-client/
2. Run the following command to use the Loader shell client to run a job:
./submit_job.sh -n <arg> -u <arg> -jobType <arg> -connectorType <arg> -frameworkType <arg>

Table 17-135 Parameters of the Loader shell client tool

Parameter	Description
-n	(Mandatory) Job name.
-u	(Mandatory) If the parameter is set to y , the job parameters are updated and the job is executed. In this scenario, parameters -jobType , -connectorType , and -frameworkType need to be set. If the parameter is set to n , the job is directly executed without updating parameters.

Parameter	Description
-jobType	<p>Job type. This parameter is mandatory when -u is set to y. import indicates the data import job. export indicates the data export job.</p>
-connectorType	<p>Connector type. This parameter is mandatory when -u is set to y. Parameters of external data sources can be modified as required.</p> <p>sftp indicates the connector is an SFTP connector.</p> <ul style="list-style-type: none"> - In a data import job, you can modify the source file input path -inputPath, the source file encode format -encodeType, and the suffix -suffixName added to the input file after the source file is imported. - In a data export job, you can modify the output path -outputPath or the name of the exported file. <p>rdb indicates the connector is a relational database connector.</p> <ul style="list-style-type: none"> - In a data import job, you can modify the database mode name -schemaName, table name -tableName, SQL statement -sql, names of columns to be imported -columns, and names of partition columns -partitionColumn. - In a data export job, you can modify the database mode name -schemaName, table name -tableName, and the temporary table name -stageTableName.

Parameter	Description
-frameworkType	<p>Data storage type on MRS. This parameter is mandatory when -u is set to y. Parameters of data storage types can be modified as required.</p> <p>hdfs indicates that the HDFS is used to store data on Hadoop.</p> <ul style="list-style-type: none"> - In a data import job, you can modify the number of started maps -extractors and the storage directory of imported data in the HDFS -outputDirectory. - In a data export job, you can modify the number of started maps -extractors, the input path of data exported from the HDFS -inputDirectory, and the file filter criteria of the data export job -fileFilter. <p>hbase indicates that HBase is used to store data on MRS. In the data import and export job, you can modify the number of started maps -extractors.</p>

----End

Task Examples

- Run a job whose name is **sftp-hdfs** without updating job parameters:
`./submit_job.sh -n sftp-hdfs -u n`
- Update the input path, encoding type, suffix, output path, and number of started maps of the data import job whose name is **sftp-hdfs**, and run the job:
`./submit_job.sh -n sftp-hdfs -u y -jobType import -connectorType sftp -inputPath /opt/tempfile/1 -encodeType UTF-8 -suffixName " -frameworkType hdfs -outputDirectory /user/user1/tttest -extractors 10`
- Update the database mode, table name, and output path of the data import job whose name is **db-hdfs**, and run the job.
`./submit_job.sh -n db-hdfs -u y -jobType import -connectorType rdb -schemaName public -tableName sq_submission -sql " -partitionColumn sqs_id -frameworkType hdfs -outputDirectory /user/user1/dbdbt`

17.16.2 loader-tool Usage Guide

Overview

loader-tool is a Loader client tool. It consists of three tools: **lt-ucc**, **lt-ucj**, **lt-ctl**.

Loader supports two modes, parameter mode and job template mode. Either mode can be used to create, update, query, and delete connectors, and to create, update, query, delete, start, and stop Loader jobs.

This section applies to MRS 3.x or later.

 **NOTE**

loader-tool implements an asynchronous interface. After a command is submitted, the command output is not returned to the console in real time. Therefore, the results of the creation, update, query, and deletion operations on a connector and the creation, update, query, deletion, start, and stop operations on a Loader job must be confirmed on the Loader WebUI or by querying server logs.

- Parameter mode:

Add a parameter invoking script with specific parameters.

- Job template mode:

Change the values of all parameters in a job template and reference the job template when invoking a script.

After a Loader client is installed, the system automatically generates job templates for various scenarios in the *Loader client installation directory*/**loader-tools-1.99.3/loader-tool/job-config/** directory. The parameters vary according to job templates. Job templates contain information about jobs and associated connectors.

Job templates are XML files. The file name format is *original data location-to-new data location.xml*, for example, **sftp-to-hdfs.xml**. If a job supports conversion step, a json conversion step configuration file with the same name exists, for example, **sftp-to-hdfs.json**.

 **NOTE**

Job templates contain the configuration information of connectors. During the connector creation and updating, only the connector information in job templates is invoked.

Scenarios

The parameters vary according to connectors or jobs.

- To modify some parameters, use the parameter mode.
- To create a connector or job, use the job template mode.

 **NOTE**

This tool currently supports the FTP, HDFS, JDBC, MySQL, Oracle, and Oracle dedicated connectors. If other types of connectors are used, you are advised to use the open-source sqoop-shell tool.

Parameters

For example, the Loader client installation directory is **/opt/client/Loader/**.

- **lt-ucc usage description**

lt-ucc is a connector configuration tool of loader-tool user-configuration-connection and is used to create, update, and delete connectors.

Table 17-136 lt-ucc script parameter description

Parameter	Description	Example Value
-help	Help information.	-
-a <arg>	Connector action. The values include create , update and delete for creating, updating, and deleting connectors respectively.	create
-at <arg>	Login authentication type. The values include kerberos and simple .	kerberos
-uk <arg>	Whether to use the keytab file.	true
-au <arg>	Login authentication username.	bar
-ap <arg>	Login authentication password. The value must be an encrypted password. The password encryption method is described as follows: sh Loader client installation directory/Loader/loader-tools-1.99.3/encrypt_tool non-encrypted user password NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.	-
-c <arg>	Login authentication principal.	bar
-k <arg>	Login authentication keytab file.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config/user.keytab
-h <arg>	Specifies the configuration file path of the MRS cluster.	-h /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config

Parameter	Description	Example Value
-l <arg>	Login template file.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml
-s <arg>	Floating IP address and port for Loader. Format: <i>floating IP address: port</i> The default port is 21351 .	127.0.0.1:21351
-w <arg>	Job template file path for obtaining job details.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml
-z <arg>	IP address and port number of ZooKeeper quorum instances. The format is <i>IP address: port</i> . Use commas (,) to separate multiple addresses and port numbers.	127.0.0.0:2181, 127.0.0.1:2181
-n <arg>	Connector name	vt_sftp_test
-t <arg>	Connector type	sftp-connector
-P <arg>	Used to update the value of an attribute. The format is - Pparam1=value1. param1 indicates the attribute name of the connector in the job template. Password parameters are required for updating SFTP and FTP connector information. -Pconnection.sftpPassword= <i>Encrypted password</i>	- Pconnection.sftpServerIp=10.6.26.11

A complete example is as follows:

```
./bin/lt-ucc -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -n vt_sftp_test -t sftp-connector -Pconnection.sftpPassword=d2NjX2NyeXB0ATQxNDU1MzVGNDM0MjQzOzQ2MzEzMzQ0MzYzMjM1MzgzMjM2MzgzMTQxMzgzMzNjQ0NDE0MjMwMzgzMzQ1NDMzMzM3MzU0MTQ0MzUzMzMDM2Mzk7OzMzMzUzMzMDMwOzU3MUFdNTA4MEJENTUwNDNEQTc2Mjg3QkJERTc4NDI3OzQyODc0NTMwN0FEQzkzNjA7NTc0MzQzNUY0MzUyNTk1MDU0NUY0NDQ1NDY0MTU1NEM1NDVGNDQ0RjRENDE0OTRFOzMwOzMzMzQzNTM2MzgzMzMTMyMzgzMzMzZlZlZlMwOw -Pconnection.sftpServerIp=10.6.26.111 -a update
```

Configuration description of a lt-ucc script job template:

Use the operation of saving SFTP data to HDFS as an example. Edit the **sftp-to-hdfs.xml** file in *Loader client installation directory/loader-tools-1.99.3/loader-tool/job-config/* directory. The connector configuration is as follows:

```
<!-- Database connection information -->
<sqoop.connection name="vt_sftp_test" type="sftp-connector">
<connection.sftpServerIp>10.96.26.111</connection.sftpServerIp>
<connection.sftpServerPort>22</connection.sftpServerPort>
<connection.sftpUser>root</connection.sftpUser>
<connection.sftpPassword>d2NjX2NyeXB0ATQxNDU1MzVGNDM0MjQzOzMzMzkzOTMwMzI0NTM5MzQzOTM1Mzk0NTMwMzIzNTM4NDZlNzQzMzIzNjQyMzMzMzMDM4MzMzMzQ1MzYzODQxMzQ7OzMyMzUzMzMDMwO0EzMTUzM0ExNTAyNDhENzE3QTRBRTICQkRBQzIFRkFEOzYyOEE4NTIDODc2MkMyNzU7NTc0MzQzNUY0MzUyNTk1MDU0NUY0NDQ1NDY0MTU1NEM1NDVGNDQ0RjRENDE0OTRFOzMwOzMzMzMzQzNTM2MzMzMTMyMzgzMzMzMzIzNzMwOw</connection.sftpPassword>
</sqoop.connection>
```

- Creation command:
./lt-ucc -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/ftp-to-hdfs.xml -a create
- Update command:
./lt-ucc -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/ftp-to-hdfs.xml -a update
- Deletion command:
./lt-ucc -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/ftp-to-hdfs.xml -a delete

- **lt-ucj usage description**

lt-ucj is a job configuration tool of loader-tool user-configuration-job and is used to create, update, and delete jobs.

Table 17-137 lt-ucj script parameter description

Parameter	Description	Example Value
-help	Help information.	-
-a <arg>	Job action. The values include create , update , and delete for creating, updating and deleting jobs respectively.	create
-at <arg>	Login authentication type. The values include kerberos and simple .	kerberos
-uk <arg>	Whether to use the keytab file.	true
-au <arg>	Login authentication username.	bar

Parameter	Description	Example Value
-ap <arg>	<p>Login authentication password. The value must be an encrypted password.</p> <p>The password encryption method is described as follows:</p> <p>sh <i>Loader client installation directory/Loader/loader-tools-1.99.3/encrypt_tool non-encrypted user password</i></p> <p>NOTE</p> <p>If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	-
-c <arg>	Login authentication principal.	bar
-k <arg>	Login authentication keytab file.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config/user.keytab
-h <arg>	Specifies the configuration file path of the MRS cluster.	-h /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config
-l <arg>	Login template file.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml
-s <arg>	<p>Floating IP address and port for Loader.</p> <p>Format: <i>floating IP address: port</i></p> <p>The default port is 21351.</p>	127.0.0.1:21351

Parameter	Description	Example Value
-w <arg>	Job template file for obtaining job details.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml
-z <arg>	IP address and port number of ZooKeeper quorum instances. The format is <i>IP address: port</i> . Use commas (,) to separate multiple addresses and port numbers.	127.0.0.0:2181, 127.0.0.1:2181
-n <arg>	Name of the job.	Sftp.to.Hdfs
-cn <arg>	Connector name	vt_sftp_test
-ct <arg>	Connector type	sftp-connector
-t <arg>	Job type. The values include IMPORT and EXPORT .	IMPORT
-trans <arg>	Job associated conversion step file.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.json
-priority <arg>	Job priority. The values include LOW , NORMAL , and HIGH .	NORMAL
-queue <arg>	Queues	default
-storageType <arg>	Storage type	HDFS
-P <arg>	Used to update the value of an attribute. The format is - Pparam1=value1. param1 indicates the attribute name of the connector in the job template. Password parameters are required for updating SFTP and FTP connector information. - Pconnection.sftpPassword= <i>Encrypted password</i>	- Pconnection.sftpServerIp=10.6.26.11

A complete example is as follows:

```
./bin/lt-ucj -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/  
job-config/login-info.xml -n Sftp.to.Hdfs -t IMPORT -ct sftp-connector -  
Poutput.outputDirectory=/user/loader/sftp-to-hdfs-test8888 -a update
```

Configuration description of a lt-ucj script job template:

Use the operation of saving SFTP data to HDFS as an example. Edit the file *loader client installation directory/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml*. The job configuration is as follows:

```
<!-- Job name, globally unique.-->  
<sqoop.job name="Sftp.to.Hdfs" type="IMPORT" queue="default" priority=" Priority NORMAL ">  
  
<!-- External data source parameter configuration -->  
<data.source connectionName="vt_sftp_test" connectionType="sftp-connector">  
<file.inputPath>/opt/houjt/hive/all</file.inputPath>  
<file.splitType>FILE</file.splitType>  
<file.filterType>WILDCARD</file.filterType>  
<file.pathFilter>*</file.pathFilter>  
<file.fileFilter>*</file.fileFilter>  
<file.encodeType>GBK</file.encodeType>  
<file.suffixName></file.suffixName>  
<file.isCompressive>FALSE</file.isCompressive>  
</data.source>  
  
<!-- MRS cluster, parameter configuration -->  
<hadoop.source storageType="HDFS" >  
<output.outputDirectory>/user/loader/sftp-to-hdfs</output.outputDirectory>  
<output.fileOprType>OVERRIDE</output.fileOprType>  
<throttling.extractors>3</throttling.extractors>  
<output.fileType>TEXT_FILE</output.fileType>  
</hadoop.source>  
  
<!-- Job associated conversion step file -->  
<sqoop.job.trans.file>/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-  
hdfs.json</sqoop.job.trans.file>  
</sqoop.job>
```

– Creation command:

```
./bin/lt-ucj -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-  
tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-  
tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml -a create
```

– Update command:

```
./bin/lt-ucj -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-  
tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-  
tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml -a update
```

– Deletion command:

```
./bin/lt-ucj -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-  
tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-  
tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml -a delete
```

- **lt-ctl usage description**

lt-ctl is a job management tool of loader-tool controller and is used to start or stop jobs, query job status and progress, and check whether jobs are running.

Table 17-138 lt-ctl script parameter description

Parameter	Description	Example Value
-help	Help information.	-

Parameter	Description	Example Value
-a <arg>	Job action. The values include status , start , stop , and is running for querying job status, starting or stopping jobs, and checking whether jobs are running.	create
-at <arg>	Login authentication type. The values include kerberos and simple .	kerberos
-uk <arg>	Whether to use the keytab file.	true
-au <arg>	Login authentication username.	bar
-ap <arg>	<p>Login authentication password. The value must be an encrypted password.</p> <p>The password encryption method is described as follows:</p> <p>sh Loader client installation directory/Loader/loader-tools-1.99.3/encrypt_tool non-encrypted user password</p> <p>NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	-
-c <arg>	Login authentication principal.	bar
-k <arg>	Login authentication keytab file.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config/user.keytab
-h <arg>	Specifies the configuration file path of the MRS cluster.	-h /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config

Parameter	Description	Example Value
-l <arg>	Login template file.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml
-n <arg>	Name of the job.	Sftp.to.Hdfs
-s <arg>	Floating IP address and port for Loader. Format: <i>floating IP address: port</i> The default port is 21351.	127.0.0.1:21351
-w <arg>	Job template file for obtaining job details.	/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml
-z <arg>	IP address and port number of ZooKeeper quorum instances. The format is <i>IP address: port</i> . Use commas (,) to separate multiple addresses and port numbers.	127.0.0.0:2181, 127.0.0.1:2181

- Command for starting jobs:
./bin/lt-ctl -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -n Sftp.to.Hdfs -a start
- Command for viewing job status:
./bin/lt-ctl -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -n Sftp.to.Hdfs -a status
- Command for checking whether jobs are running:
./bin/lt-ctl -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -n Sftp.to.Hdfs -a isrunning
- Command for stopping jobs:
./bin/lt-ctl -l /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -n Sftp.to.Hdfs -a stop

17.16.3 loader-tool Usage Example

Scenario

loader-tool can be used to create, update, query, and delete a connector or job by using a job template or setting parameters.

This section describes how to use loader-tool in the job template mode. The job of importing data from the SFTP server to HDFS is used as an example.

This section applies to MRS 3.x or later.

Prerequisites

The Loader client has been installed and configured. For details, see [Running a Loader Job by Using Commands](#).

Procedure

Step 1 Log in to the node where the client is located as the user who installs the client.

Step 2 Run the following command to go to the loader-tool directory on the Loader client, for example, `/opt/client/Loader/`:

```
cd /opt/client/Loader/loader-tools-1.99.3/loader-tool/
```

Step 3 Run the following command to modify the existing job template. For example, if the job template `sftp-to-hdfs.xml` already exists in the `/opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/`, run the following command:

```
vi /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml
```

```
<root>
<!-- Database connection information -->
<sqoop.connection name="vt_sftp_test" type="sftp-connector">
<connection.sftpServerIp>10.96.26.111</connection.sftpServerIp>
<connection.sftpServerPort>22</connection.sftpServerPort>
<connection.sftpUser>root</connection.sftpUser>
<connection.sftpPassword>d2NjX2NyeXB0ATQxNDU1MzVGNDM0MjQzOzMzMzkzOTMwMzI0NTM5MzQzOT
M1Mzk0NTMwMzIzNTM4NDEzNzQ2MzIzNjQyMzMzMzMDM4MzMzNzQ1MzYzODQxMzQ7OzMzMzUzMDMw
O0EzMTUzM0ExNTAyNDhENzE3QTRBRTICQkRBQzIjFRkFE0zYyOEE4NTIDODc2MkMyNzU7NTc0MzQzNUY0
MzUyNTk1MDU0NUY0NDQ1NDY0MTU1NEM1NDVGNzQ0RjRENDE0OTRFOzMwOzmxMzQzNTM2MzZM
TMzMzgzMzMzIzNzMWOW</connection.sftpPassword>
</sqoop.connection>

<!--Job name, globally unique.-->
<sqoop.job name="Sftp.to.Hdfs" type="IMPORT" queue="default" priority="NORMAL">
<data.source connectionName="vt_sftp_test" connectionType="sftp-connector">
<file.inputPath>/opt/houjt/hive/all</file.inputPath>
<file.splitType>FILE</file.splitType>
<file.filterType>WILDCARD</file.filterType>
<file.pathFilter>*</file.pathFilter>
<file.fileFilter>*</file.fileFilter>
<file.encodeType>GBK</file.encodeType>
<file.suffixName></file.suffixName>
<file.isCompressive>FALSE</file.isCompressive>
</data.source>

<hadoop.source storageType="HDFS" >
<output.outputDirectory>/user/loader/sftp-to-hdfs</output.outputDirectory>
<output.fileOprType>OVERRIDE</output.fileOprType>
<throttling.extractors>3</throttling.extractors>
<output.fileType>TEXT_FILE</output.fileType>
</hadoop.source>

<sqoop.job.trans.file></sqoop.job.trans.file>
</sqoop.job>
</root>
```

 NOTE

Each Loader job needs to be associated with a connector. Connectors are used to read data from external data sources when data is imported to a cluster and used to write data into external data sources when data is exported from the cluster. In the preceding example, an SFTP data source connector is configured. To configure an SFTP and FTP data source connector, a password needs to be set and encrypted. The password encryption method is described as follows:

1. Run the following command to go to the **loader-tools-1.99.3** directory. For example, if the Loader client installation directory is **/opt/hadoopclient/Loader**, run the following command:

```
cd /opt/hadoopclient/Loader/loader-tools-1.99.3
```

2. Run the following command to encrypt the non-encrypted password:

```
./encrypt_tool Unencrypted password
```

- Step 4** Run the following command to go to the directory where loader-tool is located:

```
cd /opt/client/Loader/loader-tools-1.99.3/loader-tool
```

- Step 5** Run the following command to use the lt-ucc tool to create a connector:

```
./bin/lt-ucc -l /opt/client/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml -a create
```

If no error is reported and the following information is displayed, the connector creation task is submitted successfully:

```
User login success. begin to execute task.
```

- Step 6** Run the following command to use the lt-ucj tool to create a job:

```
./bin/lt-ucj -l /opt/client/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -w /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/sftp-to-hdfs.xml -a create
```

If no error is reported and the following information is displayed, the job creation task is submitted successfully:

```
User login success. begin to execute task.
```

- Step 7** Run the following command to use the lt-ctl tool to submit the job:

```
./bin/lt-ctl -l /opt/client/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -n Sftp.to.Hdfs -a start
```

If the following information is displayed, the job is submitted successfully:

```
Start job success.
```

- Step 8** Run the following command to view the job status:

```
./bin/lt-ctl -l /opt/client/Loader/loader-tools-1.99.3/loader-tool/job-config/login-info.xml -n Sftp.to.Hdfs -a status
```

```
Job:Sftp.to.Hdfs  
Status:RUNNING  
Progress: 0.0
```

----End

17.16.4 schedule-tool Usage Guide

Overview

schedule-tool is used to submit jobs of SFTP data sources. You can modify the input path and file filtering criteria before submitting a job. You can modify the output path if the target source is HDFS.

This section applies to MRS 3.x or later.

Parameters

Table 17-139 Configuration parameters of schedule.properties

Configuration parameters	Description	Example Value
server.url	Floating IP address and port for Loader. The default port is 21351. For compatibility, multiple IP addresses and ports can be configured and need to be separated by commas (,). The first IP address and port must be those of Loader. The others can be configured based on service requirements.	10.96.26.111:213 51,127.0.0.2:2135 1
authentication.type	Login authentication mode. <ul style="list-style-type: none"> • kerberos indicates that the security mode is used and Kerberos authentication is performed. Kerberos authentication provides two authentication modes: the password mode and the keytab file mode. • simple indicates that the normal mode is used and Kerberos authentication is not performed. 	kerberos
authentication.user	User for login when the normal mode or password authentication is used. In the keytab login mode, this parameter does not need to be set.	bar

Configuration parameters	Description	Example Value
authentication.password	<p>User password for login when the password authentication mode is used. In the normal mode or keytab login mode, this parameter does not need to be set.</p> <p>The password needs to be encrypted. The encryption method is described as follows:</p> <ol style="list-style-type: none"> Go to the directory where encrypt_tool is located. For example, if the Loader client installation directory is /opt/hadoopclient/Loader, run the following command: cd /opt/hadoopclient/Loader/loader-tools-1.99.3 Run the following command to encrypt the non-encrypted password: ./encrypt_tool Unencrypted password <p>The obtained encrypted password is used as the value of authentication.password.</p> <p>NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	-
use.keytab	<p>Whether to use the keytab mode to log in.</p> <ul style="list-style-type: none"> true indicates using the keytab file to log in. false indicates using the password to log in. 	true

Configuration parameters	Description	Example Value
client.principal	User principal for accessing the Loader service when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	loader/ hadoop. <i>System domain name</i> NOTE You can log in to FusionInsight Manager, choose System > Permission > Domain and Mutual Trust , and view the value of Local Domain , which is the current system domain name.
client.keytab	Directory where the used keytab file is located when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	/opt/client/conf/ loader.keytab
krb5.conf.file	Directory where the krb5.conf file is located when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	/opt/client/conf/ krb5.conf

Table 17-140 Configuration parameters of job.properties

Configuration parameters	Description	Example Value
job.jobName	Job name.	job1
file.fileName.prefix	File name prefix.	table1
file.fileName.posfix	File name suffix.	.txt

Configuration parameters	Description	Example Value
file.filter	File filter, which filters files by matching file names. <ul style="list-style-type: none"> • true indicates that the preceding prefix or suffix is used to match all files in the input path. For details, see the example. • false indicates that the preceding prefix or suffix is used to match a file in the input path. For details, see the example. 	true
date.day	Number of delayed days, which is matched with the date in the name of an imported file. For example, if the input date is 20160202 and the number of delayed days is 3, files that contain the 20160205 date field in the input path are matched. For details, see schedule-tool Usage Example .	3
file.date.format	Log format included in the name of the file to be imported.	yyyyMMdd
parameter.date.format	Entered date format when a script is invoked, which is usually consistent with file.date.format .	yyyyMMdd
file.format.iscompressed	Whether the file to be imported is a compressed file.	false
storage.type	Storage type. The final type of the file to be imported include HDFS, HBase, and Hive.	HDFS

 NOTE

schedule-tool supports the configuration of multiple jobs at the same time. When multiple jobs are configured at the same time, **job.jobName**, **file.fileName.prefix**, and **file.fileName.posfix** in [Table 17-140](#) need to be configured with multiple values, and the values need to be separated by **commas (,)**.

Precautions

server.url must be set to a format string of two IP addresses and port numbers, and the IP addresses and ports need to be separated by **commas (,)**.

17.16.5 schedule-tool Usage Example

Scenario

After a job is created using the Loader WebUI or Loader-tool, use schedule-tool to execute the job.

This section applies to MRS 3.x or later.

Prerequisites

The Loader client has been installed and configured. For details, see [Running a Loader Job by Using Commands](#).

Procedure

- Step 1** In the directory `/opt/houjt/test03` on the SFTP server, create multiple files with **table1** as the prefix, **.txt** as the suffix, and **yyyyMMdd** as the date format in the middle of the file name.

Figure 17-86 Example

```
[root@C12-RHEL64-ZYL111 test03]# ll
total 36
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160221.txt
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160222.txt
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160223.txt
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160224.txt
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160225.txt
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160226.txt
-rw-r--r--. 1 root root 54 Feb 29 18:43 table120160227.txt
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160228.txt
-rw-r--r--. 1 root root 54 Feb 29 19:11 table120160229.txt
```

- Step 2** Create a Loader job of importing data from the SFTP server to HDFS. For details, see [Typical Scenario: Importing Data from an SFTP Server to HDFS or OBS](#).
- Step 3** Log in to the node where the client is located as the user who installs the client.
- Step 4** Run the following command to go to the **conf** directory of schedule-tool. For example, if the Loader client installation directory is `/opt/client/Loader/`, run the following command:

```
cd /opt/client/Loader/loader-tools-1.99.3/schedule-tool/conf
```

- Step 5** Run the following command to edit the `schedule.properties` file and configure the login mode:

```
vi schedule.properties
```

schedule-tool supports two login modes. Only one mode can be selected. For parameter details, see [schedule-tool Usage Guide](#).

- When the password mode is used for login, the configuration information example is as follows:

```
[server.url = 10.10.26.187:21351,127.0.0.2:21351]
[authentication.type = kerberos]
```

```
[use.keytab = false]
[authentication.user = admin]
[authentication.password= xxx]
```

- When the keytab file mode is used for login, the configuration information example is as follows:

```
[server.url = 10.10.26.187:21351,127.0.0.2:21351]
[authentication.type = kerberos]
[use.keytab = true]
[client.principal = bar]
[client.keytab = /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config/user.keytab]
[krb5.conf.file = /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/hadoop-config/krb5.conf]
```

Step 6 Run the following command to edit the job.properties file and configure job information:

vi job.properties

```
#job name
job.jobName = sftp2hdfs-schedule-tool

#Whether to update the loader configuration parameters(File filter)?This parameter is used to match the
import file name.Values are true or false.
#false means update.the file name which is get by schedule tool will be updated to Loader configuration
parameters (File filter).
#false means no update.the file name which is get by schedule tool will be updated to Loader configuration
parameters (import path).
file.filter = false

#File name = prefix + date + suffix
#Need to import the file name prefix
file.fileName.prefix=table1

#Need to import the file name suffixes
file.fileName.posfix=.txt

#Date Days.Value is an integer.
#According to the date and number of days to get the date of the import file.
date.day = 1

#Date Format.Import file name contains the date format.Format TypeE°yyyyMMdd,yyyyMMdd
HHmmss,yyyy-MM-dd,yyyy-MM-dd HH:mm:ss
file.date.format = yyyyMMdd

#Date Format.Scheduling script execution. Enter the date format.
parameter.date.format = yyyyMMdd

#Whether the import file is a compressed format.Values ??are true or false.
#true indicates that the file is a compressed format?Execution scheduling tool will extract the files.false
indicates that the file is an uncompressed.Execution scheduling tool does not unpack.
file.format.iscompressed = false

#Hadoop storage type.Values are HDFS or HBase.
storage.type = HDFS
```

According to the data provided by [Step 1](#), the filtering rules are set as follows when the **table120160221.txt** file is used as an example:

- File name prefix:
file.fileName.prefix=table1
- File name suffix:
file.fileName.posfix=.txt

- Date format included in the file name:
file.date.format = yyyyMMdd
 - Entered date parameter for invoking the script:
parameter.date.format = yyyyMMdd
 - Number of delayed days.
date.day = 1
- For example, if the input date parameter of the script is **20160220**, the result is **20160221** by using the addition.

 NOTE

If the `./run.sh 20160220 /user/loader/schedule_01` command is executed, the preceding filtering rules will be combined into a string: `"table1"+"20160221".txt = table120160221.txt`.

Step 7 Select a filtering rule according to the value of **file.filter**.

- If a file is to be exactly matched, go to [Step 8](#).
- If a series of files are to be fuzzily matched, go to [Step 9](#).

Step 8 Change the value of **file.filter** in the **job.properties** file to **false**.

Run the following commands to run the job. The task is completed.

```
cd /opt/client/Loader/loader-tools-1.99.3/schedule-tool
```

```
./run.sh 20160220 /user/loader/schedule_01
```

20160220 indicates the input date, and */user/loader/schedule_01* indicates the output path.

 NOTE

The string `table120160221.txt` obtained by combining the preceding filtering rules will be used as the file name and appended to the input path of the job. Therefore, the job will only process the uniquely matched file `table120160221.txt`.

Step 9 In the **job.properties** file, change the value of **file.filter** to **true**, and set the value of **file.fileName.prefix** to *****.

Run the following commands to run the job. The task is completed.

```
cd /opt/client/Loader/loader-tools-1.99.3/schedule-tool
```

```
./run.sh 20160220 /user/loader/schedule_01
```

20160220 indicates the input date, and */user/loader/schedule_01* indicates the output path.

 NOTE

The string `*20160221.txt` obtained by combining the preceding filtering rules will be used as the fuzzy match mode of the file filter. In the input path of the job, all files matching `*20160221.txt` will be processed by the job.

----End

17.16.6 Using loader-backup to Back Up Job Data

Scenario

After a job is created using the Loader WebUI or loader-tool, use loader-backup to back up data.

NOTE

- Only Loader jobs of data export support data backup.
- This tool is an internal Loader interface and is invoked by the upper-layer component HBase. Only the data backup from HDFS to SFTP is supported.

This section applies to MRS 3.x or later.

Prerequisites

The Loader client has been installed and configured. For details, see [Running a Loader Job by Using Commands](#).

Procedure

Step 1 Log in to the node where the client is installed as the user who installs the client. For details, see [Running a Loader Job by Using Commands](#).

Step 2 Run the following command to go to the directory where the **backup.properties** file is located. For example, if the Loader client installation directory is **/opt/client/Loader/**, run the following command:

```
cd /opt/client/Loader/loader-tools-1.99.3/loader-backup/conf
```

Step 3 Run the following command to modify the configuration parameters of **backup.properties**. For details about the parameters, see [Table 17-141](#).

vi backup.properties

```
server.url = 10.0.0.1:21351,10.0.0.2:12000
authentication.type = kerberos
authentication.user =
authentication.password=
job.jobld = 1
use.keytab = true
client.principal = loader/hadoop
client.keytab = /opt/client/conf/loader.keytab
```

Table 17-141 Configuration parameters

Configuration parameters	Description	Example Value
server.url	<p>Floating IP address and port (21351) for Loader.</p> <p>For compatibility, multiple IP addresses and ports can be configured and need to be separated by commas (,). The first IP address and port must be those of Loader (21351). The others can be configured based on service requirements.</p>	10.0.0.1:21351,10.0.0.2:12000
authentication.type	<p>Login authentication mode.</p> <ul style="list-style-type: none"> • kerberos indicates that the security mode is used and Kerberos authentication is performed. Kerberos authentication provides two authentication modes: the password mode and the keytab file mode. • simple indicates that the normal mode is used and Kerberos authentication is not performed. 	kerberos
authentication.user	<p>User for login when the normal mode or password authentication is used.</p> <p>In the keytab login mode, this parameter does not need to be set.</p>	bar

Configuration parameters	Description	Example Value
<p>authentication.password</p>	<p>User password for login when the password authentication mode is used.</p> <p>In the normal mode or keytab login mode, this parameter does not need to be set.</p> <p>The password needs to be encrypted. The encryption method is described as follows:</p> <ol style="list-style-type: none"> 1. Go to the directory where encrypt_tool is located. For example, if the Loader client installation directory is /opt/hadoopclient/Loader, run the following command: cd /opt/hadoopclient/Loader/loader-tools-1.99.3 2. Run the following command to encrypt the non-encrypted password: <i>./encrypt_tool Unencrypted password</i> <p>The obtained encrypted password is used as the value of authentication.password.</p> <p>NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	<p>-</p>
<p>job.jobId</p>	<p>ID of the job whose data is to be backed up.</p> <p>Job IDs can be viewed under created jobs on the Loader web UI.</p>	<p>1</p>

Configuration parameters	Description	Example Value
use.keytab	Whether to use the keytab mode to log in. <ul style="list-style-type: none"> • true indicates using the keytab file to log in. • false indicates using the password to log in. 	true
client.principal	User principal for accessing the Loader service when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	loader/hadoop
client.keytab	Directory where the used keytab file is located when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	/opt/client/conf/loader.keytab

Step 4 Run the following command to go to the directory where the backup script **run.sh** is located. For example, if the Loader client installation directory is **/opt/hadoopclient/Loader**, run the following command:

```
cd /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-backup
```

Step 5 Run the following command to run the backup script **run.sh** to back up Loader job data. The system backs up data to a directory at the same layer of the job output directory.

```
./run.sh Backup data input directory
```

For example, the backup data input directory is **/user/hbase/**, and the job output directory is **/opt/client/sftp/sftp1**. **sftp1** acts as a placeholder. Run the following command to back up data to the **/opt/client/sftp/hbase** directory:

```
./run.sh /user/hbase/
```

```
----End
```

17.16.7 Open Source sqoop-shell Tool Usage Guide

Overview

This section applies to MRS 3.x or later.

Sqoop-shell is a shell tool of Loader. All its functions are implemented by executing the **sqoop2-shell** script.

The sqoop-shell tool provides the following functions:

- Creating and updating connectors
- Creating and updating jobs
- Deleting connectors and jobs
- Starting jobs in the synchronous or asynchronous mode.
- Stopping jobs
- Viewing job status
- Viewing historical execution records of jobs
- Cloning connectors and jobs
- Creating and updating conversion steps
- Specifying line and field separators

The sqoop-shell tool supports the following modes:

- Interaction mode
Users execute the **sqoop2-shell** script without parameters to go to the particular interaction window of Loader. After the contents of the script are input, the tool returns the relevant information to the interaction window.
- Batch mode
The **sqoop2-shell** script has a file name as a parameter and multiple commands are stored in lines in the file. The sqoop-shell tool runs all commands in the file in sequence by executing the script. Alternatively, users can execute the **sqoop2-shell** script, to the end of which a command is attached with the **-c** parameter as the bridge. In this case, the sqoop-shell tool runs one command each time.

The sqoop-shell implements functions of Loader by running the commands in [Table 17-142](#).

Table 17-142 Command list

Com man d	Description
exit	Exits the interaction mode. This command is supported only in the interaction mode.
histor y	Views the executed commands. This command is supported only in the interaction mode.
help	Views the tool help information.
set	Sets server attributes.
show	Displays service attributes and all the metadata information of Loader.
creat e	Creates connectors and jobs.
updat e	Updates connectors and jobs.

Com man d	Description
delet e	Deletes connectors and jobs.
clone	Clones connectors and jobs.
start	Starts jobs.
stop	Stops jobs.
status	Views job status.

Commands

- The sqoop2-shell tool provides two methods to obtain login authentication information. The first method is to obtain login authentication information from the configuration file. For details about the configuration items, see [Example for Using the Open-Source sqoop-shell Tool \(SFTP-HDFS\)](#) and [Example for Using the Open-Source sqoop-shell Tool \(Oracle-HBase\)](#). The second one is to obtain the authentication information by using parameters. Two modes are available in the second method: password mode and Kerberos authentication mode.

- Command for accessing the interaction mode

Execute the **sqoop2-shell** script without parameters to go to the sqoop tool window and run the commands one by one.

Run the following command to obtain the authentication information by reading the configuration file:

```
./sqoop2-shell
```

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal
```

The following information is displayed:

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
Sqoop Shell: Type 'help' or '\h' for help.

sqoop:000>
```

- Command for entering the batch mode

Two methods are available for accessing the batch mode.

1. Execute the **sqoop2-shell** script, in which a file name is used as a parameter and multiple commands are stored in lines in this file. The sqoop-shell tool runs all commands in the file in sequence. The script must be stored in the home directory of the current user, for example, **/root/batchCommand.sh**.

Run the following command to authenticate login by reading configuration files:

```
./sqoop2-shell /root/batchCommand.sh
```

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword /root/batchCommand.sh
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal /root/batchCommand.sh
```

batchCommand.sh is the user-defined name of the text file.

2. Execute the **sqoop2-shell** script, to the end of which a command is attached with the **-c** parameter as the bridge. The sqoop-shell tool will execute the command.

Run the following command to authenticate login by reading configuration files:

```
./sqoop2-shell -c expression
```

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword -c expression
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal -c expression
```

expression is the attached statement, whose format is the same as that in the text file in the first method.

- Exit command

This command is used for exiting the interaction mode and supported only in the interaction mode.

Example:

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
Sqoop Shell: Type 'help' or '\h' for help.

sqoop:000> exit
10-5-211-9:/opt/hadoopclient/Loader/loader-tools-1.99.3/sqoop-shell#
```

- History command

This command is used for viewing the executed commands and supported only in the interaction mode.

Example:

```
sqoop:000> history
0 show connector
1 create connection -c 4
2 show connections;
3 show connection;
4 show connection -a;
5 show connections;
6 show connection;
7 show connection -x 53;
8 show connection -x 52;
9 show connection -x 2
10 show connection -x 53;
11 show connection
12 show connection -x 53
```

```
13 create job -x 53 -t import
14 show connector
15 create connection -c 5
16 show connection -x 54
17 exit
18 show connector
19 create connection -c 5
20 exit
21 show connector
22 create connection -c 6
23 create job -x 20 -t import
24 start job -j 85 -s
25 \x
26 exit
27 history
sqoop:000>
```

- **Help command**

This command is used for viewing the tool help information.

Example:

```
sqoop:000> help
For information about Sqoop, visit: http://sqoop.apache.org/docs/1.99.3/index.html

Available commands:
exit (\x ) Exit the shell
history (\H ) Display, manage and recall edit-line history
help (\h ) Display this help message
set (\st ) Set server or option Info
show (\sh ) Show server, connector, framework, connection, job, submission or option Info
create (\cr ) Create connection or job Info
delete (\d ) Delete connection or job Info
update (\up ) Update connection or job Info
clone (\cl ) Clone connection or job Info
start (\sta) Start job
stop (\stp) Stop job
status (\stu) Status job

For help on a specific command type: help command

sqoop:000>
```

- **Set command**

The set command is used for setting attributes of clients and servers and supports the following attributes:

- **server** indicates setting the connection attributes for servers.

 **NOTE**

When attribute -u is set, attributes -h, -p, and -w can be ignored.

- **option** indicates setting the client attributes.

 **NOTE**

option can be set by key values. For example, **set option --name verbose --value true**.

Attribute Type	Subattribute	Description
server	-h,--host	Service IP address.
	-p,--port	Service Port

Attribute Type	Subattribute	Description
	-w,--webapp	Tomcat application name.
	-u,--url	Sqoop service URL.
option	verbose	Redundancy mode, which indicates that more information is printed.
	poll-timeout	Sets the polling timeout duration.

Example:

```
set option --name verbose --value false
set server --host 10.0.0.1 --port 21351 --webapp loader
```

- **show** command

This command is used for displaying information, such as variable information and storage metadata information.

Attribute Type	Subattribute	Description
server	-a,--all	Displays all server attributes.
	-p,--port	Displays the service port.
	-w,--webapp	Displays the Tomcat application name.
	-h,--host	Displays the service IP address.
option	-name	Displays the attributes of the specified name.
connector	-a,--all	Displays information about all connection types.
	-c,--cid	Displays information about the connection type of a specified ID.
framework	None.	Displays metadata information about frameworks.
connection	-a,--all	Displays all connection attributes.

Attribute Type	Subattribute	Description
	-x,--xid	Displays the attributes of a specified connection.
	-n,--name	Displays the connection attributes of a specified name.
job	-a,--all	Displays information about all jobs.
	-j,--jid	Displays job information about a specified ID.
	-n,--name	Displays job information about a specified name.
submission	-j,--jid	Displays the submission record of a specified job.
	-d,--detail	Displays details.

Example:

```
show server -all
show option --name verbose
show connector -all
show framework
show connection -all
show connection -n sftp-example
show job -all
show job -j 1
show submission --jid 1
show submission --jid 1 -d
```

- Create command

This command is used for creating connectors and jobs.

Attribute Type	Subattribute	Description
connection	-c,--cid	Specifies the ID of a connector type.
	-cn,--cname	Specifies the name of a specified connector type.
job	-x,--xid	Specifies the connector ID.
	-xn,--xname	Specifies the connector name.

Attribute Type	Subattribute	Description
	-t,--type	Specifies the job type. Possible values: <ul style="list-style-type: none"> import export

- In the interaction mode, enter the attribute values one by one as prompted.

Example for creating connectors:

```
create connection -c 1
create connection -cn example
```

Example for creating jobs:

```
create job -x 1 -t import
create job -xn job_example -t export
```

- In the batch mode, run the following command to view the specific attribute and then set a value for the attribute:

create job -t import -x 1 --help

You can run the above command in either of the following ways:

Save the command to a text file and attach this file to the end of the **sqoop-shell** script, and run the following command:

```
./sqoop2-shell batchCommand.sh
```

Attach a command with the **-c** parameter to the end of the **sqoop-shell** script and run the following command:

```
./sqoop2-shell -c expression
```

For details about command execution, refer to previous description in this section. The following shows two complete commands:

Example for creating connectors:

```
create connection -c 4 --connector-connection-sftpPassword xxxxx --connector-connection-sftpServerIp 10.0.0.1 --connector-connection-sftpServerPort 22 --connector-connection-sftpUser root--name testConnection
```

Example for creating jobs:

```
create job -t import -x 1 --connector-file-inputPath /opt/tempfile --connector-file-fileFilter * --framework-output-outputDirectory /user/loader/1 --framework-output-storageType HDFS --framework-throttling-extractorSize 120 --framework-output-fileType TEXT_FILE --connector-file-splitType FILE -queue default -priority low -name newJob
```

- In the batch mode, you can attach a statement using the **-c** parameter as the bridge.

Example for creating connectors:

```
./sqoop2-shell -c "create connection -c 4 --connector-connection-sftpPassword xxxxx --connector-connection-sftpServerIp 10.0.0.1 --connector-connection-sftpServerPort 22 --connector-connection-sftpUser root--name testConnection"
```

- **update** command

This command is used for updating connectors and jobs.

Attribute Type	Subattribute	Description
connection	-x,--xid	Specifies the connector ID. NOTE When the connectors are updated, the password must be set.
job	-j,--jid	Specifies the job ID.

- Interaction mode

Example for updating connectors:

```
update connection --xid 1
```

Example for updating jobs:

```
update job --jid 1
```

- Batch mode

Example for updating connectors:

```
update connection -x 6 --connector-connection-sftpServerPort 21 - --name sfp_130--connector-connection-sftpPassword xxxx
```

Example for updating jobs:

Example 1:

```
update job -jid 1 -name sftp2hdfs --connector-file-fileFilter *.txt
```

Example 2:

```
./sqoop2-shell -uk true -k /opt/loader/user.keytab -s user /opt/loader/testupdate.txt
./sqoop2-shell -uk true -k /opt/loader/user.keytab -s user -c "update job --jid 24 --name oracle-hive --connector-table-sql 'SELECT * FROM range_example WHERE replace(datadt,\'-\',\'.\')='20240801' and \${CONDITIONS}'"
```

 **NOTE**

When updating a job, you can write the update commands in a file, for example, **/opt/loader/testupdate.txt** (the file name can be customized), or specify the commands using **--connector-table-sql**, in which the **sql** command must be enclosed in single quotation marks ('). For details, see example 2. Involved commands include **connector-table-sql**, **connector-table-columns**, **connector-table-partitionColumn**, **connector-table-conditions**, **connector-table-queryCondition**.

• **delete** command

This command is used for deleting connectors and jobs.

Attribute Type	Subattribute	Description
connection	-x,--xid	Specifies the connector ID.
	-n,--name	Specifies the connector name.
job	-j,--jid	Specifies the job ID.
	-n,--name	Specifies the job name.

Example:

```
delete connection -x 1
delete connection --name abc
delete job -j 1
delete job -n qwerty
```

- **clone** command

This command is used for cloning connectors and jobs.

Attribute Type	Subattribute	Description
connection	-x,--xid	Specifies the connector ID. NOTE The password and connector name must be entered when the connectors are cloned.
job	-j,--jid	Specifies the job ID.

Example:

```
clone job -j 1
```

- **start** command

This command is used for starting jobs.

Attribute Type	Subattribute	Description
job	-j,--jid	Specifies the job ID.
	-n,--name	Specifies the job name.
	-s,--synchronous	Whether to start jobs in the synchronous mode or not.

Example for starting jobs in the asynchronous mode:

```
start job -j 1
start job -n abc
```

Example for starting jobs in the synchronous mode:

```
start job -j 1 -s
start job --name abc --synchronous
```

- **stop** command

This command is used for stopping jobs.

Attribute Type	Subattribute	Description
job	-j,--jid	Specifies the job ID.

Attribute Type	Subattribute	Description
	-n,--name	Specifies the job name.

Example:

```
stop job -j 1
stop job -n abc
```

- Status command

This command is used for viewing job status.

Attribute Type	Subattribute	Description
job	-j,--jid	Specifies the job ID.

When **-s** parameter is attached to the command, the result only contains the enumerated value of job status.

Example:

```
status job -j 1
status job -j 1 -s
```

Extended Attributes of Create Command

For the scenario in which HDFS exchanges data with the SFTP server or RDB, MRS extends the create command attributes on the basis of the open source sqoop-shell tool, so as to specify line and field separators and conversion steps when jobs are created.

Table 17-143 Extended Attributes of Create Command

Property	Description
fields-terminated-by	Default field separator.
lines-terminated-by	Default line separator.
input-fields-terminated-by	Inputs the step field separator. If the step field separator is not specified, the value equals to fields-terminated-by by default.
input-lines-terminated-by	Inputs the step line separator. If the step line separator is not specified, the value equals to lines-terminated-by by default.
output-fields-terminated-by	Outputs the step field separator. If the step field separator is not specified, the value equals to fields-terminated-by by default.

Property	Description
output-lines-terminated-by	Outputs the step line separator. If the step line separator is not specified, the value equals to lines-terminated-by by default.
trans	Specifies the conversion steps. The value is the directory where the conversion step file is located. When the relative directory of file is specified, the file is by default stored in the directory where the sqoop2-shell script is located. When the attribute is set, the other extended attributes can be ignored.

Interconnecting Sqoop1 with MRS

Step 1 Download the open source Sqoop from <http://www.apache.org/dyn/closer.lua/sqoop/1.4.7>.

Step 2 Save the downloaded **sqoop-1.4.7.bin__hadoop-2.6.0.tar.gz** package to the **/opt/sqoop** directory on the Master node in the MRS cluster and decompress the package.

```
tar zxvf sqoop-1.4.7.bin__hadoop-2.6.0.tar.gz
```

Step 3 Go to the directory where the package is decompressed and modify the configuration.

```
cd /opt/sqoop/sqoop-1.4.7.bin__hadoop-2.6.0/conf
```

```
cp sqoop-env-template.sh sqoop-env.sh
```

```
vi sqoop-env.sh
```

Add the following configurations:

```
export HADOOP_COMMON_HOME=/opt/client/HDFS/hadoop
```

```
export HADOOP_MAPRED_HOME=/opt/client/HDFS/hadoop
```

```
export HIVE_HOME=/opt/Bigdata/MRS_1.9.X/install/FusionInsight-Hive-3.1.0/hive  
(Enter the actual path.)
```

```
export HIVE_CONF_DIR=/opt/client/Hive/config
```

```
export HCAT_HOME=/opt/client/Hive/HCatalog
```

Step 4 Add the system variable **SQOOP_HOME** to **PATH**.

```
vi /etc/profile
```

Add the following information:

```
export SQOOP_HOME=/opt/sqoop/sqoop-1.4.7.bin__hadoop-2.6.0
```

```
export PATH=$PATH:$SQOOP_HOME/bin
```

Step 5 Run the following command to copy the **jline-2.12.jar** file to the **lib** file.

```
cp /opt/share/jline-2.12/jline-2.12.jar /opt/sqoop/  
sqoop-1.4.7.bin__hadoop-2.6.0/lib
```

Step 6 Run the following command to add the following configuration to the file.

```
vim $JAVA_HOME/jre/lib/security/java.policy  
permission javax.management.MBeanTrustPermission "register";
```

Step 7 Run the following command to interconnect sqoop1 with MRS.

```
source /etc/profile  
----End
```

17.16.8 Example for Using the Open-Source sqoop-shell Tool (SFTP-HDFS)

Scenario

Taking importing data from SFTP to HDFS as an example, this section introduces how to use the sqoop-shell tool to create and start Loader jobs in the interaction mode and batch mode.

This section applies to MRS 3.x or later.

Prerequisites

The Loader client has been installed and configured. For details, see [Running a Loader Job by Using Commands](#).

Example for the Interaction Mode

Step 1 Log in to the node where the Loader client is installed as the user who installs the client.

Step 2 Run the following command to go to the **conf** directory of the sqoop-shell tool. For example, if the Loader client installation directory is **/opt/client/Loader/**, run the following command:

```
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell/conf
```

Step 3 Run the following command to configure authentication information:

```
vi client.properties  
server.url=10.0.0.1:21351  
# simple or kerberos  
authentication.type=simple  
# true or false  
use.keytab=true  
  
authentication.user=  
authentication.password=  
  
client.principal=hdfs/hadoop@<system domain name>
```

```
# keytab file
client.keytab.file=./conf/login/hdfs.keytab
```

 **NOTE**

Log in to FusionInsight Manager and choose **System > Permission > Domain and Mutual Trust**. The value of **Local Domain** is the current system domain name.

Table 17-144 Configuration parameters

Configuration parameters	Description	Example Value
server.url	Floating IP address and port (21351) for Loader. For compatibility, multiple IP addresses and ports can be configured and need to be separated by commas (,). The first IP address and port must be those of Loader (21351). The others can be configured based on service requirements.	10.0.0.1:21351
authentication.type	Login authentication mode. <ul style="list-style-type: none"> • kerberos indicates that the security mode is used and Kerberos authentication is performed. Kerberos authentication provides two authentication modes: the password mode and the keytab file mode. • simple indicates that the normal mode is used and Kerberos authentication is not performed. 	kerberos
authentication.user	User for login when the normal mode or password authentication is used. In the keytab login mode, this parameter does not need to be set.	bar

Configuration parameters	Description	Example Value
authentication.password	<p>User password for login when the password authentication mode is used.</p> <p>In the normal mode or keytab login mode, this parameter does not need to be set.</p> <p>The password needs to be encrypted. The encryption method is described as follows:</p> <ol style="list-style-type: none"> 1. Go to the directory where encrypt_tool is located. For example, if the Loader client installation directory is /opt/hadoopclient/Loader, run the following command: cd /opt/hadoopclient/Loader/loader-tools-1.99.3 2. Run the following command to encrypt the non-encrypted password: <i>./encrypt_tool Unencrypted password</i> <p>The obtained encrypted password is used as the value of authentication.password.</p> <p>NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	-
use.keytab	<p>Whether to use the keytab mode to log in.</p> <ul style="list-style-type: none"> • true indicates using the keytab file to log in. • false indicates using the password to log in. 	true

Configuration parameters	Description	Example Value
client.principal	User principal for accessing the Loader service when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	loader/hadoop
client.keytab.file	Directory where the used keytab file is located when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	/opt/client/conf/loader.keytab

Step 4 Run the following command to go to the interaction mode:

```
source /opt/client/bigdata_env
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell
./sqoop2-shell
```

The preceding commands obtain authentication information by reading the configuration file.

Alternatively, you can also use the password or Kerberos authentication.

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal
```

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
Sqoop Shell: Type 'help' or '\h' for help.
sqoop:000>
```

Step 5 Run the following command to view the corresponding ID of the current connector:

```
show connector
```

The following information is displayed:

```
+-----+-----+-----+-----+
| Id | Name | Version | Class |
+-----+-----+-----+-----+
| 1 | generic-jdbc-connector | 2.0.6-SNAPSHOT | org.apache.sqoop.connector.jdbc.GenericJdbcConnector |
| 2 | ftp-connector | 2.0.5-SNAPSHOT | org.apache.sqoop.connector.ftp.FtpConnector |
| 3 | hdfs-connector | 2.0.5-SNAPSHOT |
```



```
Compression:
Output configuration
Storage type:
 0 : HDFS
 1 : HBASE_BULKLOAD
 2 : HBASE_PUTLIST
 3 : HIVE
Choose: 0
File type:
 0 : TEXT_FILE
 1 : SEQUENCE_FILE
 2 : BINARY_FILE
Choose: 0
Compression format:
 0 : NONE
 1 : DEFAULT
 2 : DEFLATE
 3 : GZIP
 4 : BZIP2
 5 : LZ4
 6 : SNAPPY
Choose:
Output directory: /user/loader/test
File operate type:
 0 : OVERRIDE
 1 : RENAME
 2 : APPEND
 3 : IGNORE
 4 : ERROR
Choose: 0

Throttling resources

Extractors: 2
Extractor size:
New job was successfully created with validation status FINE and persistent id 85
sqoop:000>
```

The preceding information indicates that the job ID is 85.

Step 8 Run the following command to start the job:

```
start job -j job ID -s
```

For example, if the job ID is 85, run the following command:

```
start job -j 85 -s
```

Displaying the **SUCCEEDED** information indicates that the job is started successfully.

```
Submission details
Job ID: 85
Server URL: https://10.0.0.0:21351/loader/
Created by: admin
Creation date: 2016-07-20 16:25:38 GMT+08:00
Lastly updated by: admin
2016-07-20 16:25:38 GMT+08:00: BOOTING - Progress is not available
2016-07-20 16:25:46 GMT+08:00: BOOTING - 0.00 %
2016-07-20 16:25:53 GMT+08:00: BOOTING - 0.00 %
2016-07-20 16:26:08 GMT+08:00: RUNNING - 90.00 %
2016-07-20 16:26:08 GMT+08:00: RUNNING - 90.00 %
2016-07-20 16:26:27 GMT+08:00: SUCCEEDED
```

----End

Example for the Batch Mode

Step 1 Log in to the node where the Loader client is installed as the user who installs the client.

Step 2 Run the following command to go to the **conf** directory of the sqoop-shell tool. For example, if the Loader client installation directory is **/opt/client/Loader/**, run the following command:

```
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell/conf
```

Step 3 Run the following command to configure authentication information:

```
vi client.properties
server.url=10.0.0.1:21351
# simple or kerberos
authentication.type=simple
# true or false
use.keytab=true

authentication.user=
authentication.password=

client.principal=hdfs/hadoop@<system domain name>

# keytab file
client.keytab.file=./conf/login/hdfs.keytab
```

Table 17-145 Configuration parameters

Configuration parameters	Description	Example Value
server.url	Floating IP address and port (21351) for Loader. For compatibility, multiple IP addresses and ports can be configured and need to be separated by commas (,). The first IP address and port must be those of Loader (21351). The others can be configured based on service requirements.	10.0.0.1:21351
authentication.type	Login authentication mode. <ul style="list-style-type: none"> kerberos indicates that the security mode is used and Kerberos authentication is performed. Kerberos authentication provides two authentication modes: the password mode and the keytab file mode. simple indicates that the normal mode is used and Kerberos authentication is not performed. 	kerberos

Configuration parameters	Description	Example Value
authentication.user	<p>User for login when the normal mode or password authentication is used.</p> <p>In the keytab login mode, this parameter does not need to be set.</p>	bar
authentication.password	<p>User password for login when the password authentication mode is used.</p> <p>In the normal mode or keytab login mode, this parameter does not need to be set.</p> <p>The password needs to be encrypted. The encryption method is described as follows:</p> <ol style="list-style-type: none"> 1. Go to the directory where encrypt_tool is located. For example, if the Loader client installation directory is /opt/hadoopclient/Loader, run the following command: cd /opt/hadoopclient/Loader/loader-tools-1.99.3 2. Run the following command to encrypt the non-encrypted password: ./encrypt_tool Unencrypted password <p>The obtained encrypted password is used as the value of authentication.password.</p> <p>NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	-

Configuration parameters	Description	Example Value
use.keytab	Whether to use the keytab mode to log in. <ul style="list-style-type: none"> • true indicates using the keytab file to log in. • false indicates using the password to log in. 	true
client.principal	User principal for accessing the Loader service when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	loader/hadoop
client.keytab.file	Directory where the used keytab file is located when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	/opt/client/conf/loader.keytab

Step 4 Run the following command to go to the directory where the **sqoop2-shell** script is located and create a text file in the directory, such as **batchCommand.sh**:

```
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell
```

```
vi batchCommand.sh
```

An example of **batchCommand.sh** is displayed as follows:

```
View parameters
create connection -c 6 --help

// Create a connector
create connection -c 6 -name sftp-connection --connector-connection-sftpServerIp 10.0.0.1 --connector-connection-sftpServerPort 22 --connector-connection-sftpUser root --connector-connection-sftpPassword
xxxxx

Create a job
create job -t import -x 20 --connector-file-inputPath /opt/tempfile --connector-file-fileFilter * --framework-output-outputDirectory /user/loader/1 --framework-output-storageType HDFS --framework-throttling-extractorSize 120 --framework-output-fileType TEXT_FILE --connector-file-splitType FILE -name test

Start a job
start job -j 85 -s
```

xxxxx is the password for the connector.

Step 5 Run the following command and the sqoop-shell tool will run the preceding commands in sequence:

```
./sqoop2-shell batchCommand.sh
```

The commands above authenticate login by reading configuration files. Alternatively, you can attach the authentication information to the command, that is, use the password mode or Kerberos mode to authenticate login.

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword batchCommand.sh
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal batchCommand.sh
```

Displaying the **SUCCEEDED** information indicates that the job is started successfully.

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
sqoop:000> create connection -c 6 --help
usage: Show connection parameters:
  --connector-connection-sftpPassword <arg>
  --connector-connection-sftpServerIp <arg>
  --connector-connection-sftpServerPort <arg>
  --connector-connection-sftpUser <arg>
  --framework-security-maxConnections <arg>
  --name <arg>
====> FINE
sqoop:000> create connection -c 6 -name sftp-connection --connector-connection-sftpServerIp 10.0.0.1 --
connector-connection-sftpServerPort 22 --connector-connection-sftpUser root --connector-connection-
sftpPassword xxxxx
Creating connection for connector with id 6
New connection was successfully created with validation status FINE and persistent id 20
====> FINE
sqoop:000> create job -t import -x 20 --connector-file-inputPath /opt/tempfile --connector-file-fileFilter * --
framework-output-outputDirectory /user/loader/1 --framework-output-storageType HDFS --framework-
throttling-extractorSize 120 --framework-output-fileType TEXT_FILE --connector-file-splitType FILE -name
test
Creating job for connection with id 20
New job was successfully created with validation status FINE and persistent id 85
====> FINE

Submission details
Job ID: 85
Server URL: https://10.0.0.0:21351/loader/
Created by: admin
Creation date: 2016-07-20 16:25:38 GMT+08:00
Lastly updated by: admin
2016-07-20 16:25:38 GMT+08:00: BOOTING - Progress is not available
2016-07-20 16:25:46 GMT+08:00: BOOTING - 0.00 %
2016-07-20 16:25:53 GMT+08:00: BOOTING - 0.00 %
2016-07-20 16:26:08 GMT+08:00: RUNNING - 90.00 %
2016-07-20 16:26:08 GMT+08:00: RUNNING - 90.00 %
2016-07-20 16:26:27 GMT+08:00: SUCCEEDED
```

Step 6 In the batch mode, the **-c** parameter can be used to attach a command. sqoop-shell can execute only the attached command at a time.

Run the following command to create a connection:

```
./sqoop2-shell -c "create connection -c 6 -name sftp-connection --connector-
connection-sftpServerIp 10.0.0.1 --connector-connection-sftpServerPort 22 --
connector-connection-sftpUser root --connector-connection-sftpPassword
xxxxx"
```

You can also use the password mode or Kerberos mode to attach the authentication information to the command.

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword -c "create connection -c 6 -name sftp-connection --connector-connection-sftpServerIp 10.0.0.1 --connector-connection-sftpServerPort 22 --connector-connection-sftpUser root --connector-connection-sftpPassword xxxxx"
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal -c "create connection -c 6 -name sftp-connection --connector-connection-sftpServerIp 10.0.0.1 --connector-connection-sftpServerPort 22 --connector-connection-sftpUser root --connector-connection-sftpPassword xxxxx"
```

Displaying the **FINE** information indicates the connection is created successfully.

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
sqoop:000> create connection -c 6 -name sftp-connection --connector-connection-sftpServerIp 10.0.0.1 --connector-connection-sftpServerPort 22 --connector-connection-sftpUser root --connector-connection-sftpPassword xxxxx
Creating connection for connector with id 6
New connection was successfully created with validation status FINE and persistent id 20
====> FINE
```

----End

17.16.9 Example for Using the Open-Source sqoop-shell Tool (Oracle-HBase)

Scenario

Taking **Importing Data from Oracle to HBase** as an example, this section introduces how to use the sqoop-shell tool to create and start Loader jobs in the interaction mode and batch mode.

This section applies to MRS 3.x or later.

Prerequisites

The Loader client has been installed and configured. For details, see [Running a Loader Job by Using Commands](#).

Example for the Interaction Mode

- Step 1** Log in to the node where the Loader client is installed as the user who installs the client.
- Step 2** Run the following command to go to the **conf** directory of the sqoop-shell tool. For example, if the Loader client installation directory is **/opt/client/Loader/**, run the following command:

```
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell/conf
```
- Step 3** Run the following command to configure authentication information:

```
vi client.properties
```

```
server.url=10.0.0.1:21351
# simple or kerberos
authentication.type=simple
# true or false
use.keytab=true

authentication.user=
authentication.password=

client.principal=oracle/hadoop@<system domain name>

# keytab file
client.keytab.file=./conf/login/oracle.keytab
```

 **NOTE**

Log in to FusionInsight Manager and choose **System > Permission > Domain and Mutual Trust**. The value of **Local Domain** is the current system domain name.

Table 17-146 Configuration parameters

Configuration parameters	Description	Example Value
server.url	Floating IP address and port (21351) for Loader. For compatibility, multiple IP addresses and ports can be configured and need to be separated by commas (,) . The first IP address and port must be those of Loader (21351). The others can be configured based on service requirements.	10.0.0.1:21351
authentication.type	Login authentication mode. <ul style="list-style-type: none"> • kerberos indicates that the security mode is used and Kerberos authentication is performed. Kerberos authentication provides two authentication modes: the password mode and the keytab file mode. • simple indicates that the normal mode is used and Kerberos authentication is not performed. 	kerberos
authentication.user	User for login when the normal mode or password authentication is used. In the keytab login mode, this parameter does not need to be set.	bar

Configuration parameters	Description	Example Value
authentication.password	<p>User password for login when the password authentication mode is used.</p> <p>In the normal mode or keytab login mode, this parameter does not need to be set.</p> <p>The password needs to be encrypted. The encryption method is described as follows:</p> <ol style="list-style-type: none"> 1. Go to the directory where encrypt_tool is located. For example, if the Loader client installation directory is /opt/hadoopclient/Loader, run the following command: cd /opt/hadoopclient/Loader/loader-tools-1.99.3 2. Run the following command to encrypt the non-encrypted password: <i>./encrypt_tool Unencrypted password</i> <p>The obtained encrypted password is used as the value of authentication.password.</p> <p>NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	-
use.keytab	<p>Whether to use the keytab mode to log in.</p> <ul style="list-style-type: none"> • true indicates using the keytab file to log in. • false indicates using the password to log in. 	true

Configuration parameters	Description	Example Value
client.principal	User principal for accessing the Loader service when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	loader/hadoop
client.keytab.file	Directory where the used keytab file is located when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	/opt/client/conf/loader.keytab

Step 4 Run the following command to go to the interaction mode:

```
source /opt/client/bigdata_env
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell
./sqoop2-shell
```

The preceding commands obtain authentication information by reading the configuration file.

Alternatively, you can also use the password or Kerberos authentication.

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal
```

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
Sqoop Shell: Type 'help' or '\h' for help.
sqoop:000>
```

Step 5 Run the following command to view the corresponding ID of the current connector:

```
show connector
```

The following information is displayed:

Id	Name	Version	Class
1	generic-jdbc-connector	2.0.7-SNAPSHOT	org.apache.sqoop.connector.jdbc.GenericJdbcConnector
2	ftp-connector	2.0.5-SNAPSHOT	org.apache.sqoop.connector.ftp.FtpConnector
3	hdfs-connector	2.0.5-SNAPSHOT	org.apache.sqoop.connector.hdfs.HdfsConnector

4 oracle-connector	2.0.1-SNAPSHOT	org.apache.sqoop.connector.oracle.OracleConnector
5 mysql-fastpath-connector	2.0.1-SNAPSHOT	org.apache.sqoop.connector.mysql.MySqlConnector
6 sftp-connector	2.0.6-SNAPSHOT	org.apache.sqoop.connector.sftp.SftpConnector
7 oracle-partition-connector	2.0.6-SNAPSHOT	org.apache.sqoop.connector.oracle.partition.OraclePartitionConnector

The preceding information indicates that the Oracle connector ID is 4.

Step 6 Run the following command to create connectors and enter the specific connector information as prompted:

create connection -c *connector ID*

For example, if the connector ID is 4, run the following command:

create connection -c 4

```
sqoop:000> create connection -c 4
Creating connection for connector with id 4
Please fill following values to create new connection object
Name: oracle14

Oracle connection configuration

JDBC connection string: jdbc:oracle:thin:@189.120.84.106:1521:orcl
Username: oracledba
Password: *****
JDBC connection properties:
There are currently 0 values in the map:
entry#
New connection was successfully created with validation status FINE and persistent id 3
sqoop:000>
```

The preceding information indicates that the connection ID is 3.

Step 7 Based on the connection ID, run the following command to create jobs:

create job -x *connection ID* -t import --trans *absolute path of job-config/oracle-hbase.json*

For example, if the connection ID is 3, run the following command:

create job -x 3 -t import --trans /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-config/oracle-hbase.json

The following information is displayed:

```
sqoop:000> create job -x 3 -t import --trans /opt/hadoopclient/Loader/loader-tools-1.99.3/loader-tool/job-
config/oracle-to-hbase.json
Creating job for connection with id 3
Please fill following values to create new job object
Name: run

Database target

Table name: test
Columns:
Conditions:
Data split method:
 0 : ROWID
 1 : PARTITION
Choose:
Table Partitions:
Data split allocation method:
```

```
0 : ROUNDROBIN
1 : SEQUENTIAL
2 : RANDOM
Choose:
JDBC fetch size:

Output configuration

Storage type:
0 : HDFS
1 : HBASE_BULKLOAD
2 : HBASE_PUTLIST
3 : HIVE
4 : SPARK
Choose: 1
HBase instance: HBase
Clear data before import : false

Throttling resources

Extractors: 10
Extractor size:
New job was successfully created with validation status FINE and persistent id 7
sqoop:000>
```

The preceding information indicates that the job ID is 7.

Step 8 Run the following command to start the job:

```
start job -j job ID -s
```

For example, if the job ID is 7, run the following command:

```
start job -j 7 -s
```

Displaying the **SUCCEEDED** information indicates that the job is started successfully.

```
Submission details
Job ID: 7
Server URL: https://10.0.0.0:21351/loader/
Created by: admintest
Creation date: 2019-12-04 16:37:34 CST
Lastly updated by: admintest
2019-12-04 16:37:34 CST: BOOTING - Progress is not available
2019-12-04 16:37:42 CST: BOOTING - 0.00 %
2019-12-04 16:37:42 CST: BOOTING - 0.00 %
2019-12-04 16:37:57 CST: RUNNING - 0.00 %
2019-12-04 16:38:12 CST: RUNNING - 45.00 %
2019-12-04 16:38:12 CST: RUNNING - 45.00 %
2019-12-04 16:38:27 CST: SUCCEEDED
```

----End

Example for the Batch Mode

Step 1 Log in to the node where the Loader client is installed as the user who installs the client.

Step 2 Run the following command to go to the **conf** directory of the sqoop-shell tool. For example, if the Loader client installation directory is **/opt/client/Loader/**, run the following command:

```
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell/conf
```

Step 3 Run the following command to configure authentication information:

vi client.properties

```
server.url=10.0.0.1:21351
# simple or kerberos
authentication.type=simple
# true or false
use.keytab=true

authentication.user=
authentication.password=

client.principal=hdfs/hadoop.<system domain name>@<system domain name>

# keytab file
client.keytab.file=./conf/login/hdfs.keytab
```

Table 17-147 Configuration parameters

Configuration parameters	Description	Example Value
server.url	Floating IP address and port (21351) for Loader. For compatibility, multiple IP addresses and ports can be configured and need to be separated by commas (,) . The first IP address and port must be those of Loader (21351). The others can be configured based on service requirements.	10.0.0.1:21351
authentication.type	Login authentication mode. <ul style="list-style-type: none"> • kerberos indicates that the security mode is used and Kerberos authentication is performed. Kerberos authentication provides two authentication modes: the password mode and the keytab file mode. • simple indicates that the normal mode is used and Kerberos authentication is not performed. 	kerberos
authentication.user	User for login when the normal mode or password authentication is used. In the keytab login mode, this parameter does not need to be set.	bar

Configuration parameters	Description	Example Value
authentication.password	<p>User password for login when the password authentication mode is used.</p> <p>In the normal mode or keytab login mode, this parameter does not need to be set.</p> <p>The password needs to be encrypted. The encryption method is described as follows:</p> <ol style="list-style-type: none"> 1. Go to the directory where encrypt_tool is located. For example, if the Loader client installation directory is /opt/hadoopclient/Loader, run the following command: cd /opt/hadoopclient/Loader/loader-tools-1.99.3 2. Run the following command to encrypt the non-encrypted password: <i>./encrypt_tool Unencrypted password</i> <p>The obtained encrypted password is used as the value of authentication.password.</p> <p>NOTE If a non-encrypted password contains special characters, the special characters must be escaped. For example, the dollar sign (\$) is a special character and can be escaped using single quotation marks ('). If a non-encrypted password contains single quotation marks, use double quotation marks to escape the single quotation marks. If a non-encrypted password contains double quotation marks, use backslashes (\) to escape the double quotation marks. For details, see the shell escape character rules.</p>	-
use.keytab	<p>Whether to use the keytab mode to log in.</p> <ul style="list-style-type: none"> • true indicates using the keytab file to log in. • false indicates using the password to log in. 	true

Configuration parameters	Description	Example Value
client.principal	User principal for accessing the Loader service when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	loader/hadoop
client.keytab.file	Directory where the used keytab file is located when the keytab authentication mode is used. In the normal mode or password login mode, this parameter does not need to be set.	/opt/client/conf/loader.keytab

Step 4 Run the following command to go to the directory where the **sqoop2-shell** script is located and create a text file in the directory, such as **batchCommand.sh**:

```
cd /opt/client/Loader/loader-tools-1.99.3/sqoop-shell
```

```
vi batchCommand.sh
```

An example of **batchCommand.sh** is displayed as follows:

```
View parameters
create connection -c 4 --help

// Create a connector
create connection -c 4 -name oracle-connection --connector-connection-oracleServerIp 10.0.0.1 --connector-connection-oracleServerPort 22 --connector-connection-oracleUser root --connector-connection-oraclePassword xxxxx

Create a job
create job -t import -x 3 --connector-file-inputPath /opt/tempfile --connector-file-fileFilter * --framework-output-directory /user/loader/1 --framework-output-storageType HBase --framework-throttling-extractorSize 120 --framework-output-fileType TEXT_FILE --connector-file-splitType FILE -name test

Start a job
start job -j 7 -s
```

xxxxx is the password for the connector.

Step 5 Run the following command and the sqoop-shell tool will run the preceding commands in sequence:

```
./sqoop2-shell batchCommand.sh
```

The commands above authenticate login by reading configuration files. Alternatively, you can attach the authentication information to the command, that is, use the password mode or Kerberos mode to authenticate login.

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword batchCommand.sh
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal batchCommand.sh
```

Displaying the **SUCCEEDED** information indicates that the job is started successfully.

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
sqoop:000> create connection -c 4 --help
usage: Show connection viparameters:
  --connector-connection-oraclePassword <arg>
  --connector-connection-oracleServerIp <arg>
  --connector-connection-oracleServerPort <arg>
  --connector-connection-oracleUser <arg>
  --framework-security-maxConnections <arg>
  --name <arg>
====> FINE
sqoop:000> create connection -c 4 -name oracle-connection --connector-connection-oracleServerIp 10.0.0.1
--connector-connection-oracleServerPort 22 --connector-connection-oracleUser root --connector-connection-
oraclePassword xxxxx
Creating connection for connector with id 4
New connection was successfully created with validation status FINE and persistent id 3
====> FINE
sqoop:000> create job -t import -x 3 --connector-file-inputPath /opt/tempfile --connector-file-fileFilter * --
framework-output-outputDirectory /user/loader/1 --framework-output-storageType HDFS --framework-
throttling-extractorSize 120 --framework-output-fileType TEXT_FILE --connector-file-splitType FILE -name
test
Creating job for connection with id 3
New job was successfully created with validation status FINE and persistent id 7
====> FINE
Submission details
Job ID: 7
Server URL: https://10.0.0.0:21351/loader/
Created by: admintest
Creation date: 2019-12-04 16:37:34 CST
Lastly updated by: admintest
2019-12-04 16:37:34 CST: BOOTING - Progress is not available
2019-12-04 16:37:42 CST: BOOTING - 0.00 %
2019-12-04 16:37:42 CST: BOOTING - 0.00 %
2019-12-04 16:37:57 CST: RUNNING - 0.00 %
2019-12-04 16:38:12 CST: RUNNING - 45.00 %
2019-12-04 16:38:12 CST: RUNNING - 45.00 %
2019-12-04 16:38:27 CST: SUCCEEDED
```

Step 6 In the batch mode, the **-c** parameter can be used to attach a command. sqoop-shell can execute only the attached command at a time.

Run the following command to create a connection:

```
./sqoop2-shell -c "create connection -c 4 -name oracle-connection --
connector-connection-oracleServerIp 10.0.0.1 --connector-connection-
oracleServerPort 22 --connector-connection-oracleUser root --connector-
connection-oraclePassword xxxxx"
```

You can also use the password mode or Kerberos mode to attach the authentication information to the command.

Run the following command to authenticate login using the password mode:

```
./sqoop2-shell -uk false -u username -p encryptedPassword -c "create
connection -c 4 -name oracle-connection --connector-connection-
oracleServerIp 10.0.0.1 --connector-connection-oracleServerPort 22 --
connector-connection-oracleUser root --connector-connection-oraclePassword
xxxxx"
```

Run the following command to authenticate login using the Kerberos mode:

```
./sqoop2-shell -uk true -k user.keytab -s userPrincipal -c "create connection -c 4
-name oracle-connection --connector-connection-oracleServerIp 10.0.0.1 --
```

```
connector-connection-oracleServerPort 22 --connector-connection-oracleUser
root --connector-connection-oraclePassword xxxxx"
```

Displaying the **FINE** information indicates the connection is created successfully.

```
Welcome to sqoop client
Use the username and password authentication mode
Authentication success.
sqoop:000> create connection -c 4 -name oracle-connection --connector-connection-oracleServerIp 10.0.0.1
--connector-connection-oracleServerPort 22 --connector-connection-oracleUser root --connector-connection-
oraclePassword xxxxx
Creating connection for connector with id 4
New connection was successfully created with validation status FINE and persistent id 3
====> FINE
```

----End

17.17 Loader Log Overview

Log Description

Log path: The default storage path of Loader log files is `/var/log/Bigdata/loader/Log category`.

- runlog: `/var/log/Bigdata/loader/runlog` (run logs)
- scriptlog: `/var/log/Bigdata/loader/scriptlog/` (script execution logs)
- catalina: `/var/log/Bigdata/loader/catalina` (Tomcat startup and stop logs)
- audit: `/var/log/Bigdata/loader/audit` (audit logs)

Log archive rule:

The automatic compression and archiving function are enabled for Loader run logs and audit logs. By default, when the size of a log file exceeds 10 MB, the log file is automatically compressed into a log file named in the following rule: `<Original log file name>-<yyyy-mm-dd_hh-mm-ss>.[ID].log.zip`. A maximum of 20 latest compressed files are reserved. The number of compressed files can be configured on the Manager portal.

Table 17-148 Loader log list

Log Type	Log File Name	Description
Run log	loader.log	Loader system log file that records most of the logs generated when the TelcoFS system is running.
	loader-omm-***-pid***-gc.log.*.current	Loader process GC log file
	sqoopInstanceCheck.log	Loader instance health check log file

Log Type	Log File Name	Description
Audit log	default.audit	Loader operation audit log file that records operations such as adding, deleting, modifying, and querying jobs and user login
Tomcat log	catalina.out	Tomcat run log file.
	catalina. <yyyy-mm-dd >.log	Tomcat run log file
	host-manager. <yyyy-mm-dd >.log	Tomcat run log file
	localhost_access_log. <yyyy-mm-dd >.txt	Tomcat run log file
	manager <yyyy-mm-dd >.log	Tomcat run log file
	localhost. <yyyy-mm-dd >.log	Tomcat run log file
Script log	postInstall.log	Loader installation script log file Log file generated during the execution of the Loader installation script (postInstall.sh)
	preStart.log	Pre-startup script log file of the Loader service During startup of the Loader service, a series of preparation operations are first performed (by executing preStart.sh), such as generating the keytab file. This log file records information about these operations
	loader_ctl.log	Log file generated when Loader executes the service start and stop script (sqoop.sh)

Log Level

Table 17-149 describes the log levels provided by Loader. The priorities of log levels are ERROR, WARN, INFO, and DEBUG in descending order. Logs whose levels are higher than or equal to the specified level are printed. The number of printed logs decreases as the specified log level increases.

Table 17-149 Log levels

Level	Description
ERROR	Error information about the current event processing.
WARN	Exception information about the current event processing.
INFO	Normal running status information about the system and events.
DEBUG	System information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of Loader by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Save the configuration. In the dialog box that is displayed, click **OK**. Then restart the service for the configuration to take effect.

----End

Log Formats

The following table lists the Loader log formats.

Table 17-150 Log formats

Log Type	Format	Example
Run log	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> <Thread that generates the log> <Message in the log> <Location of the log event>	2015-06-29 14:54:35,553 INFO [localhost-startStop-1] ConnectionRequestHandler initialized org.apache.sqoop.handler.ConnectionRequestHandler.<init>(ConnectionRequestHandler.java:100)

Log Type	Format	Example
Audit log	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level> default <Message in the log> <Location of the log event>	2015-06-29 15:35:40,969 INFO default: UserName=admin, UserIP=10.52.0.111, Time=2015-06-29 15:35:40,969, Operation=submit, Resource=submission@21, Result=Failure, Detail={ [reason:GET_SFTP_SESSION_FAILED:Failed to get sftp session - 10.162.0.35 (caused by: Auth cancel)]; [config:null]}

17.18 Example: Using Loader to Import Data from OBS to HDFS

Scenario

If you need to import a large volume of data from the external cluster to the internal cluster, import it from OBS to HDFS.

Prerequisites

- You have prepared service data.
- You have created an analysis cluster.

Procedure

Step 1 Upload service data to your OBS file system.

Step 2 Obtain the AK/SK information and create an OBS and HDFS link.

For details, see [Loader Link Configuration](#).

Step 3 Access the Loader page.

If Kerberos authentication is enabled in the analysis cluster, refer to instructions in [Accessing the Hue Web UI](#).

Step 4 Click **New Job**.

Step 5 In **Information**, set parameters.

1. In **Name**, enter a job name. For example, **obs2hdfs**.
2. In **From link**, select the OBS link you create.
3. In **To link**, select the HDFS link you create.

Step 6 In **From**, set source link parameters.

1. In **Bucket Name**, enter a name of the OBS file system.
2. In **Input directory or file**, enter a detailed location of service data in the file system.
If it is a single file, enter a complete path containing the file name. If it is a directory, enter the complete path of the directory.
3. In **File format**, enter the type of the service data file.

For details, see [obs-connector](#).

Step 7 In **To**, set destination link parameters.

1. In **Output directory**, enter the directory for storing service data in HDFS.
If Kerberos authentication is enabled in the cluster, the current user accessing Loader needs to have the permission to write data to the directory.
2. In **File format**, enter the type of the service data file.
The type must correspond to the type in [Step 6.3](#).
3. In **Compression codec**, enter a compression algorithm. For example, if you do not compress data, select **NONE**.
4. In **Overwrite**, select **True**.
5. Click **Show Senior Parameter** and set **Line Separator**.
6. Set **Field Separator**.

For details, see [hdfs-connector](#).

Step 8 In **Task Config**, set job running parameters.

1. In **Extractors**, enter the number of Map tasks.
2. In **Loaders**, enter the number of Reduce tasks.
If the destination link is an HDFS link, **Loaders** is hidden.
3. In **Max error records in single split**, enter an error record threshold.
4. In **Dirty data directory**, enter a directory for saving dirty data, for example, /**user/sqoop/obs2hdfs-dd**.

Step 9 Click **Save and execute**.

On the **Manage jobs** page, view the job running result. You can click **Refresh** to obtain the latest job status.

----End

17.19 Common Issues About Loader

17.19.1 Data Cannot Be Saved in Internet Explorer 10 or 11

Question

When Internet Explorer 10 or 11 is used to access the Loader page and submit data, an error is reported.

Answer

- Symptom
After data is saved and submitted, an error similar to "Invalid query parameter jobgroup id. cause: [jobgroup]" is reported.
- Cause
The POST requests are converted into GET requests after receiving the HTTP 307 response in some Internet Explorer 11 versions. As a result, POST data cannot be delivered to the server.
- Solution
Use Google Chrome.

17.19.2 Differences Among Connectors Used During the Process of Importing Data from the Oracle Database to HDFS

Question

Three types of connectors are available for importing data from the Oracle database to HDFS using Loader. That is, generic-jdbc-connector, oracle-connector, and oracle-partition-connector. Which one should I select? What are the differences between them?

Answers

- generic-jdbc-connector
Reads data from the Oracle database in JDBC mode. It is applicable to databases that support JDBC.
In this mode, data loading performance of Loader is subject to data distribution in a partition column. When data skew occurs (data has only one value or several values) in a partition column, a few Maps process a significant portion of data. As a result, the index becomes invalid, causing a sharp decline in SQL query performance.
generic-jdbc-connector supports view import and export, but **oracle-partition-connector** and **oracle-connector** do not support. Therefore, only this connector can be used to import views.
- Both **oracle-partition-connector** and **oracle-connector**
can use the ROWID of Oracle for partitioning. **oracle-partition-connector** is self-developed and **oracle-connector** is an open-source edition. The two types of connectors share similar performance.
oracle-connector requires more system table permissions. The following lists the read permissions required by the system tables of **oracle-connector** and **oracle-connector**.
 - **oracle-connector**: dba_tab_partitions, dba_constraints, dba_tables t, dba_segments, v\$instance, dba_objects, v\$instance, SYS_CONTEXT function, dba_extents, and dba_tab_subpartitions
 - **oracle-partition-connector**: DBA_OBJECTS and DBA_EXTENTSCompared with **generic-jdbc-connector**, **oracle-partition-connector** and **oracle-connector** have the following advantages:

- a. Load balancing: Number and scope of data segments are determined by the storage structure (data blocks) of the source table rather than the data on the source table. In terms of granularity, a data block can occupy a partition.
- b. Stable performance: Invalid index faults caused by data skew and bound variable snooping can be completely eliminated.
- c. Fast query speed: Using data segmentation delivers a higher query speed than that of using index.
- d. Excellent horizontal scalability: The number of generated segments increases with the increase of data volume. In this case, ideal performance can be delivered when you increase the number of concurrent tasks. Contrarily, decreasing concurrent tasks saves resources.
- e. Simplified data segmentation logic: Problems like precision loss, type compatibility, and bound variables can be prevented.
- f. Enhanced usability: Users do not need to create partition columns and tables for Loader.

18 Using Kudu

18.1 Using Kudu from Scratch

Kudu is a columnar storage manager developed for the Apache Hadoop platform. Kudu shares the common technical properties of Hadoop ecosystem applications. It is horizontally scalable and supports highly available operations.

Prerequisites

The cluster client has been installed. For example, the client is installed in the `/opt/hadoopclient` directory. The client directory in the following operations is only an example. Change it to the actual installation directory.

Procedure

Step 1 Log in to the node where the client is installed as the client installation user.

Run the `su - omm` command to switch to user `omm`.

Step 2 Run the following command to go to the client installation directory:

```
cd /opt/hadoopclient
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 Run the Kudu command line tool.

Run the command line tool of the Kudu component to view help information.

```
kudu -h
```

The command output is as follows:

```
Usage: kudu <command> [<args>]

<command> can be one of the following:
  cluster  Operate on a Kudu cluster
  diagnose Diagnostic tools for Kudu servers and clusters
  fs       Operate on a local Kudu filesystem
  hms      Operate on remote Hive Metastores
```

local_replica	Operate on local tablet replicas via the local filesystem
master	Operate on a Kudu Master
pbcc	Operate on PBC (protobuf container) files
perf	Measure the performance of a Kudu cluster
remote_replica	Operate on remote tablet replicas on a Kudu Tablet Server
table	Operate on Kudu tables
tablet	Operate on remote Kudu tablets
test	Various test actions
tserver	Operate on a Kudu Tablet Server
wal	Operate on WAL (write-ahead log) files

 **NOTE**

The Kudu command line tool does not support DDL and DML operations, but provides the refined query function for the **cluster**, **master**, **tserver**, **fs**, and **table** parameters.

Common operations:

- Check the tables in the current cluster.
kudu table list *KuduMaster instance IP1:7051, KuduMaster instance IP2:7051, KuduMaster instance IP3:7051*
- Query the configurations of the KuduMaster instance of the Kudu service.
kudu master get_flags *KuduMaster instance IP:7051*
- Query the schema of a table.
kudu table describe *KuduMaster instance IP1:7051, KuduMaster instance IP2:7051, KuduMaster instance IP3:7051 Table name*
- Delete a table.
kudu table delete *KuduMaster instance IP1:7051, KuduMaster instance IP2:7051, KuduMaster instance IP3:7051 Table name*

 **NOTE**

To obtain the IP address of the KuduMaster instance, choose **Components > Kudu > Instances** on the cluster details page.

----End

18.2 Accessing the Kudu Web UI

You can view Kudu job information on the Kudu web UI.


Prerequisites

Kudu has been installed in a cluster.

Accessing KuduMaster WebUI (MRS 3.x or Later)

- Step 1** Log in to Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
- Step 2** Choose **Cluster > Services > Spark**.
- Step 3** In the **Dashboard** page of Kudu, click **KuduMaster(KuduMaster)** on the right side of **KuduMaster WebUI**. The KuduMaster web UI is displayed.

Figure 18-1 KuduMaster WebUI



----End

Accessing KuduMaster WebUI (Versions Earlier Than MRS 3.x)

- Step 1** Access Manager. For details, see [Accessing MRS Manager \(Versions Earlier Than MRS 3.x\)](#).
- Step 2** choose **Services > Kudu**.
- Step 3** In **KuduMaster WebUI** of **Kudu Summary**, click **KuduMaster(KuduMaster)**. The KuduMaster web UI is displayed.

----End

19 Using MapReduce

19.1 Configuring the Log Archiving and Clearing Mechanism

Scenario

Job and task logs are generated during execution of a MapReduce application.

- Job logs are generated by the MRApplicationMaster, which record details about the start and running time of jobs and each task, Counter value, and other information. After being analyzed by HistoryServer, the job logs are used to view job execution details.
- A task log records the log information generated by each task running in a container. By default, task logs are stored only on the local disk of each NodeManager. After the log aggregation function is enabled, the NodeManager merges local task logs and writes them into HDFS after job execution completes.

The job logs and task logs of the MapReduce are stored on HDFS (when the log aggregation function is enabled). If the mechanism for periodically archiving and deleting log files is not configured for a cluster with a large number of computation tasks, the log files will occupy large memory space of HDFS and increase the cluster load.

Log archive is implemented by Hadoop Archives. The number (number of Map tasks) of concurrent archiving tasks started by the Hadoop Archives is related to the total size of log files to be archived. The formula is as follows: Number of concurrent archive tasks = Total size of log files to be archived/Size of archive files.

Configuration

Go to the **All Configurations** page of the MapReduce service. For details, see [Modifying Cluster Service Configuration Parameters](#).

Enter the parameter name in the search box, change the parameter value, and save the configuration. On the **Dashboard** tab page of the Mapreduce service,

choose **More > Synchronize Configuration**. After the synchronization is complete, restart the Mapreduce service.

- Job log parameters:

Table 19-1 Parameter description

Parameter	Description	Default Value
mapreduce.jobhistory.cleanner.enable	Whether to enable the job log file deletion function.	true
mapreduce.jobhistory.cleanner.interval-ms	Period for starting a log file cleanup. Only log files whose retention period is longer than the time specified by mapreduce.jobhistory.max-age-ms can be deleted.	86,400,000 ms (1 day)
mapreduce.jobhistory.max-age-ms	Log files whose retention period is longer than the retention period in milliseconds specified by this parameter will be deleted.	1,296,000,000 ms (15 days)

- Task log parameters:

Table 19-2 Parameter description

Parameter	Description	Default Value
yarn.log-aggregation.archive.files.minimum	Indicates the minimum number of archived MapReduce job log files. The archiving task starts when the number of files in the yarn.nodemanager.remote-app-log-dir folder is greater than or equal to the value of this parameter. This parameter applies to MRS 3.x.	5,000
yarn.log-aggregation.archive-check-interval-seconds	Indicates the MapReduce job log archiving interval, in seconds. Log files are archived only when the number of log files reaches the value of yarn.log-aggregation.archive.files.minimum . The archiving function is disabled when the period is set to 0 or -1 . This parameter applies to MRS 3.x.	-1

Parameter	Description	Default Value
yarn.log-aggregation.retain-seconds	Indicates the retention period on HDFS for archiving the MapReduce job logs. The value -1 indicates that log files are stored permanently.	1,296,000
yarn.log-aggregation.retain-check-interval-seconds	Indicates the check period (in seconds) of the MapReduce job log deletion task. If this parameter is set to -1 , the check period is one tenth of the log retention period.	86400

 NOTE

If task logs occupy too much HDFS storage space, modify the **mapreduce.jobhistory.max-age-ms** and **yarn.log-aggregation.retain-check-interval-seconds** configuration items to control the storage duration of task logs.

19.2 Reducing Client Application Failure Rate

Scenario

When the network is unstable or the cluster I/O and CPU are overloaded, client applications might encounter running failures.

Configuration

Adjust the following parameters in the **mapred-site.xml** configuration file on the client to reduce the client application failure rate:

 NOTE

The **mapred-site.xml** configuration file is in the **conf** directory of the client installation path, for example, **/opt/client/Yarn/config**.

Table 19-3 Parameter description

Parameter	Description	Default Value
mapreduce.reduce.shuffle.max-host-failures	Indicates the number of allowed failures of an MR task to read remote shuffle data in the Reduce process. When the number is set to be over 5, the client application failure rate can be reduced. This parameter applies to MRS 3.x.	5

Parameter	Description	Default Value
mapreduce.client.submit.file.replication	Indicates the backup of job files on HDFS. MR tasks are dependent on the job files during running. When the number of backups is set to be over 10, the client application failure rate can be reduced.	10

19.3 Transmitting MapReduce Tasks from Windows to Linux

Scenarios

If you want to transmit a job from Windows to Linux, set **mapreduce.app-submission.cross-platform** to **true**. If this parameter is unavailable for a cluster or its value is **false**, the function of transmitting MapReduce tasks from Windows to Linux is not supported. In this case, perform the following operations to add this parameter or change its value to enable this function:

 **NOTE**

This section applies to MRS 3.x or later.

Configuration Description

Adjust the following parameter in the **mapred-site.xml** configuration file on the client to enable the running of MapReduce tasks: The **mapred-site.xml** configuration file is in the **config** directory of the client installation path, for example, **/opt/client/Yarn/config**.

Table 19-4 Parameters

Parameter	Description	Default Value
mapreduce.app-submission.cross-platform	Indicates whether to support running of MapReduce tasks after they are transmitted from Windows to Linux. When the parameter value is true , the running of MapReduce tasks is supported. When the parameter value is false , the running of MapReduce tasks is not supported.	true

19.4 Configuring the Distributed Cache

Scenarios

 NOTE

This section applies to MRS 3.x or later.

Distributed caching is useful in the following scenarios:

Rolling Upgrade

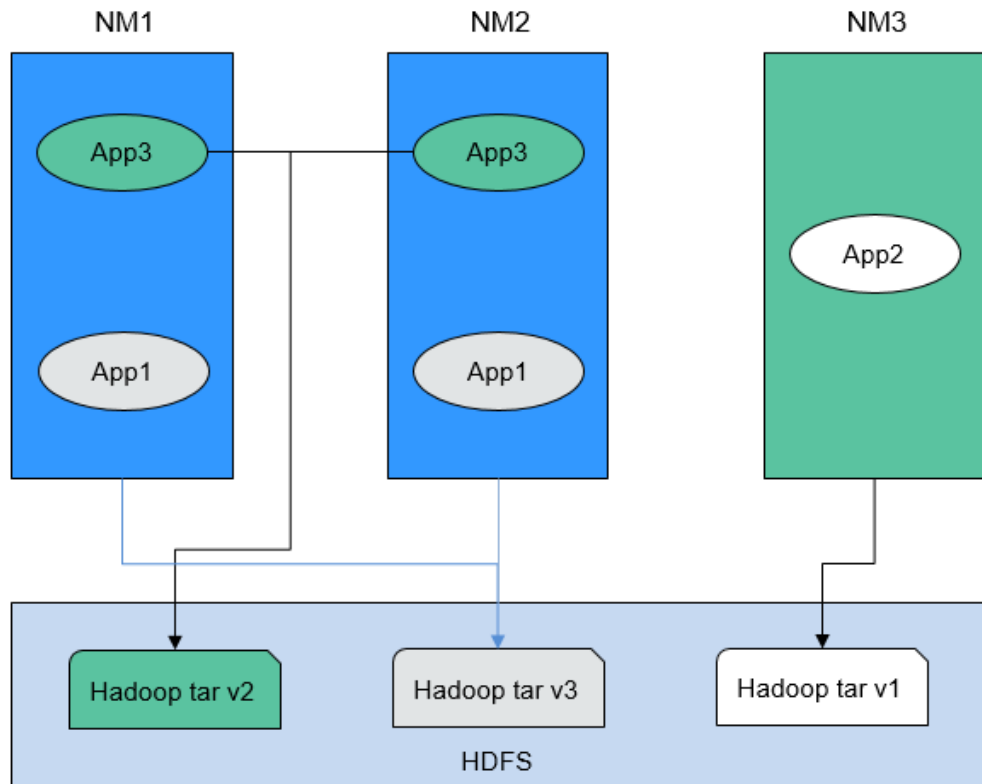
During the upgrade, applications must keep the text content (JAR file or configuration file) unchanged. The content is not based on Yarn of the current version, but on the version when it is submitted. This is a challenging issue. Generally, applications (such as MapReduce, Hive, and Tez) need to be installed locally. Libraries need to be installed on all cluster servers (clients and servers). When a rolling upgrade or downgrade starts in the cluster, the version of the locally installed library changes during application running. During the rolling upgrade, only a few NodeManagers are upgraded first. These NodeManagers obtain the software of the latest version. This leads to inconsistent behavior and can result in run-time errors.

Co-existence of Multiple Yarn Versions

Cluster administrators may run tasks that use multiple versions of Yarn and Hadoop JARs in a cluster. However, this task is difficult to be implemented because the JARs have been localized and have only one version.

The MapReduce application framework can be deployed through the distributed cache and does not depend on the static version copied during installation. Therefore, you can store multiple versions of Hadoop in HDFS and configure the **mapred-site.xml** file to specify the default version used by the task. You can run different versions of MapReduce by setting proper configuration attributes without using the versions deployed in the cluster.

Figure 19-1 Clusters with NodeManagers and Applications of multiple versions



As shown in [Figure 19-1](#), the application can use Hadoop JARs in HDFS instead of the local version. Therefore, during the rolling upgrade, even if NodeManager has been upgraded, the application can still run Hadoop of the earlier version.

Configuration Description

Step 1 Save the MapReduce **.tar** package of the specified version to a directory that can be accessed by applications in HDFS, as shown in the following command.

```
$HADOOP_HOME/bin/hdfs dfs -put hadoop-x.tar.gz /mapred/framework/
```

Step 2 Set parameters in the **mapred-site.xml** file based on [Table 19-5](#).

Table 19-5 Distributed cache parameters

Parameter	Description	Default Value
mapreduce.application.framework.path	Indicates the URL directing to the archive location. NOTE This property can also create an alias for the archive if the URL fragment identity name is specified as follows. In this example, the alias is set to mr-framework . <property> <name>mapreduce.application.framework.path</name> <value>hdfs:/mapred/framework/hadoop-x.tar.gz#mr-framework</value> </property>	NA

Parameter	Description	Default Value
mapreduce.application.classpath	<p>Indicates the parameter property, which contains the MapReduce JARs in the class directory.</p> <p>NOTE For example, the alias mr-framework used in the framework path is used to match the directory.</p> <pre><property> <name>mapreduce.application.classpath</name> <value>\$PWD/mr-framework/hadoop/share/hadoop/mapreduce/ *:\$PWD/mr-framework/hadoop/share/hadoop/mapreduce/lib/ *:\$PWD/mr-framework/hadoop/share/hadoop/common*:\$PWD/mr- framework/hadoop/share/hadoop/common/lib*:\$PWD/mr- framework/hadoop/share/hadoop/yarn*:\$PWD/mr-framework/ hadoop/share/hadoop/yarn/lib*:\$PWD/mr-framework/hadoop/share/ hadoop/hdfs*:\$PWD/mr-framework/hadoop/share/hadoop/ hdfs/lib*/etc/hadoop/conf/secure</value></property></pre>	N/A

You can upload MapReduce tarballs of multiple versions to HDFS. Different **mapred-site.xml** files indicate different locations. After that, you can run tasks for a specific **mapred-site.xml** file. The following is an example of running an MapReduce task for the MapReduce tarball of the *x* version:

```
hadoop jar share/hadoop/mapreduce/hadoop-mapreduce-examples-*.jar pi -conf
etc/hadoop-x/mapred-site.xml 10 10
```

----End

19.5 Configuring the MapReduce Shuffle Address

Scenario

When the MapReduce shuffle service is started, it attempts to bind an IP address based on local host. If the MapReduce shuffle service is required to connect to a specific IP address, no configuration is available. The following description allows you to configure a connection to a specific IP address.

Configuration

To bind a specific IP address to the MapReduce shuffle service, set the following parameters in the **mapred-site.xml** configuration file of the node where the NodeManager instance resides:

Table 19-6 Parameter description

Parameter	Description	Default Value
mapreduce.shuffle.address	<p>Indicates the specified address to run the shuffle service. The format is <i>IP:PORT</i>. The default value is empty. If this parameter is left empty, the local host IP address is bound. The default port number is 13562.</p> <p>NOTE If the value of <i>PORT</i> is different from that of mapreduce.shuffle.port, the mapreduce.shuffle.port value does not take effect.</p>	-

19.6 Configuring the Cluster Administrator List

Scenario

This function is used to specify the MapReduce cluster administrator.

The cluster administrator list is specified by **mapreduce.cluster.administrators**. The cluster administrator **admin** has all operation permissions.

Configuration

On the **All Configurations** page of the MapReduce service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).

Table 19-7 Parameter description

Parameter	Description	Default Value
mapreduce.cluster.acls.enabled	Indicates whether to enable permission control on Job History Server.	true
mapreduce.cluster.administrators	Indicates the administrator list of the MapReduce cluster. You can configure both users and user groups. Multiple users or user groups are separated by commas (,), and users and user groups are separated by spaces, for example, userA,userB groupA,groupB. The value * indicates all users or user groups.	<p>For versions earlier than MRS 3.x: mapred</p> <p>For MRS 3.x or later: mapred supergroup, System_administrator_186</p>

19.7 Introduction to MapReduce Logs

Log Description

Log paths:

- JobhistoryServer: `/var/log/Bigdata/mapreduce/jobhistory` (run log) and `/var/log/Bigdata/audit/mapreduce/jobhistory` (audit log)
- Container: `/srv/BigData/hadoop/data1/nm/containerlogs/application_{appid}/container_{$contid}`

NOTE

The logs of running tasks are stored in the preceding paths. After the running is complete, the system determines whether to aggregate the logs to an HDFS directory based on the YARN configuration. For details, see [Common YARN Parameters](#).

Log archive rule:

The automatic compression and archive function is enabled for MapReduce logs. By default, a log file is automatically compressed when the size of the log file is greater than 50 MB. The name of the compressed log file is in the following format: `<Name of the original log>-<yyyy-mm-dd_hh-mm-ss>.[NO.].log.zip`. A maximum of 100 latest compressed files are reserved. The number of compressed files can be configured on the parameter configuration page.

In MapReduce, JobhistoryServer cleans the old log files stored in HDFS periodically. The default storage directory is `/mr-history/done`. `mapreduce.jobhistory.max-age-ms` is used to set the cleanup interval. The default value of this parameter is 1,296,000,000 ms, which indicates 15 days.

Table 19-8 MapReduce log list

Type	Name	Description
Run log	jhs-daemon-start-stop.log	Startup log file of the daemon process
	hadoop-<SSH_USER>-jhshadaemon-<hostname>.log	Run log file of the daemon process
	hadoop-<SSH_USER>-<process_name>-<hostname>.out	Log that records the MapReduce running environment information
	historyserver-<SSH_USER>-<DATE>-<PID>-gc.log	Log that records the garbage collection of the MapReduce service
	jhs-haCheck.log	Log that records the active and standby status of MapReduce instances

Type	Name	Description
	yarn-start-stop.log	Log that records the startup and stop of the MapReduce service
	yarn-prestart.log	Log that records cluster operations before the MapReduce service startup
	yarn-postinstall.log	Work log before the MapReduce service startup and after the installation
	yarn-cleanup.log	Log that records the cleanup logs about the uninstallation of the MapReduce service
	mapred-service-check.log	Log that records the health check details of the MapReduce service
	container_{\$contid}	Container log
	hadoop-<SSH_USER>-<process_name>-<hostname>.log	MR run log
	mapred-switch-jhs.log	MR active/standby switchover log
	env.log	Environment information log before the instance is started or stopped
Audit log	mapred-audit-jobhistory.log	MapReduce operation audit log
	SecurityAuth.audit	MapReduce security audit log

Log Level

Table 19-9 describes the log levels supported by MapReduce. The log levels are FATAL, ERROR, WARN, INFO, and DEBUG from high priority to low. Logs whose levels are higher than or equal to the specified level are printed. The number of printed logs decreases as the specified log level increases.

Table 19-9 Log level

Level	Description
FATAL	Logs of this level record critical error information about the current event processing.
ERROR	Logs of this level record error information about the current event processing.
WARN	Logs of this level record unexpected alarm information about the current event processing.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of the MapReduce service. For details, see [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the left menu bar, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.

 **NOTE**

The configurations take effect immediately without restarting the service.

----End

Log Format

The following table lists the MapReduce log formats.

Table 19-10 Log format

Type	Format	Example
Run log	<i><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></i>	2020-01-26 14:18:59,109 INFO main Client environment:java.compiler=<N A> org.apache.zookeeper.Environment.logEnv(Environment.java:100)

Type	Format	Example
Audit log	<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>	2020-01-26 14:24:43,605 INFO main-EventThread USER=omm OPERATION=refreshAdminAcl s TARGET=AdminService RESULT=SUCCESS org.apache.hadoop.yarn.server. resourcemanager.RMAuditLog ger\$LogLevel \$6.printLog(RMAuditLogger.ja va:91)

19.8 MapReduce Performance Tuning

19.8.1 Optimization Configuration for Multiple CPU Cores

Scenario

Optimization can be performed when the number of CPU cores is large, for example, the number of CPU cores is three times the number of disks.

Procedure

You can set the following parameters in either of the following ways:

- Configuration on the server:
On the **All Configurations** page of the Yarn service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).
- Configuration on the client:
Modify the corresponding configuration file on the client.

NOTE

- Path of configuration files on the HDFS client: *Client installation directory*/HDFS/hadoop/etc/hadoop/hdfs-site.xml
- Path of configuration files on the Yarn client: *Client installation directory*/HDFS/hadoop/etc/hadoop/yarn-site.xml.
- Path of configuration files on the MapReduce client: *Client installation directory*/HDFS/hadoop/etc/hadoop/mapred-site.xml.

Table 19-11 Settings of multiple CPU cores

Conf igitration	Descriptio n	Parameter	Defa ult Valu e	Serv er/ Clie nt	Impact	Remarks
Num ber of slots in a node container	The combinatio n of the following parameter s determines the number of concurrent tasks (Map and Reduce tasks) of each node: <ul style="list-style-type: none"> • yarn.no demanager.resour ce.memory-mb • mapred uce.map.memory-mb • mapred uce.reduce.memory.m b 	yarn.nodemanager.resour ce.memory-mb NOTE For versions earlier than MRS 3.x, configure this parameter on the MRS console. For MRS 3.x or later: You need to configure this parameter on FusionInsight Manager.	Versi ons earlier than MRS 3.x: 8192 MRS 3.x or later: 16384	Server	If data needs to be read from and written into disks for all tasks (Map/Reduce tasks), a disk may be accessed by multiple processes at the same time, which leads to poor disk I/O performance. To ensure disk I/O performance, the number of concurrent access requests from a client to a disk cannot exceed 3.	The maximum number of concurrent containers must be [2.5 x Number of disks configured in Hadoop].
		mapreduce.map.memory.m b NOTE You need to set this parameter in the configuration file on the client in the <i>Client installation directory</i> /HDFS/hadoop/etc/hadoop/mapred-site.xml path.	4096	Client		

Conf igura tion	Descriptio n	Parameter	Defa ult Valu e	Serv er/ Clie nt	Impact	Remarks
		mapreduce.re duce.memory. mb NOTE You need to set this parameter in the configuration file on the client in the <i>Client</i> <i>installation</i> <i>directory/</i> HDFS/ hadoop/etc/ hadoop/ mapred- site.xml path.	4096	Clie nt		

Conf igure tion	Descriptio n	Parameter	Defa ult Valu e	Serv er/ Clie nt	Impact	Remarks
Map outp ut and com press ion	The Map task output before being written into disks can be compre ssed. This can save disk space, offer faster data write, and reduce the data traffic delivered to Reducer. You need to configure the following parameter s on the client: <ul style="list-style-type: none"> • mapred uce.ma p.outpu t.compr ess: The Map task output can be compre ssed before it is transmi tted over the network . It is a per-job 	mapreduce.m ap.output.co mpress NOTE You need to set this parameter in the configuration file on the client in the <i>Client installation directory/ HDFS/ hadoop/etc/ hadoop/ mapred- site.xml</i> path.	true	Clie nt	The disk I/O is the bottleneck. Therefore, use a compression algorithm with a high compression rate.	Snappy is used. The benchmar k test results show that Snappy delivers high performa nce and efficiency.
		mapreduce.m ap.output.co mpress.codec NOTE You need to set this parameter in the configuration file on the client in the <i>Client installation directory/ HDFS/ hadoop/etc/ hadoop/ mapred- site.xml</i> path.	org.a pach e.had oop.i o.co mpre ss.Lz4 Code c	Clie nt		

Conf igura tion	Descriptio n	Parameter	Defa ult Valu e	Serv er/ Clien t	Impact	Remarks
	configur ation. <ul style="list-style-type: none"> • mapred uce.ma p.outpu t.compr ess.cod ec: the codec used for data compre ssion 					
Spills	mapreduce .map.sort.s pill.percent	mapreduce.m ap.sort.spill.pe rcent NOTE You need to set this parameter in the configuration file on the client in the <i>Client installation directory/ HDFS/ hadoop/etc/ hadoop/ mapred- site.xml</i> path.	0.8	Clien t	Disk I/Os are the bottleneck. You can set the value of mapreduce.ta sk.io.sort.mb to minimize the memory spilled to the disk.	-

Conf iguretion	Descriptio n	Parameter	Defa ult Value	Serv er/ Client	Impact	Remarks
Data pack et size	When the HDFS client writes data to a data node, the data will be accumulated until a packet is generated. Then, the packet is transmitted over the network. dfs.client-write-packet-size specifies the data packet size. It can be specified by each job.	dfs.client-write-packet-size NOTE You need to set this parameter in the configuration file on the client in the <i>Client installation directory/HDFS/hadoop/etc/hadoop/hdfs-site.xml/</i> path.	262144	Client	The data node receives data packets from the HDFS client and writes data into disks through single threads. When disks are in the concurrent write state, increasing the data packet size can reduce the disk seek time and improve the I/O performance.	dfs.client-write-packet-size = 262144

19.8.2 Determining the Job Baseline

Scenario

The performance optimization effect is verified by comparing actual values with the baseline data. Therefore, determining optimal job baseline is critical to performance optimization.

When determining the job baseline, comply with the following rules:

- Making full use of cluster resources
- Setting the number of Map and Reduce tasks appropriately
- Setting the runtime of each task appropriately

Procedure

- **Rule 1: Making full use of cluster resources**

Enable all nodes to handle tasks as actively as they can when a job is executed. Maximizing the number of concurrent tasks helps make full use of resources. You can achieve this purpose by adjusting the data volume to be processed and the number of Map and Reduce tasks.

You can set **mapreduce.job.reduces** to control the number of Reduce tasks.

The number of Map tasks depends on the InputFormat type and whether the data file to be processed can be split. By default, TextFileInputFormat allocates Map tasks based on the number of blocks, that is, one Map task for each block. You can adjust the following parameters to improve resource utilization.

Parameter portal:

On the **All Configurations** page of the Yarn service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).

Parameter	Description	Default Value
mapreduce.input.fileinputformat.split.maxsize	Indicates the maximum size of the data block into which the Map input information is to be split. The shard size can be calculated based on its size customized by the user and the block size of each file. The formula is as follows: splitSize = Math.max(minSize, Math.min(maxSize, blockSize)) If maxSize is bigger than blockSize , a block is a shard. If maxSize is smaller than blockSize , a block will be split into multiple shards. If the size of the remaining data in a block is smaller than splitSize , the remaining data will be treated as a separated shard.	-
mapreduce.input.fileinputformat.split.minsize	Indicates the minimum size of a data shard.	0

- **Principle 2: Setting Reduce tasks to be executed in one round.**

Avoid the following scenarios:

- Most of Reduce tasks are completed in the first round, but there is still one Reduce task left running. The execution of the last Reduce task extends the runtime of the job. Therefore, reduce the number of Reduce tasks to enable all of them to run at the same time.

- All Map tasks are completed, but there are still Reduce tasks running on some nodes. In this case, the cluster resources are not fully utilized. You need to increase the number of Reduce tasks to enable each node to handle tasks.
- **Rule 3: Setting the runtime of each task appropriately**
If each Map or Reduce task of a job takes only a few seconds, most time of the job is wasted on scheduling tasks and starting and stopping processes. Therefore, you need to increase the data volume to be processed in each task. The preferred processing time for each task is 1 minute.

You can configure the following parameters to adjust the processing time in a task.

Parameter portal:

On the **All Configurations** page of the Yarn service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).


Parameter	Description	Default Value
mapreduce.input.fileinputformat.split.maxsize	Indicates the maximum size of the data block into which the Map input information is to be split. The shard size can be calculated based on its size customized by the user and the block size of each file. The formula is as follows: splitSize = Math.max(minSize, Math.min(maxSize, blockSize)) If maxSize is bigger than blockSize , a block is a shard. If maxSize is smaller than blockSize , a block will be split into multiple shards. If the size of the remaining data in a block is smaller than splitSize , the remaining data will be treated as a separated shard.	-
mapreduce.input.fileinputformat.split.minsize	Indicates the minimum size of a data shard.	0

19.8.3 Streamlining Shuffle

Scenario

During the shuffle procedure of MapReduce, the Map task writes intermediate data into disks, and the Reduce task copies and adds the data to the reduce function. Hadoop provides lots of parameters for the optimization.

Figure 19-2 Shuffle process



Procedure

1. Improving Performance in Map Phase

- Determine the memory used by Map.

To determine whether Map has sufficient memory, check the number of GCs and the ratio of the GC time over the total task time in counters of completed jobs. Normally, the GC time cannot exceed 10% of the task time (that is, GC time elapsed (ms)/CPU time spent (ms) < 10%).

You can improve Map performance by adjusting the following parameters.

Parameter portal:

On the **All Configurations** page of the Yarn service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).

Table 19-12 Parameter description

Parameter	Description	Default Value
mapreduce.map.memory.mb	Memory restriction of a Map task.	4096

Parameter	Description	Default Value
mapreduce.map.java.opts	JVM parameter of the Map subtask. If this parameter is set, it will replace the mapred.child.java.opts parameter. If -Xmx is not set, the value of Xmx is calculated based on mapreduce.map.memory.mb and mapreduce.job.heap.memory-mb.ratio .	<p>For versions earlier than MRS 3.x: -Xmx2048M -Djava.net.preferIPv4Stack=true</p> <p>For MRS cluster 3.x and later versions:</p> <ul style="list-style-type: none"> Clusters with Kerberos authentication enabled: -Djava.net.preferIPv4Stack=true -Djava.net.preferIPv6Addresses=false -Djava.security.krb5.conf=\${BIGDATA_HOME}/common/runtime/krb5.conf -Dbeetle.application.home.path=\${BIGDATA_HOME}/common/runtime/security/config Clusters with Kerberos authentication disabled: -Djava.net.preferIPv4Stack=true -Djava.net.preferIPv6Addresses=false -Dbeetle.application.home.path=\${BIGDATA_HOME}/common/runtime/security/config

It is recommended that the **-Xmx** in **mapreduce.map.java.opts** is 0.8 times the value of **mapreduce.map.memory.mb**.

- Using Combiner

Combiner is an optional procedure in the Map phase, in which the intermediate results with the same key value are combined. Generally, set

the reduce class to combiner. Combiner helps reduce the intermediate result output of Map, thereby consuming less network bandwidth during the shuffle process. You can use the following API to set a combiner class for a specific job.

Table 19-13 Combiner API

Class	API	Description
org.apache.hadoop.mapreduce.Job	public void setCombinerClass(Class<? extends Reducer> cls)	API used to set a combiner class for a specific job.

2. Improving Performance in Copy Phase

- Compress data.

Compress the intermediate output of Map. Data compression reduces the data to be transferred over the network. However, data compression and decompression consume more CPU. Determine whether to compress the intermediate results of Map based on site requirements. If a task is bandwidth-intensive, data compression improves processing performance. As for the bulkload optimization, compression of the intermediate output improves the performance by 60%.

To improve copy performance, set **mapreduce.map.output.compress** to **true** and **mapreduce.map.output.compress.codec** to **org.apache.hadoop.io.compress.SnappyCodec**.

3. Improving Performance in Merge Phase

To improve merge performance, configure the following parameters to reduce the number of times that Reduce writes data to disks.

Parameter portal:

On the **All Configurations** page of the Yarn service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).

Table 19-14 Parameter description

Parameter	Description	Default Value
mapreduce.reduce.merge.inmem.threshold	Threshold of the number of files for the in-memory merge process. When the accumulated number of files reaches the threshold, the process of in-memory merge and spilling to disks is initiated. If the value is less than or equal to 0 , the threshold does not take effect and the merge is triggered only based on the RAMFS memory usage.	1000
mapreduce.reduce.shuffle.merge.percent	Usage threshold for initiating in-memory merge, indicating the percentage of memory allocated to the Map outputs (defined by mapreduce.reduce.shuffle.input.buffer.percent).	0.66
mapreduce.reduce.shuffle.input.buffer.percent	Percentage of memory to be allocated from the maximum heap size to storing Map outputs during the Shuffle.	0.70
mapreduce.reduce.input.buffer.percent	Percentage of memory (relative to the maximum heap size) to retain Map outputs during the Reduce. When the Shuffle is completed, all remaining Map outputs in memory must use less than this threshold before the Reduce begins.	0.0

19.8.4 AM Optimization for Big Tasks

Scenario

A big job containing 100,000 Map tasks fails. It is found that the failure is triggered by the slow response of ApplicationMaster (AM).

When the number of tasks increases, the number of objects managed by the AM increases, which requires much more memory for management. The default memory heap for AM is 1 GB.

Procedure

You can improve the AM performance by setting the following parameters.

Navigation path for setting parameters:

Adjust the following parameters in the **mapred-site.xml** configuration file on the client to adjust the following parameters: The **mapred-site.xml** configuration file is in the **conf** directory of the client installation path, for example, **/opt/client/Yarn/config**.

Parameter	Description	Default Value
yarn.app.mapreduce.am.resource.mb	This parameter must be greater than the heap size specified by yarn.app.mapreduce.am.command-opts . Unit: MB	1536
yarn.app.mapreduce.am.command-opts	Indicates the JVM startup parameters loaded to MapReduce ApplicationMaster.	For versions earlier than MRS 3.x: -Xmx1024m -XX:CMSFullGCsBeforeCompaction=1 -XX:+UseConcMarkSweepGC -XX:+CMSParallelRemarkEnabled -XX:+UseCMSCompactAtFullCollection -verbose:gc MRS 3.x or later: -Xmx1024m -XX:+UseConcMarkSweepGC -XX:+CMSParallelRemarkEnabled -verbose:gc -Djava.security.krb5.conf=\${KRB5_CONFIG} -Dhadoop.home.dir=\${BIGDATA_HOME}/FusionInsight_HD_xxx/install/FusionInsight-Hadoop-xxx/hadoop

19.8.5 Speculative Execution

Scenario

If a cluster has hundreds or thousands of nodes, the hardware or software fault of a node may prolong the execution time of the entire task (as most tasks are already completed, the system is still waiting for the task running on the faulty node). Speculative execution allows a task to be executed on multiple machines. You can disable speculative execution for small clusters.

Procedure

Navigation path for setting parameters:

On the **All Configurations** page of the Yarn service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).

Parameter	Description	Default Value
mapreduce.map.speculative	Sets whether to execute multiple instances of some map tasks concurrently. true indicates that speculative execution is enabled.	false
mapreduce.reduce.speculative	Sets whether to execute multiple instances of some reduce tasks concurrently. true indicates that speculative execution is enabled.	false

19.8.6 Using Slow Start

Scenario

The Slow Start feature specifies the proportion of Map tasks to be completed before Reduce tasks are started. If the Reduce tasks are started too early, resources will be occupied, thereby reducing task running efficiency. However, if the Reduce tasks are started at an appropriate time, resource usage during shuffle and task running efficiency will be improved. For example, the MapReduce job includes 15 Map tasks and a cluster can start 10 Map tasks, there are 5 Map tasks remained after a round of Map tasks is completed and the cluster has available resources. In this case, you can configure the value of Slow Start to a value less than 1 (for example, 0.8), then the Reduce tasks can make use of the remaining cluster resources.

Procedure

Parameter portal:

On the **All Configurations** page of the MapReduce service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).

Parameter	Description	Default Value
mapreduce.job.reduce.slowstart.completedmaps	Fraction of the number of Maps in the job which should be completed before Reduces are scheduled for the job. By default, the Reduce tasks start when all the Map tasks are completed.	1.0

19.8.7 Optimizing Performance for Committing MR Jobs

Scenario

By default, if an MR job generates a large number of output files, it takes a long time for the job to commit the temporary outputs of a task to the final output

directory in the commit phase. In large clusters, the time-consuming commit process of jobs greatly affects the performance.

In this case, you can set the **mapreduce.fileoutputcommitter.algorithm.version** to **2** to improve the performance in the commit phase of MR jobs.

Procedure

Navigation path for setting parameters:

On the **All Configurations** page of the Yarn service, enter a parameter name in the search box. For details, see [Modifying Cluster Service Configuration Parameters](#).

Table 19-15 Parameter description

Parameter	Description	Default Value
mapreduce.fileoutputcommitter.algorithm.version	Indicates the algorithm version submitted by a job. The value is 1 or 2 . NOTE 2 is the recommended algorithm version. This algorithm enables tasks to directly commit the output results of each task to the final result output directory, reducing the time for the results of large jobs are committed.	2

19.9 Common Issues About MapReduce

19.9.1 Why Does a MapReduce Task Stay Unchanged for a Long Time?

Question

MapReduce job is not progressing for long time

Answer

This is because of less memory. When the memory is less, the time taken by the job to copy the map output increases significantly.

In order to reduce the waiting time, increase the heap memory.

The job configuration should be tuned according to number of mappers and data size processed by each mapper. Based on the input data size, tune the following configurations accordingly for feasible performance.

- `mapreduce.reduce.memory.mb`
- `mapreduce.reduce.java.opts`

Example: If the data size is 5 GB with 10 mappers, then the ideal heap memory would be 1.5 GB. Increase the heap memory size according with the increase in data size.

19.9.2 Why the Client Hangs During Job Running?

Question

Why is the client unavailable when the MR ApplicationMaster or ResourceManager is moved to the D state during job running?

Answer

When a task is running, the MR ApplicationMaster or ResourceManager is moved to D state (uninterrupted sleep state) or T state (stopped state). The client waits to return the task running state, but the MR ApplicationMaster does not return. Therefore, the client remains in the waiting state.

To avoid the preceding scenario, use the `ipc.client.rpc.timeout` configuration item in the `core-site.xml` file to set the client timeout interval.

The value of this parameter is millisecond. The default value is `0`, indicating that no timeout occurs. The client timeout interval ranges from 0 ms to 2,147,483,647 ms.

NOTE

- If the Hadoop process is in the D state, restart the node where the process is located.
- The `core-site.xml` configuration file is stored in the `conf` directory of the client installation path, for example, `/opt/client/Yarn/config`.

19.9.3 Why Cannot HDFS_DELEGATION_TOKEN Be Found in the Cache?

Question

In security mode, why delegation token HDFS_DELEGATION_TOKEN is not found in the cache?

Answer

In MapReduce, by default HDFS_DELEGATION_TOKEN will be canceled after the job completion. So if the token has to be re- used for the next job then the token will not be found in the cache.

To re-use the same token in subsequent job set the below parameter for the MR job configuration. When it is false the user can re-sue the same token.

```
jobConf.setBoolean("mapreduce.job.complete.cancel.delegation.tokens", false);
```

19.9.4 How Do I Set the Task Priority When Submitting a MapReduce Task?

Question

How do I set the job priority when submitting a MapReduce task?

Answer

You can add the parameter `-Dmapreduce.job.priority=<priority>` in the command to set task priority when submitting MapReduce tasks on the client. The format is as follows:

```
yarn jar <jar> [mainClass] -Dmapreduce.job.priority=<priority> [path1] [path2]
```

The parameters in the command are described as follows:

- `<jar>`: specifies the name of the JAR package to be run.
- `[mainClass]`: specifies the **main** method of the class for an application project in a JAR file.
- `<priority>`: specifies the priority of a task. The value can be **VERY_HIGH**, **HIGH**, **NORMAL**, **LOW**, or **VERY_LOW**.
- `[path1]`: specifies the data input path.
- `[path2]`: specifies the data output path.

For example, set the `/opt/client/HDFS/hadoop/share/hadoop/mapreduce/hadoop-mapreduce-examples*.jar` file to a high-priority task.

```
yarn jar /opt/client/HDFS/hadoop/share/hadoop/mapreduce/hadoop-mapreduce-examples*.jar wordcount -Dmapreduce.job.priority=VERY_HIGH /DATA.txt /out/
```

19.9.5 Why Physical Memory Overflow Occurs If a MapReduce Task Fails?

Question

The HBase bulkload task has 210,000 Map tasks and 10,000 Reduce tasks. The MapReduce task fails to be executed, and the physical memory of ApplicationMaster overflows.

```
For more detailed output, check the application tracking page:https://bigdata-55:8090/cluster/app/application\_1449841777199\_0003  
Then click on links to logs of each attempt.  
Diagnostics: Container [pid=21557,containerID=container_1449841777199_0003_02_000001] is running beyond physical memory limits  
Current usage: 1.0 GB of 1 GB physical memory used; 3.6 GB of 5 GB virtual memory used. Killing container.  
Dump of the process-tree for container_1449841777199_0003_02_000001 :  
|- PID PPID PGRP PID SESSID CMD_NAME USER_MODE TIME(MILLIS) SYSTEM_TIME(MILLIS)  
VMEM_USAGE(BYTES) RSSMEM_USAGE(PAGES) FULL_CMD_LINE  
|- 21584 21557 21557 21557 (java) 12342 1627 3871748096 271331 ${BIGDATA_HOME}/jdk1.8.0_51/bin/java  
-Djava.io.tmpdir=/srv/BigData/hadoop/data1/nm/localdir/usercache/hbase/appcache/application_1449841777199_0003/container_1449841777199_0003_02_000001/tmp -Dlog4j.configuration=container-log4j.properties  
-Dyarn.app.container.log.dir=/srv/BigData/hadoop/data1/nm/containerlogs/
```

```
application_1449841777199_0003/container_1449841777199_0003_02_000001 -  
Dyarn.app.container.log.filesize=0 -Dhadoop.root.logger=INFO,CLA  
-Dhadoop.root.logfile=syslog -Xmx784m org.apache.hadoop.mapreduce.v2.app.MRAppMaster  
|- 21557 21547 21557 21557 (bash) 0 0 13074432 368 /bin/bash -c ${BIGDATA_HOME}/jdk1.8.0_51//bin/  
java  
-Djava.io.tmpdir=/srv/BigData/hadoop/data1/nm/localdir/usercache/hbase/appcache/  
application_1449841777199_0003/container_1449841777199_0003_02_000001/tmp -  
Dlog4j.configuration=container-log4j.properties  
-Dyarn.app.container.log.dir=/srv/BigData/hadoop/data1/nm/containerlogs/  
application_1449841777199_0003/container_1449841777199_0003_02_000001 -  
Dyarn.app.container.log.filesize=0 -Dhadoop.root.logger=INFO,CLA  
-Dhadoop.root.logfile=syslog -Xmx784m org.apache.hadoop.mapreduce.v2.app.MRAppMaster 1>/srv/  
BigData/hadoop/data1/nm/containerlogs/application_1449841777199_0003/  
container_1449841777199_0003_02_000001/stdout  
2>/srv/BigData/hadoop/data1/nm/containerlogs/application_1449841777199_0003/  
container_1449841777199_0003_02_000001/stderr  
Container killed on request. Exit code is 143  
Container exited with a non-zero exit code 143  
Failing this attempt. Failing the application.
```

Answer

This is a performance specification problem. The root cause of the MapReduce task execution failure is the memory overflow of ApplicationMaster, that is, the NodeManager kills the task due to the physical memory overflow.

Solutions:

Increase the memory of ApplicationMaster and optimize the following parameters in the **mapred-site.xml** configuration file on the client:

- **yarn.app.mapreduce.am.resource.mb**
- **yarn.app.mapreduce.am.command-opts**. The recommended value of **-Xmx** is $0.8 \times \text{yarn.app.mapreduce.am.resource.mb}$.

Specification:

ApplicationMaster supports 24,000 concurrent containers when the configuration is as follows:

- **yarn.app.mapreduce.am.resource.mb=2048**
- In **yarn.app.mapreduce.am.command-opts**, **-Xmx** is **1638m**.

19.9.6 After the Address of MapReduce JobHistoryServer Is Changed, Why the Wrong Page is Displayed When I Click the Tracking URL on the ResourceManager WebUI?

Question

After the address of MapReduce JobHistoryServer is changed, why the wrong page is displayed when I click the tracking URL on the ResourceManager WebUI?

Answer

JobHistoryServer address (mapreduce.jobhistory.address / mapreduce.jobhistory.webapp.<https.>address) is the parameter of MapReduce. The MapReduce client will submit the address together with jobs to

ResourceManager. After ResourceManager completing the jobs, the parameter is saved in RMStateStore as the target address for viewing history job information.

If the JobHistoryServer address is changed, update the address in the configuration file of the MapReduce client in time. If the address is not updated, the page of earlier JobHistoryServer is displayed when you click the tracking URL of the new job. The target address of information about MapReduce jobs running before the change of address cannot be changed, so the wrong page is also displayed when you click the tracking URL. You can check the history information by accessing the new JobHistoryServer address.

19.9.7 MapReduce Job Failed in Multiple NameService Environment

Question

MapReduce or Yarn job fails in multiple nameService environment using viewFS.

Answer

When using viewFS only the mount directories are accessible, so the most possible cause is that the path configured is not in one of the mounted paths. For example:

```
<property>
<name>fs.defaultFS</name>
<value>viewfs://ClusterX/</value>
</property>
<property>
<name>fs.viewfs.mounttable.ClusterX.link./folder1</name>
<value>hdfs://NS1/folder1</value>
</property>
<property>
<name>fs.viewfs.mounttable.ClusterX.link./folder2</name>
<value>hdfs://NS2/folder2</value>
</property>
```

For all the MR properties which depends on HDFS, should use the paths inside mount folders.

Incorrect:

```
<property>
<name>yarn.app.mapreduce.am.staging-dir</name>
<value>/tmp/hadoop-yarn/staging</value>
</property>
```

As the root folder (/) is not accessible in viewFS.

Correct:

```
<property>
<name>yarn.app.mapreduce.am.staging-dir</name>
<value>/folder1/tmp/hadoop-yarn/staging</value>
</property>
```

19.9.8 Why a Fault MapReduce Node Is Not Blacklisted?

Question

MapReduce task fails and the ratio of fault nodes to all nodes is smaller than the blacklist threshold configured by `yarn.resourcemanager.am-scheduling.node-blacklisting-disable-threshold`. Why the fault node not be blacklisted?

Answer

If the blacklisted percentage exceeds the threshold, all blacklisted nodes are released. Traditionally, the blacklist percentage is the ratio of fault nodes to all nodes in the cluster. Currently, each node has a label expression. Therefore, the blacklist percentage needs to be calculated based on the number of nodes related to valid node label expressions. In other way, the blacklist percentage is the ratio of fault nodes related to valid node label expressions.

Assume that there are 100 nodes in the cluster, including 10 nodes (labelA) related to valid node label expressions. Assume that all nodes related to valid node label expressions are faulty and default blacklist threshold is 0.33. In traditional calculation method, $10/100 = 0.1$, which is far smaller than the threshold (0.33). In this case, the 10 nodes will never get released. Therefore, MapReduce always cannot obtain nodes and applications cannot run properly. In practice, the blacklist percentage needs to be calculated based on the total number of nodes related to valid node label expressions: $10/10 = 1$ is greater than the blacklist threshold and all nodes are released.

Therefore, even the ratio of fault nodes to all nodes in the cluster is below the threshold, all nodes in the blacklist are released.

20 Using OpenTSDB

20.1 Using an MRS Client to Operate OpenTSDB Metric Data

You can perform an interactive operation on an MRS cluster client. For a cluster with Kerberos authentication enabled, the user must belong to the **opentsdb**, **hbase**, **opentsdbgroup**, and **supergroup** groups and have the HBase permission.

Prerequisites

- The password of user **admin** has been obtained. The password of user **admin** is specified by the user during MRS cluster creation.
- The cluster client has been installed, for example, in the **/opt/client** directory. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite. For details about how to update a client, see [Updating a Client \(Versions Earlier Than 3.x\)](#).

Using a Client

Step 1 If Kerberos authentication is enabled for the current cluster, log in to MRS Manager and create a user that belongs to the **opentsdb**, **hbase**, **opentsdbgroup**, and **supergroup** groups and has the HBase permission, for example, **opentsdbuser**. If Kerberos authentication is disabled for the current cluster, skip this step.

Step 2 Prepare a client based on service conditions and log in to the node where the client is installed.

For example, if you have updated the client on the Master2 node, log in to the Master2 node to use the client. For details, see [Updating a Client \(Versions Earlier Than 3.x\)](#).

Step 3 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 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step.

- If the user is a human-machine user, run the **kinit opentsdbuser** command to authenticate the user.
- If the user is a machine-machine user, download the user authentication credential file, and save and decompress it to obtain the user's **user.keytab** and **krb5.conf** files. Go to the decompressed **user.keytab** directory, and run the **kinit -kt user.keytab opentsdbuser** command to authenticate the user.

Step 7 Operate the OpenTSDB data. For details, see [Operating Data](#).

----End

Operating Data

- Viewing help information

Run the **tsdb** command to print all commands supported by OpenTSDB, for example, **fsck**, **import**, **mkmetric**, **query**, **tsd**, **scan**, **search**, **uid**, and **version**.

Command output:

```
tsdb: error: unknown command "  
usage: tsdb <command> [args]  
Valid commands: fsck, import, mkmetric, query, tsd, scan, search, uid, version
```

- Creating an OpenTSDB metric

Run the **tsdb mkmetric** command to create a metric. For example, run the **tsdb mkmetric sys.cpu.user** command to create a metric named **sys.cpu.user**.

Command output:

```
Start run net.opentsdb.tools.UidManager, args: assign metrics sys.cpu.user  
metrics sys.cpu.user: [0, 0, 6]
```

- Importing data to the OpenTSDB metric

a. Prepare a metric file, for example, the **importData.txt** file that contains following information.

```
sys.cpu.user 1356998400 41 host=web01 cpu=0  
sys.cpu.user 1356998401 42 host=web01 cpu=0  
sys.cpu.user 1356998402 44 host=web01 cpu=0  
sys.cpu.user 1356998403 47 host=web01 cpu=0  
sys.cpu.user 1356998404 42 host=web01 cpu=0  
sys.cpu.user 1356998405 42 host=web01 cpu=0
```

b. Run the **tsdb import** command to import metric data. For example, run the **tsdb import importData.txt** command to import the **importData.txt** file.

```
Start run net.opentsdb.tools.TextImporter, args: importData.txt  
2019-06-26  
15:45:22,091 INFO [main] TextImporter:
```

```
reading from file:importData.txt
2019-06-26
15:45:22,102 INFO [main] TextImporter:
Processed importData.txt in 11 ms, 6 data points (545.5 points/s)
2019-06-26
15:45:22,102 INFO [main] TextImporter:
Total: imported 6 data points in 0.012s (504.0 points/s)
```

- Querying the OpenTSDB metric

Run the **tsdb uid metrics** command to obtain the metric stored in OpenTSDB. For example, run the **tsdb uid metrics sys.cpu.user** command to query the data of the **sys.cpu.user** metric.

Command output:

```
Start run net.opentsdb.tools.UidManager, args: metrics sys.cpu.user
metrics sys.cpu.user: [0, 0, 6]
```

To obtain more information, run the **tsdb uid** command.

```
Start run net.opentsdb.tools.UidManager, args:
Not enough arguments
Usage: uid <subcommand> args
Sub commands:
grep [kind] <RE>: Finds matching IDs.
assign <kind> <name> [names]: Assign an ID for the given name(s).
rename <kind> <name> <newname>: Renames this UID.
delete <kind> <name>: Deletes this UID.
fsck: [fix] [delete_unknown] Checks the consistency of UIDs.
    fix          - Fix errors. By default errors are logged.
    delete_unknown - Remove columns with unknown qualifiers.
                    The "fix" flag must be supplied as well.
[kind] <name>: Lookup the ID of this name.
[kind] <ID>: Lookup the name of this ID.
metasync: Generates missing TSUID and UID meta entries, updates created timestamps
metapurge: Removes meta data entries from the UID table
treesync: Process all timeseries meta objects through tree rules
treepurge <id> [definition]: Purge a tree and/or the branches from storage. Provide an integer Tree
ID and optionally add "true" to delete the tree definition
Example values for [kind]: metrics, tagk (tag name), tagv (tag value).
--config=PATH Path to a configuration file (default: Searches for file see docs).
--idwidth=N Number of bytes on which the Uniqueid is encoded.
--ignore-case Ignore case distinctions when matching a regexp.
--table=TABLE Name of the HBase table where to store the time series (default: tsdb).
--uidtable=TABLE Name of the HBase table to use for Unique IDs (default: tsdb-uid).
--verbose Print more logging messages and not just errors.
--zkbasedir=PATH Path under which is the znode for the -ROOT- region (default: /hbase).
--zkquorum=SPEC Specification of the ZooKeeper quorum to use (default: localhost).
-i Short for --ignore-case.
-v Short for --verbose.
```

- Scanning the OpenTSDB metric data

Run the **tsdb query** command to query the imported metric data in batches. The command format is as follows: **tsdb query <START-DATE> <END-DATE> <aggregator> <metric> <tagk=tagv>**. For example, run the **tsdb query 0 1h-ago sum sys.cpu.user host=web01** command.

```
Start run net.opentsdb.tools.CliQuery, args: 0 1h-ago sum sys.cpu.user host=web01
sys.cpu.user 1356998400000 41 {host=web01, cpu=0}
sys.cpu.user 1356998401000 42 {host=web01, cpu=0}
sys.cpu.user 1356998402000 44 {host=web01, cpu=0}
sys.cpu.user 1356998403000 47 {host=web01, cpu=0}
sys.cpu.user 1356998404000 42 {host=web01, cpu=0}
sys.cpu.user 1356998405000 42 {host=web01, cpu=0}
```

 NOTE

- <START-DATE>: start time of the metric to be queried
 - <END-DATE>: end time of the metric to be queried
 - <aggregator>: aggregation mode of the data query
 - <metric>: name of the metric to be queried
 - <tagk=tagv>: key and value of a tag
- Deleting the imported OpenTSDB metric
Run the **tsdb uid delete** command to delete the imported metric and its value. For example, to delete the **sys.cpu.user** metric, run the **tsdb uid delete metrics sys.cpu.user** command.

```
Start run net.opentsdb.tools.UidManager, args: delete metrics sys.cpu.user
```

20.2 Running the curl Command to Operate OpenTSDB

Writing Data

For example, to write data of a metric named **testdata**, whose timestamp is **1524900185**, value is **true**, tag is **key** and **value**, run the following command:

```
curl -ki -X POST -d '{"metric":"testdata", "timestamp":1524900185, "value":"true", "tags": {"key":"value"}}' https://<tsd_ip>:4242/api/put?sync
```

<tsd_ip>: indicates the IP address of the TSD instance of OpenTSDB to which data is to be written.

```
HTTP/1.1 204 No Content  
Content-Type: application/json; charset=UTF-8  
Content-Length:0
```

Querying Data

For example, to query summary information about the **testdata** metric in the past three years, run the following command:

```
curl -ks https://<tsd_ip>:4242/api/query?start=3y-ago&m=sum:testdata | python -m json.tool
```

- <tsd_ip>: indicates the IP address or host name of the TSD instance of OpenTSDB that needs to be accessed.
- <start=3y-ago&m=sum:testdata>: Translates the **&** symbol, which may not be identified in the request.
- (Optional) <python -m json.tool>: Converts the response request to the JSON format.

```
[  
  {  
    "aggregateTags": [],  
    "dps": {  
      "1524900185": 1  
    },  
    "metric": "testdata",  
    "tags": {  
      "key": "value"  
    }  
  }  
]
```

Querying tsd Status

For example, to query information about the client connected to HBase, run the following command:

```
curl -ks https://<tsd_ip>:4242/api/stats/region_clients | python -m json.tool
```

<tsd_ip>: indicates the IP address of the TSD instance of OpenTSDB that needs to be accessed.

```
[
  {
    "dead": false,
    "endpoint": "/xx.xx.xx.xx:16020",
    "inflightBreached": 0,
    "pendingBatchedRPCs": 0,
    "pendingBreached": 0,
    "pendingRPCs": 0,
    "rpcResponsesTimedout": 0,
    "rpcResponsesUnknown": 0,
    "rpcid": 78,
    "rpcsInFlight": 0,
    "rpcsSent": 79,
    "rpcsTimedout": 0,
    "writesBlocked": 0
  }
]
```

21 Using Oozie

21.1 Using Oozie from Scratch

Oozie is an open-source workflow engine that is used to schedule and coordinate Hadoop jobs.

Oozie can be used to submit a wide array of jobs, such as Hive, Spark2x, Loader, MapReduce, Java, DistCp, Shell, HDFS, SSH, SubWorkflow, Streaming, and scheduled jobs.

This section describes how to use the Oozie client to submit a MapReduce job.

Prerequisites

The client has been installed in a directory, for example, **/opt/client**. The client directory in the following operations is only an example. Change it based on site requirements.

Procedure

Step 1 Log in to the node where the client is installed as the client installation user.

Step 2 Run the following command to go to the client installation directory, for example, **/opt/client**:

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 Check the cluster authentication mode.

- If the cluster is in security mode, run the following command to authenticate the user: *UserOozie* indicates the user who submits tasks.

```
kinit UserOozie
```

- If the cluster is in normal mode, go to [Step 5](#).

Step 5 Upload the Oozie configuration file and JAR package to HDFS.

```
hdfs dfs -mkdir /user/UserOozie
```

```
hdfs dfs -put -f /opt/client/Oozie/oozie-client-*/examples /user/UserOozie/
```

 NOTE

- /opt/client is the client installation directory. Change it based on site requirements.
- **UserOozie** indicates the name of the user who submits jobs.

Step 6 Run the following commands to modify the job execution configuration file:

```
cd /opt/client/Oozie/oozie-client-*/examples/apps/map-reduce/
```

```
vi job.properties
```

```
nameNode=hdfs://hacluster
resourceManager=10.64.35.161:8032 (10.64.35.161 is the service plane IP address of the Yarn
resourceManager (active) node, and 8032 is the port number of yarn.resourcemanager.port)
queueName=default
examplesRoot=examples
user.name=admin
oozie.wf.application.path=${nameNode}/user/${user.name}/${examplesRoot}/apps/map-reduce#
HDFS upload path
outputDir=map-reduce
oozie.wf.rerun.failnodes=true
```

Step 7 Run the following command to execute the Oozie job:

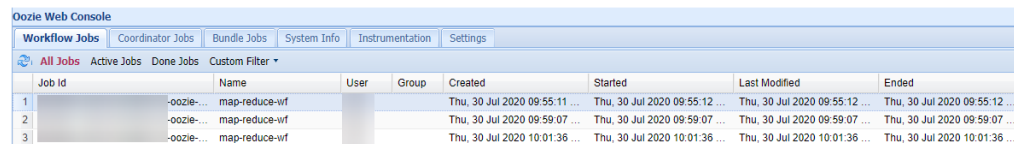
```
oozie job -oozie https://Host name of the Oozie role:21003/oozie/ -config
job.properties -run
```

```
[root@kwephispra44947 map-reduce]# oozie job -oozie https://kwephispra44948:21003/oozie/ -config
job.properties -run
.....
job: 0000000-200730163829770-oozie-omm-W
```

Step 8 Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).

Step 9 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie**, click the hyperlink next to **Oozie WebUI** to go to the Oozie page, and view the task execution result on the Oozie web UI.

Figure 21-1 Task execution result



Job Id	Name	User	Group	Created	Started	Last Modified	Ended
1	-oozie-... map-reduce-wf			Thu, 30 Jul 2020 09:55:11 ...	Thu, 30 Jul 2020 09:55:12 ...	Thu, 30 Jul 2020 09:55:12 ...	Thu, 30 Jul 2020 09:55:12 ...
2	-oozie-... map-reduce-wf			Thu, 30 Jul 2020 09:59:07 ...	Thu, 30 Jul 2020 09:59:07 ...	Thu, 30 Jul 2020 09:59:07 ...	Thu, 30 Jul 2020 09:59:07 ...
3	-oozie-... map-reduce-wf			Thu, 30 Jul 2020 10:01:36 ...	Thu, 30 Jul 2020 10:01:36 ...	Thu, 30 Jul 2020 10:01:36 ...	Thu, 30 Jul 2020 10:01:36 ...

----End

21.2 Using the Oozie Client

Scenario

This section describes how to use the Oozie client in an O&M scenario or service scenario.

Prerequisites

- The client has been installed in a directory, for example, **/opt/client**. The client directory in the following operations is only an example. Change it based on site requirements.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login.

Using the Oozie Client

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 installation directory (change it to the actual installation directory):

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 Check the cluster authentication mode.

- If the cluster is in security mode, run the following command to authenticate the user: *exampleUser* indicates the name of the user who submits tasks.

```
kinit exampleUser
```

- If the cluster is in normal mode, go to [Step 5](#).

Step 5 Perform the following operations to configure Hue:

1. Configure the Spark2x environment (skip this step if the Spark2x task is not involved):

```
hdfs dfs -put /opt/client/Spark2x/spark/jars/*.jar /user/oozie/share/lib/spark2x/
```

When the JAR package in the HDFS directory **/user/oozie/share** changes, you need to restart the Oozie service.

2. Upload the Oozie configuration file and JAR package to HDFS.

```
hdfs dfs -mkdir /user/exampleUser
```

```
hdfs dfs -put -f /opt/client/Oozie/oozie-client-*/examples /user/exampleUser/
```

 NOTE

- *exampleUser* indicates the name of the user who submits tasks.
- If the user who submits the task and other files except **job.properties** are not changed, client installation directory **Oozie/oozie-client-*/examples** can be repeatedly used after being uploaded to HDFS.
- Resolve the JAR file conflict between Spark and Yarn about Jetty.
hdfs dfs -rm -f /user/oozie/share/lib/spark/jetty-all-9.2.22.v20170606.jar
- In normal mode, if **Permission denied** is displayed during the upload, run the following commands:
su - omm
source /opt/client/bigdata_env
hdfs dfs -chmod -R 777 /user/oozie
exit

----End

21.3 Using Oozie Client to Submit an Oozie Job

21.3.1 Submitting a Hive Job

Scenario

This section describes how to use the Oozie client to submit a Hive job.

Hive jobs are divided into the following types:

- Hive job
Hive job that is connected in JDBC mode
- Hive2 job
Hive job that is connected in Beeline mode

This section describes how to submit a Hive job using the Oozie client.

 NOTE

- The procedure for submitting a Hive2 job using the Oozie client is the same as that for submitting a Hive job. You only need to change **/Hive** in the procedure to **/Hive2**.
For example, the directory of Hive jobs is **/opt/client/Oozie/oozie-client-*/examples/apps/hive/**, and that of Hive2 jobs is **/opt/client/Oozie/oozie-client-*/examples/apps/hive2/**.
- You are advised to download the latest client.

Prerequisites

- The Hive and Oozie components and clients have been installed and are running properly.
- You have created or obtained the human-machine account and password for accessing the Oozie service.

 **NOTE**

- This user must belong to the **hadoop**, **supergroup**, and **hive** groups and be assigned with the Oozie role operation permission. If the multi-instance function is enabled for Hive, the user must belong to a specific Hive instance group, for example, **hive3**.
- This user must also be assigned the **manager_viewer** role at least.
- You have obtained the URL of the Oozie server (any instance) in the running state, for example, **https://10.1.130.10:21003/oozie**.
- You have obtained the name of the Oozie server, for example, **10-1-130-10**.
- You have obtained the IP address of the active Yarn ResourceManager, for example, **10.1.130.11**.

Procedure

Step 1 Log in to the node where the Oozie client is installed as the client installation user.

Step 2 Run the following command to obtain the installation environment. **/opt/client** is an example client installation path.

```
source /opt/client/bigdata_env
```

Step 3 Check the cluster authentication mode.

- If the cluster is in security mode, run the **kinit** command to authenticate users.

For example, the **oozieuser** user is authenticated using the following command:

```
kinit oozieuser
```

- If the cluster is in normal mode, go to [Step 4](#).

Step 4 Run the following command to go to the example directory:

```
cd /opt/client/Oozie/oozie-client-*/examples/apps/hive/
```

[Table 21-1](#) lists the files that you need to pay attention to in the directory.

Table 21-1 File description

File	Description
hive-site.xml	Configuration file of a Hive job
job.properties	Parameter definition file of a workflow
script.q	SQL script of a Hive job
workflow.xml	Rule definition file of a workflow

Step 5 Run the following command to edit the **job.properties** file:

```
vi job.properties
```

Perform the following modifications:

Change the value of **userName** to the name of the human-machine user who submits the job, for example, **userName=oozieuser**.

Step 6 Run the **oozie job** command to run the workflow file:

```
oozie job -oozie https://Host name of the Oozie role:21003/oozie/ -config  
job.properties -run
```

 **NOTE**

- The command parameters are described as follows:
 - oozie** URL of the Oozie server that executes a job
 - config** Workflow property file
 - run** Executing a workflow
- If a job ID, for example, **job: 0000021-140222101051722-oozie-omm-W**, is displayed after the workflow file is executed, the job is successfully submitted. You can view the execution results on the Oozie management page.

Log in to the Oozie web UI at **https://IP address of the Oozie role:21003/oozie** as user **oozieuser**.

On the Oozie web UI, you can view the submitted workflow information based on the job ID in the table on the page.

----End

21.3.2 Submitting a Spark2x Job

Scenario

This section describes how to submit a Spark2x job using the Oozie client.

 **NOTE**

You are advised to download the latest client.

Prerequisites

- The Spark2x and Oozie components and clients have been installed and are running properly.

If the current client is an earlier version, you need to download and install the client again.
- You have created or obtained the human-machine account and password for accessing the Oozie service.

 **NOTE**

- This user must belong to the **hadoop**, **supergroup**, and **hive** groups and be assigned with the Oozie role operation permission. If the multi-instance function is enabled for Hive, the user must belong to a specific Hive instance group, for example, **hive3**.
- This user must also be assigned the **manager_viewer** role at least.
- You have obtained the URL of the Oozie server (any instance) in the running state, for example, **https://10.1.130.10:21003/oozie**.
- You have obtained the name of the Oozie server, for example, **10-1-130-10**.

- You have obtained the IP address of the active Yarn ResourceManager, for example, **10.1.130.11**.

Procedure

Step 1 Log in to the node where the Oozie client is installed as the client installation user.

Step 2 Run the following command to obtain the installation environment. **/opt/client** is an example client installation path.

```
source /opt/client/bigdata_env
```

Step 3 Check the cluster authentication mode.

- If the cluster is in security mode, run the **kinit** command to authenticate users.

For example, the **oozieuser** user is authenticated using the following command:

```
kinit oozieuser
```

- If the cluster is in normal mode, go to [Step 4](#).

Step 4 Run the following command to go to the example directory:

```
cd /opt/client/Oozie/oozie-client-*/examples/apps/spark2x/
```

[Table 21-2](#) lists the files that you need to pay attention to in the directory.

Table 21-2 File description

File	Description
job.properties	Parameter definition file of a workflow
workflow.xml	Rule definition file of a workflow
lib	Directory of the JAR file on which a workflow depends

Step 5 Run the following command to edit the **job.properties** file:

```
vi job.properties
```

Perform the following modifications:

Change the value of **userName** to the name of the human-machine user who submits the job, for example, **userName=oozieuser**.

Step 6 Run the **oozie job** command to run the workflow file:

```
oozie job -oozie https://Host name of the Oozie role:21003/oozie/ -config job.properties -run
```

 NOTE

- The command parameters are described as follows:
 - oozie** URL of the Oozie server that executes a job
 - config** Workflow property file
 - run** Executing a workflow
- If a job ID, for example, **job: 0000021-140222101051722-oozie-omm-W**, is displayed after the workflow file is executed, the job is successfully submitted. You can view the execution results on the Oozie management page.

Log in to the Oozie web UI at **https://IP address of the Oozie role:21003/oozie** as user **oozieuser**.

On the Oozie web UI, you can view the submitted workflow information based on the job ID in the table on the page.

----End

21.3.3 Submitting a Loader Job

Scenario

This section describes how to submit a Loader job using the Oozie client.

 NOTE

You are advised to download the latest client.

Prerequisites

- The Hive and Oozie components and clients have been installed and are running properly.
- You have created or obtained the human-machine account and password for accessing the Oozie service.

 NOTE

- This user must belong to the **hadoop**, **supergroup**, and **hive** groups and be assigned with the Oozie role operation permission. If the multi-instance function is enabled for Hive, the user must belong to a specific Hive instance group, for example, **hive3**.
- This user must also be assigned the **manager_viewer** role at least.
- You have obtained the URL of the Oozie server (any instance) in the running state, for example, **https://10.1.130.10:21003/oozie**.
- You have obtained the name of the Oozie server, for example, **10-1-130-10**.
- You have obtained the IP address of the active Yarn ResourceManager, for example, **10.1.130.11**.
- You have created a Loader job to be scheduled and obtained the job ID.

Procedure

Step 1 Log in to the node where the Oozie client is installed as the client installation user.

Step 2 Run the following command to obtain the installation environment. **/opt/client** is an example client installation path.

```
source /opt/client/bigdata_env
```

Step 3 Check the cluster authentication mode.

- If the cluster is in security mode, run the **kinit** command to authenticate users.

For example, the **oozieuser** user is authenticated using the following command:

```
kinit oozieuser
```

- If the cluster is in normal mode, go to [Step 4](#).

Step 4 Run the following command to go to the example directory:

```
cd /opt/client/Oozie/oozie-client-*/examples/apps/sqoop/
```

[Table 21-3](#) lists the files that you need to pay attention to in the directory.

Table 21-3 File description

File	Description
job.properties	Parameter definition file of a workflow
workflow.xml	Rule definition file of a workflow

Step 5 Run the following command to edit the **job.properties** file:

```
vi job.properties
```

Perform the following modifications:

Change the value of **userName** to the name of the human-machine user who submits the job, for example, **userName=oozieuser**.

Step 6 Run the following command to edit the **workflow.xml** file:

```
vi workflow.xml
```

Perform the following modifications:

Change the value of **command** to the ID of the Loader job to be scheduled, for example, **1**.

Upload the **workflow.xml** file to the HDFS path in the **job.properties** file.

```
hdfs dfs -put -f workflow.xml /user/userName/examples/apps/sqoop
```

Step 7 Run the **oozie job** command to run the workflow file:

```
oozie job -oozie https://Host name of the Oozie role:21003/oozie/ -config job.properties -run
```

 NOTE

- The command parameters are described as follows:
 - oozie** URL of the Oozie server that executes a job
 - config** Workflow property file
 - run** Executing a workflow
- If a job ID, for example, **job: 0000021-140222101051722-oozie-omm-W**, is displayed after the workflow file is executed, the job is successfully submitted. You can view the execution results on the Oozie management page.

Log in to the Oozie web UI at **https://IP address of the Oozie role:21003/oozie** as user **oozieuser**.

On the Oozie web UI, you can view the submitted workflow information based on the job ID in the table on the page.

----End

21.3.4 Submitting a DistCp Job

Scenario

This section describes how to submit a DistCp job using the Oozie client.

 NOTE

You are advised to download the latest client.

Prerequisites

- The HDFS and Oozie components and clients have been installed and are running properly.

If the current client is an earlier version, you need to download and install the client again.
- You have created or obtained the human-machine account and password for accessing the Oozie service.

 NOTE

- This user must belong to the **hadoop**, **supergroup**, and **hive** groups and be assigned with the Oozie role operation permission. If the multi-instance function is enabled for Hive, the user must belong to a specific Hive instance group, for example, **hive3**.
- This user must also be assigned the **manager_viewer** role at least.
- You have obtained the URL of the Oozie server (any instance) in the running state, for example, **https://10.1.130.10:21003/oozie**.
- You have obtained the name of the Oozie server, for example, **10-1-130-10**.
- You have obtained the IP address of the active Yarn ResourceManager, for example, **10.1.130.11**.

Procedure

- Step 1** Log in to the node where the Oozie client is installed as the client installation user .

Step 2 Run the following command to obtain the installation environment. `/opt/client` is an example client installation path.

```
source /opt/client/bigdata_env
```

Step 3 Check the cluster authentication mode.

- If the cluster is in security mode, run the **kinit** command to authenticate users.

For example, the **oozieuser** user is authenticated using the following command:

```
kinit oozieuser
```

- If the cluster is in normal mode, go to [Step 4](#).

Step 4 Run the following command to go to the example directory:

```
cd /opt/client/Oozie/oozie-client-*/examples/apps/distcp/
```

[Table 21-4](#) lists the files that you need to pay attention to in the directory.

Table 21-4 File description

File	Description
job.properties	Parameter definition file of a workflow
workflow.xml	Rule definition file of a workflow

Step 5 Run the following command to edit the **job.properties** file:

```
vi job.properties
```

Perform the following modifications:

Change the value of **userName** to the name of the human-machine user who submits the job, for example, **userName=oozieuser**.

Step 6 Whether DistCp is not deployed across security clusters.

- If yes, go to [Step 7](#).
- If no, go to [Step 9](#).

Step 7 Establish cross-Manager mutual trust between two clusters.

Step 8 Run the following commands to back up and modify the **workflow.xml** file:

```
cp workflow.xml workflow.xml.bak
```

```
vi workflow.xml
```

Modify the following content:

```
<workflow-app xmlns="uri:oozie:workflow:1.0" name="distcp-wf">
  <start to="distcp-node"/>
  <action name="distcp-node">
    <distcp xmlns="uri:oozie:distcp-action:1.0">
      <resource-manager>${resourceManager}</resource-manager>
      <name-node>${nameNode}</name-node>
      <prepare>
        <delete path="hdfs://target_ip:target_port/user/${userName}/${examplesRoot}/output-data/$
```

```
{outputDir}"/>
  </prepare>
  <configuration>
    <property>
      <name>mapred.job.queue.name</name>
      <value>${queueName}</value>
    </property>
    <property>
      <name>oozie.launcher.mapreduce.job.hdfs-servers</name>
      <value>hdfs://source_ip:source_port,hdfs://target_ip:target_port</value>
    </property>
  </configuration>
  <arg>${nameNode}/user/${userName}/${examplesRoot}/input-data/text/data.txt</arg>
  <arg>hdfs://target_ip:target_port/user/${userName}/${examplesRoot}/output-data/${outputDir}/
data.txt</arg>
  </distcp>
  <ok to="end"/>
  <error to="fail"/>
</action>
<kill name="fail">
  <message>DistCP failed, error message[${wf.errorMessage(wf.lastErrorNode())}]</message>
</kill>
<end name="end"/>
</workflow-app>
```

target_ip:target_port is the HDFS active NameNode address of the other trusted cluster, for example, **10.10.10.233:25000**.

source_ip:source_port indicates the HDFS active NameNode address of the source cluster, for example, **10.10.10.223:25000**.

Change the two IP addresses and port numbers based on the site requirements.

Step 9 Run the **oozie job** command to run the workflow file:

oozie job -oozie https://Host name of the Oozie role:21003/oozie/ -config job.properties -run

NOTE

- The command parameters are described as follows:
 - oozie URL of the Oozie server that executes a job
 - config Workflow property file
 - run Executing a workflow
- If a job ID, for example, **job: 0000021-140222101051722-oozie-omm-W**, is displayed after the workflow file is executed, the job is successfully submitted. You can view the execution results on the Oozie management page.

Log in to the Oozie web UI at **https://IP address of the Oozie role:21003/oozie** as user **oozieuser**.

On the Oozie web UI, you can view the submitted workflow information based on the job ID in the table on the page.

----End

21.3.5 Submitting Other Jobs

Scenario

In addition to Hive, Spark2x, and Loader jobs, MapReduce, Java, Shell, HDFS, SSH, SubWorkflow, Streaming, and scheduled jobs can be submitted using the Oozie client.

 NOTE

You are advised to download the latest client.

Prerequisites

- The Oozie component and its client have been installed and are running properly.
- You have created or obtained the human-machine account and password for accessing the Oozie service.

 NOTE

- Shell job:
This user must belong to the **hadoop** and **supergroup** groups and be assigned the Oozie role operation permission. The Shell script must have the execution permission on each NodeManager.
- SSH job:
This user must belong to the **hadoop** and **supergroup** groups and be assigned the Oozie role operation permission. The mutual trust configuration is complete.
- Other jobs:
This user must belong to the **hadoop** and **supergroup** groups and be assigned the Oozie role operation permission and other required permissions.
- This user must also be assigned the **manager_viewer** role at least.
- You have obtained the URL of the Oozie server (any instance) in the running state, for example, **https://10.1.130.10:21003/oozie**.
- You have obtained the name of the Oozie server, for example, **10-1-130-10**.
- You have obtained the IP address of the active Yarn ResourceManager, for example, **10.1.130.11**.

Procedure

Step 1 Log in to the node where the Oozie client is installed as the client installation user.

Step 2 Run the following command to obtain the installation environment. **/opt/client** is an example client installation path.

```
source /opt/client/bigdata_env
```

Step 3 Check the cluster authentication mode.

- If the cluster is in security mode, run the **kinit** command to authenticate users.

For example, the **oozieuser** user is authenticated using the following command:

```
kinit oozieuser
```

- If the cluster is in normal mode, go to [Step 4](#).

Step 4 Go to the example directory based on the type of the task you submit.

Table 21-5 List of example directories

Job Type	Example Directory
MapReduce job	<i>Client installation directory</i> /Oozie/oozie-client-*/ examples/apps/map-reduce
Java job	<i>Client installation directory</i> /Oozie/oozie-client-*/ examples/apps/java-main
Shell job	<i>Client installation directory</i> /Oozie/oozie-client-*/ examples/apps/shell
Streaming job	<i>Client installation directory</i> /Oozie/oozie-client-*/ examples/apps/shell
SubWorkflow job	<i>Client installation directory</i> /Oozie/oozie-client-*/ examples/apps/subwf
SSH job	<i>Client installation directory</i> /Oozie/oozie-client-*/ examples/apps/ssh
Scheduled job	<i>Client installation directory</i> /Oozie/oozie-client-*/ examples/apps/cron

 **NOTE**

The examples of other jobs contain HDFS job examples.

Table 21-6 lists the files that you need to pay attention to in the example directory.

Table 21-6 File description

File	Description
job.properties	Parameter definition file of a workflow
workflow.xml	Rule definition file of a workflow
lib	Directory of the JAR file on which a workflow depends
coordinator.xml	Scheduled job configuration file which can be used to set a scheduled policy. The file is in the cron directory.
oozie_shell.sh	Shell script file required for submitting shell jobs. The file is in the shell directory.

Step 5 Run the following command to edit the **job.properties** file:

```
vi job.properties
```

Perform the following modifications:

Change the value of **userName** to the name of the human-machine user who submits the job, for example, **userName=oozieuser**.

Step 6 Run the **oozie job** command to run the workflow file:

```
oozie job -oozie https://Host name of the oozie role:21003/oozie -config File  
path of job.properties -run
```

Example:

```
oozie job -oozie https://10-1-130-10:21003/oozie -config  
/opt/client/Oozie/oozie-client-*/examples/apps/map-reduce/job.properties -  
run
```

 **NOTE**

- The command parameters are described as follows:
 - oozie** URL of the Oozie server that executes a job
 - config** Workflow property file
 - run** Executing a workflow
- If a job ID, for example, **job: 0000021-140222101051722-oozie-omm-W**, is displayed after the workflow file is executed, the job is successfully submitted. You can view the execution results on the Oozie management page.

Log in to the Oozie web UI at **https://IP address of the Oozie role:21003/oozie** as user **oozieuser**.

On the Oozie web UI, you can view the submitted workflow information based on the job ID in the table on the page.

----End

21.4 Using Hue to Submit an Oozie Job

21.4.1 Creating a Workflow

Scenario

You can submit an Oozie job on the Hue management page, but a workflow must be created before the job is submitted.

Prerequisites

Before using Hue to submit an Oozie job, configure the Oozie client and upload the sample configuration file and JAR file to the specified HDFS directory. For details, see [Using the Oozie Client](#).

Procedure

Step 1 Prepare a user who has operation permissions on the corresponding components.

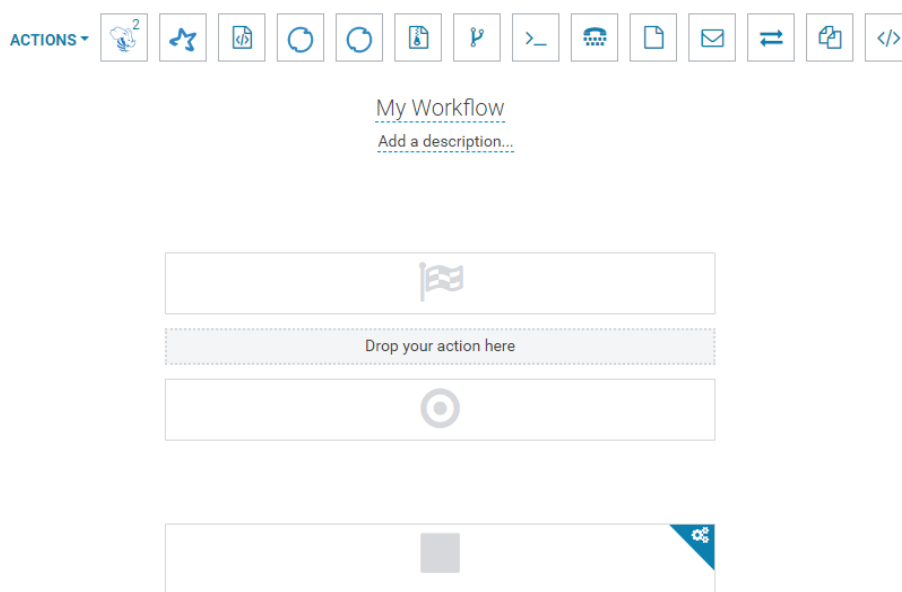
For example, log in to FusionInsight Manager as user **admin** and choose **System** in the top menu bar. On the **System** page that is displayed, choose **User** under **Permission** in the navigation pane on the left. On the displayed **User** page, click

Create. On the **Create** page, set **Username** to **hueuser** and **User Type** to **Human-Machine**, set the password and confirm it, set **User Group** to **hive, hadoop,** and **supergroup**, set **Primary Group** to **hive**, set **Role** to **System_administrator**, and click **OK**.

Step 2 Log in to FusionInsight Manager as the user created in **Step 1** (change the password upon your first login), choose **Cluster > Services > Hue**, and click the link next to **Hue WebUI** to go to the Hue WebUI page.

Step 3 In the navigation tree on the left, click  and choose **Workflow** to open the Workflow editor.

Step 4 Select **Actions** from the **DOCUMENTS** drop-down list, select the job type to be created and drag it to the operation area.



For submitting different job types, follow instructions in the following sections:

- [Submitting a Hive2 Job](#)
- [Submitting a Spark2x Job](#)
- [Submitting a Java Job](#)
- [Submitting a Loader Job](#)
- [Submitting a MapReduce Job](#)
- [Submitting a Sub-workflow Job](#)
- [Submitting a Shell Job](#)
- [Submitting an HDFS Job](#)
- [Submitting a Streaming Job](#)
- [Submitting a DistCp Job](#)

----End

21.4.2 Submitting a Workflow Job


21.4.2.1 Submitting a Hive2 Job

Scenario

This section describes how to submit an Oozie job of the Hive2 type on the Hue web UI.

Procedure


Step 1 Create a workflow. For details, see [Creating a Workflow](#).

Step 2 On the workflow editing page, select  next to **HiveServer2 Script** and drag it to the operation area.

Step 3 In the **HiveServer2 Script** dialog box that is displayed, configure the script path in the HDFS, for example, `/user/admin/examples/apps/hive2/script.q`, and click **Add**.

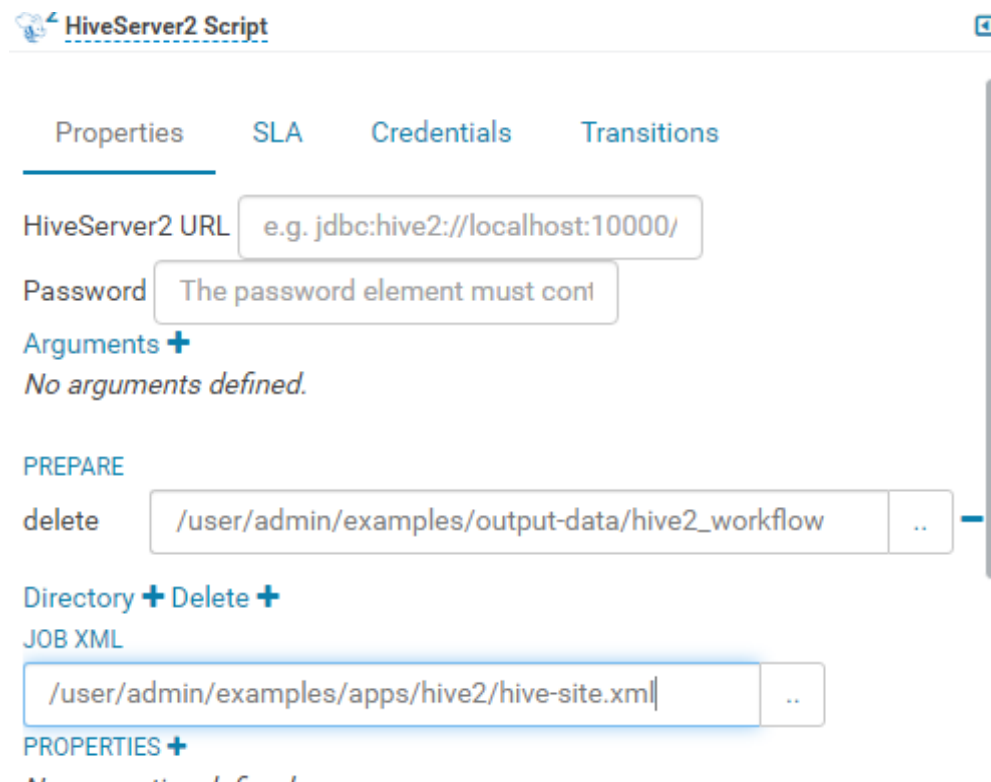
Step 4 Click **PARAMETER+** to add input and output parameters.

For example, if the input parameter is **INPUT=/user/admin/examples/input-data/table**, the output parameter is **OUTPUT=/user/admin/examples/output-data/hive2_workflow**.




Step 5 Click the configuration button  in the upper right corner. On the configuration page that is displayed, click **Delete +** to delete a directory, for example, `/user/admin/examples/output-data/hive2_workflow`.

Step 6 Configure the job XML, for example, to the HDFS path `/user/admin/examples/apps/hive2/hive-site.xml`.




NOTE

If the preceding parameters and values are modified, you can query them in **Oozie client installation directory//oozie-client-*/conf/hive-site.xml**.

Step 7 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Hive2-Workflow**.

Step 8 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End


21.4.2.2 Submitting a Spark2x Job

Scenario


This section describes how to submit an Oozie job of the Spark2x type on Hue.




Procedure

Step 1 Create a workflow. For details, see [Creating a Workflow](#).

- Step 2** On the workflow editing page, select  next to **Spark program** and drag it to the operation area.
- Step 3** In the Spark window that is displayed, set the value of **Files**, for example, to **hdfs://hacluster/user/admin/examples/apps/spark2x/lib/oozie-examples.jar**. Set the value of **jar/py name**, for example, to **org.apache.oozie.example.SparkFileCopy**, and click **Add**.
- Step 4** Set the value of **Main class**, for example, **org.apache.oozie.example.SparkFileCopy**.
- Step 5** Click **PARAMETER+** to add related input and output parameters.
For example, add the following parameters:
- **hdfs://hacluster/user/admin/examples/input-data/text/data.txt**
 - **hdfs://hacluster/user/admin/examples/output-data/spark_workflow**
- Step 6** In the **Options list** text box, specify Spark parameters, for example, **--conf spark.yarn.archive=hdfs://hacluster/user/spark2x/jars/8.1.0.1/spark-archive-2x.zip --conf spark.eventLog.enabled=true --conf spark.eventLog.dir=hdfs://hacluster/spark2xJobHistory2x**.

 **NOTE**

The version 8.1.0.1 is used as an example. You can log in to FusionInsight Manager, click  in the upper right corner, choose **About** from the drop-down list, and view the FusionInsight Manager version in the dialog box that is displayed.

- Step 7** Click the configuration button  in the upper right corner. Set the value of **Spark Master**, for example, to **yarn-cluster**. Set the value of **Mode**, for example, **cluster**.
- Step 8** On the configuration page that is displayed, click **Delete +** to delete a directory, for example, **hdfs://hacluster/user/admin/examples/output-data/spark_workflow**.
- Step 9** Click **PROPERTIES+** and add **sharelib** used by Oozie. Enter the attribute name **oozie.action.sharelib.for.spark** in the left text box and the attribute value **spark2x** in the right text box.
- Step 10** Click  in the upper right corner of the Oozie editor.
If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Spark-Workflow**.
- Step 11** After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End

21.4.2.3 Submitting a Java Job

Scenario


This section describes how to submit an Oozie job of the Java type on the Hue web UI.

Procedure


Step 1 Create a workflow. For details, see [Creating a Workflow](#).

Step 2 On the workflow editing page, select  next to **Java program** and drag it to the operation area.

Step 3 In the **Jar program** window that is displayed, set the value of **Jar name**, for example, `/user/admin/examples/apps/java-main/lib/oozie-examples-5.1.0.jar`. Set the value of **Main class**, for example, `org.apache.oozie.example.DemoJavaMain`. Click **Add**.

Step 4 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Java-Workflow**.

Step 5 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End

21.4.2.4 Submitting a Loader Job

Scenario

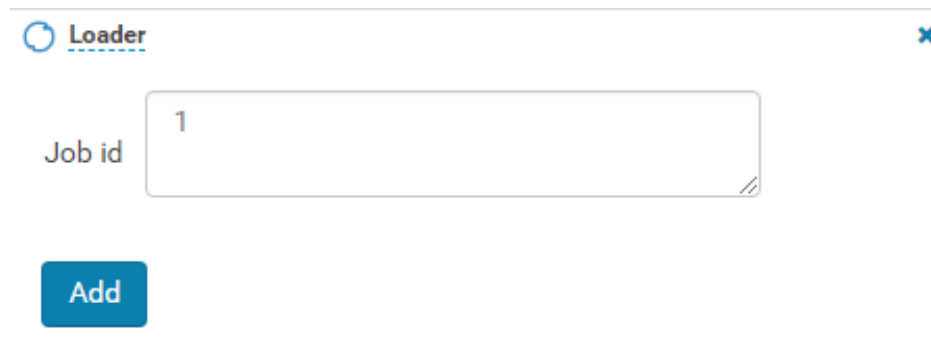
This section describes how to submit an Oozie job of the Loader type on the Hue web UI.

Procedure

Step 1 Create a workflow. For details, see [Creating a Workflow](#).

Step 2 On the workflow editing page, select  next to **Loader** and drag it to the operation area.


Step 3 In the **Loader** window that is displayed, set **Job id**, for example, to **1**. Click **Add**.




 **NOTE**

Job id is the ID of the Loader job to be orchestrated and can be obtained from the Loader page.

You can create a Loader job to be scheduled and obtain its job ID. For details, see [Using Loader](#).

Step 4 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Loader-Workflow**.

Step 5 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End

21.4.2.5 Submitting a MapReduce Job

Scenario

This section describes how to submit an Oozie job of the MapReduce type on the Hue web UI.

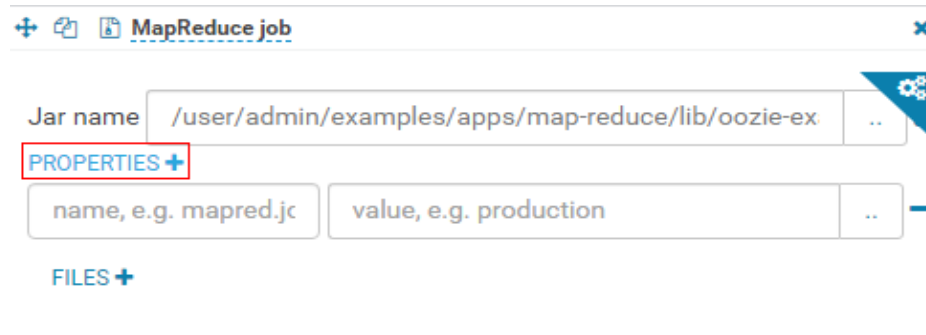
Procedure

Step 1 Create a workflow. For details, see [Creating a Workflow](#).

Step 2 On the workflow editing page, select  next to **MapReduce job** and drag it to the operation area.


Step 3 In the displayed **MapReduce job** dialog box, set **Jar name**, for example, to `/user/admin/examples/apps/map-reduce/lib/oozie-examples-5.1.0.jar`. Click **Add**.

Step 4 Click **PROPERTIES+** to add input and output properties.



For example, set the value of **mapred.input.dir** to **/user/admin/examples/input-data/text** and set the value of **mapred.output.dir** to **/user/admin/examples/output-data/map-reduce_workflow**.

Step 5 Click the configuration button  in the upper right corner. On the configuration page that is displayed, click **Delete +** to delete a directory, for example, **/user/admin/examples/output-data/map-reduce_workflow**.

Step 6 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **MapReduce-Workflow**.

Step 7 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End


21.4.2.6 Submitting a Sub-workflow Job

Scenario


This section describes how to submit an Oozie job of the Sub-workflow type on the Hue web UI.


Procedure

Step 1 Create a workflow. For details, see [Creating a Workflow](#).


Step 2 On the workflow editing page, select  next to **Sub workflow** and drag it to the operation area.

Step 3 In the **Sub workflow** dialog box that is displayed, set **Sub-workflow**, for example, to **Java-Workflow** (one of the created workflows) from the drop-down list box, and click **Add**.



Step 4 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Subworkflow-Workflow**.

Step 5 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End


21.4.2.7 Submitting a Shell Job

Scenario

This section describes how to submit an Oozie job of the Shell type on the Hue web UI.

Procedure


Step 1 Create a workflow. For details, see [Creating a Workflow](#).

Step 2 On the workflow editing page, select  next to **Shell** and drag it to the operation area.

Step 3 In the **Shell** window that is displayed, set **Shell command**, for example, to **oozie_shell.sh**, and click **Add**.

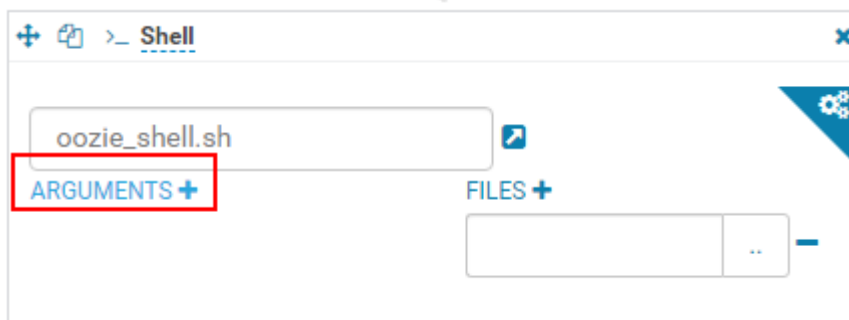
Step 4 Click **FILE+** to add the Shell command execution file or Oozie example execution file. You can select a file stored in HDFS or a local file.

- If the file is stored in HDFS, select the path of the **.sh** file, for example, **user/hueuser/shell/oozie_shell.sh**.




- If you select a local file, click **Upload a file** on the **Choose a file** page to upload the local file. After the file is uploaded, select the file.

Step 5 If the shell file to be executed needs to transfer parameters, click **ARGUMENTS+** to set parameters.



NOTE

The sequence of transferring parameters must be the same as that in the shell script.

Step 6 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Shell-Workflow**.

Step 7 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

 **NOTE**

- When configuring a shell command as a Linux command, specify it as the original command instead of the shortcut key command. For example, do not set **ls -l** to **ll**. You can configure it as the shell command **ls**, and add a parameter **-l**.
- When uploading the shell script to HDFS on Windows, make sure that the shell script format is Unix. If the format is incorrect, the shell job fails to be submitted.

----End


21.4.2.8 Submitting an HDFS Job

Scenario

This section describes how to submit an Oozie job of the HDFS type on the Hue web UI.


Procedure

Step 1 Create a workflow. For details, see [Creating a Workflow](#).


Step 2 On the workflow editing page, select  next to **Fs** and drag it to the operation area.

Step 3 In the **Fs** window that is displayed, click **Add**.

Step 4 Click **CREATE DIRECTORY+** to add the HDFS directories to be created, for example, **/user/admin/examples/output-data/mkdir_workflow** and **/user/admin/examples/output-data/mkdir_workflow1**.

Step 5 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **HDFS-Workflow**.

Step 6 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End

21.4.2.9 Submitting a Streaming Job

Scenario

This section describes how to submit an Oozie job of the Streaming type on the Hue web UI.

Procedure

Step 1 Create a workflow. For details, see [Creating a Workflow](#).

Step 2 On the workflow editing page, select  next to **Streaming** and drag it to the operation area.


Step 3 In the **Streaming** window that is displayed, set **Mapper**, for example, to `/bin/cat`. Set **Reducer**, for example, to `/usr/bin/wc`. Click **Add**.

Step 4 Click **FILE+** to add the files required for running, for example, `/user/oozie/share/lib/mapreduce-streaming/hadoop-streaming-xxx.jar` and `/user/oozie/share/lib/mapreduce-streaming/oozie-sharelib-streaming-5.1.0.jar`.

Step 5 Click the configuration button  in the upper right corner. On the configuration page that is displayed, click **Delete+** to delete a directory, for example, `/user/admin/examples/output-data/streaming_workflow`.

Step 6 Click **PROPERTIES+** to add the following properties:

- Enter the property name `mapred.input.dir` in the left box and enter the property value `/user/admin/examples/input-data/text` in the right box.
- Enter the property name `mapred.output.dir` in the left box and enter the attribute value `/user/admin/examples/output-data/streaming_workflow` in the right box.

Step 7 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Streaming-Workflow**.

Step 8 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End


21.4.2.10 Submitting a DistCp Job

Scenario

This section describes how to submit an Oozie job of the DistCp type on the Hue web UI.

Procedure

Step 1 Create a workflow. For details, see [Creating a Workflow](#).


Step 2 On the workflow editing page, select  next to **Distcp** and drag it to the operation area.

Step 3 Determine whether the current DistCp operation is performed across clusters.

- If yes, go to [Step 4](#).
- If no, go to [Step 7](#).

Step 4 Establish cross-Manager mutual trust between two clusters.


Step 5 In the **Distcp** window that is displayed, set the value of **Source**, for example, to `hdfs://hacluster/user/admin/examples/input-data/text/data.txt`. Set **Destination**, for example, to `hdfs://target_ip:target_port/user/admin/examples/output-data/distcp-workflow/data.txt`. Click **Add**.

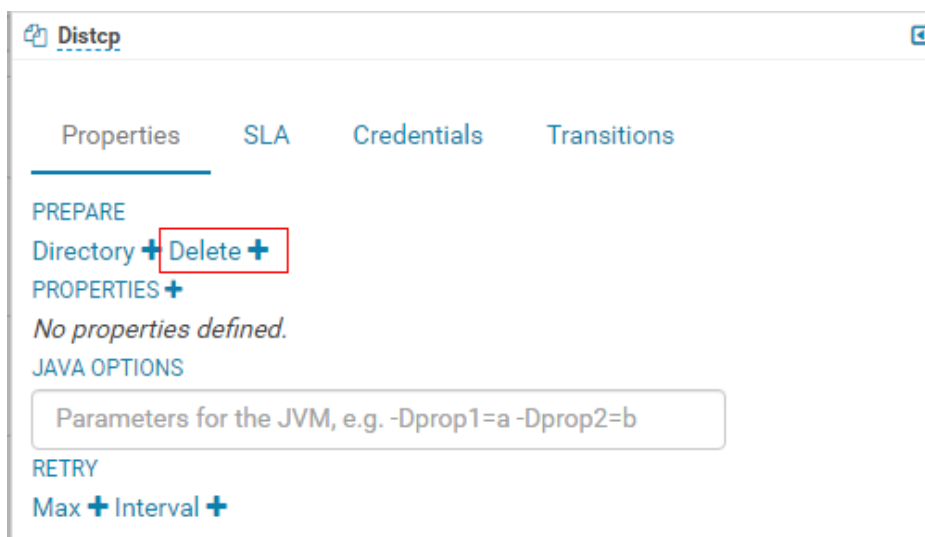
Step 6 Click the configuration button  in the upper right corner. On the **Properties** tab page, click **PROPERTIES+**, enter the attribute name `oozie.launcher.mapreduce.job.hdfs-servers` in the text box on the left, enter the attribute value `hdfs://source_ip:source_port,hdfs://target_ip:target_port` in the text box on the right, and go to [Step 8](#).


 **NOTE**

source_ip: service address of the HDFS NameNode in the source cluster
source_port: port number of the HDFS NameNode in the source cluster.
target_ip: service address of the HDFS NameNode in the target cluster
target_port: port number of the HDFS NameNode in the target cluster.


Step 7 In the **Distcp** window that is displayed, set the value of **Source**, for example, to `/user/admin/examples/input-data/text/data.txt`. Set **Destination**, for example, to `/user/admin/examples/output-data/distcp-workflow/data.txt`. Click **Add**.

Step 8 Click  in the upper right corner. On the configuration page that is displayed, click **Delete+** and add the directory to be deleted, for example, `/user/admin/examples/output-data/distcp-workflow`.



Step 9 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Distcp-Workflow**.

Step 10 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End

21.4.2.11 Example of Mutual Trust Operations

Scenario

This section guides you to enable unidirectional password-free mutual trust when Oozie nodes are used to execute shell scripts of external nodes through SSH jobs.

Prerequisites

You have installed Oozie, and it can communicate with external nodes (nodes connected using SSH).

Procedure

Step 1 Ensure that the user used for SSH connection exists on the external node, and the user directory `~/.ssh` exists.

Step 2 Log in to the node where Oozie is located as user **omm** and check whether the `~/.ssh/id_rsa.pub` file exists.

- If yes, go to [Step 3](#).
- If no, run the following command to generate a public-private key pair:

```
ssh-keygen -t rsa
```

Step 3 Log in to the node where the Oozie instance resides as user **omm** and run the following command to configure mutual trust:

```
ssh-copy-id -i ~/.ssh/id_rsa.pub User who runs SSH tasks@IP address of the node where SSH tasks run
```

You need to enter the password of the user who runs SSH tasks as prompted.

NOTE

- The user of the node where Shell resides (external node) must have the permission to execute shell scripts and access all directories and files involved in the Shell scripts.
- If there are multiple Oozie nodes, perform [Step 2](#) to [Step 3](#) on all Oozie nodes.

Step 4 Log in to other Oozie nodes as user **omm** and repeat [Step 2](#) to [Step 3](#).

----End

21.4.2.12 Submitting an SSH Job

Scenario

This section guides you to submit an Oozie job of the SSH type on the Hue web UI.


Due to security risks, SSH jobs cannot be submitted by default. To use the SSH function, you need to manually enable it.

Procedure

- Step 1** Enable the SSH function. (Skip this step if the `oozie.job.ssh.enable` parameter is unavailable for the current cluster.)
1. On FusionInsight Manager, choose **Cluster > Services > Oozie** and click the **Configurations** tab and then **All Configurations**. In the navigation pane on the left, choose **oozie(Role) > Security**, change the value of `oozie.job.ssh.enable` to **true**, and click **Save**. In the displayed dialog box, click **OK** to save the configuration.
 2. On the **Dashboard** page of Oozie, choose **More > Restart Service** in the upper-right corner to restart Oozie.

Step 2 Create a workflow. For details, see [Creating a Workflow](#).


Step 3 For details about how to add the trust relationship, see [Example of Mutual Trust Operations](#).

Step 4 On the workflow editing page, select the **Ssh** button  and drag it to the operation area.

Step 5 In the **Ssh** window that is displayed, set **User and Host** and **Ssh command** commands and click **Add**.

NOTE

User in the **User and Host** configuration item indicates a non-**omm** user. For example, the value of this configuration item can be set to `root@100.x.x.x`.

Step 6 Click  in the upper right corner of the Oozie editor.

If you need to modify the job name before saving the job (default value: **My Workflow**), click the name directly for modification, for example, **Ssh-Workflow**.

Step 7 After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End



21.4.2.13 Submitting a Hive Script


Scenario



This section describes how to submit a Hive job on the Hue web UI.

Procedure

- Step 1** Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).

- Step 2** In the navigation tree on the left, click  and choose **Workflow** to open the Workflow editor.
- Step 3** Click **Documents**, click  to select a Hive script from the operation list, and drag it to the operation page.
- Step 4** In the **HiveServer2 Script** dialog box that is displayed, select the saved Hive script. For details about how to save the Hive script, see [Using HiveQL Editor on the Hue Web UI](#). Select a script and click **Add**.



- Step 5** Configure the Job XML, for example, to the HDFS path `/user/admin/examples/apps/hive2/hive-site.xml`. For details, see [Submitting a Hive2 Job](#).
- Step 6** Click  in the upper right corner of the Oozie editor.
- Step 7** After the configuration is saved, click , and submit the job.

After the job is submitted, you can view the related contents of the job, such as the detailed information, logs, and processes, on Hue.

----End

21.4.3 Submitting a Coordinator Periodic Scheduling Job


Scenario

This section describes how to submit a job of the periodic scheduling type on the Hue web UI.

Prerequisites

Required workflow jobs have been configured before the coordinator task is submitted.

Procedure

- Step 1** Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).
- Step 2** In the navigation tree on the left, click  and choose **Schedule** to open the Coordinator editor.
- Step 3** On the job editing page, click **My Schedule** to change the job name.


Step 4 Click **Choose a workflow...** to select the workflow to be orchestrated.

My Schedule

Add a description...


Which workflow to schedule?

Choose a workflow...

Step 5 After you select the workflow, set the job execution frequency as prompted. If the workflow to be executed needs to transfer parameters, click **+ Add parameter** to set parameters and click  in the upper right corner to save the job.

 **NOTE**

The time may be several hours different from the local time due to time zone conversion.

Step 6 Click  in the upper right corner of the editor, set the start value and end value of the time range for executing the scheduled job, and click **Submit** to submit the job.

 **NOTE**

Because the time zone is changed, the difference between the time and the local time may be several hours.

----End

21.4.4 Submitting a Bundle Batch Processing Job

Scenario





In the case that multiple scheduled jobs exist at the same time, you can manage the jobs in batches over the Bundle task. This section describes how to submit a job of the batch type on the Hue web UI.

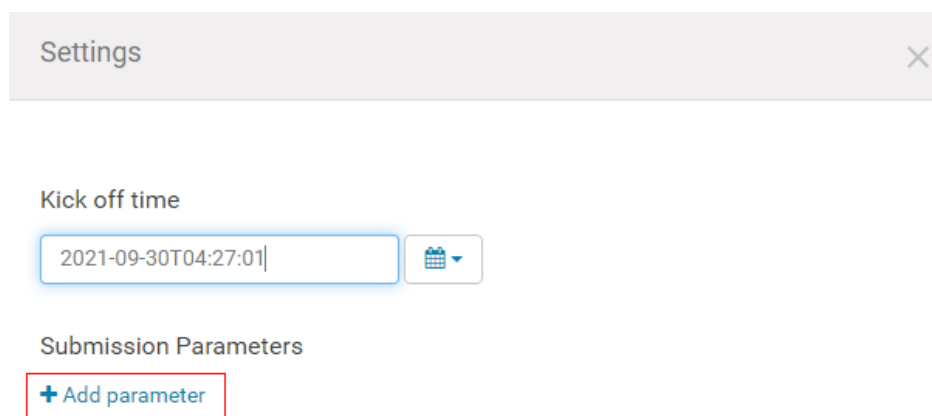
Prerequisites

Required related workflow and Coordinator jobs have been configured before the Bundle batch processing job is submitted.

Procedure


Step 1 Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).

- Step 2** In the navigation tree on the left, click  and choose **Bundle** to open the Bundle editor.
- Step 3** On the job editing page, click **My Bundle** to change the job name.
- Step 4** Click **+Add a coordinator** to select the Coordinator job to be orchestrated.
- Step 5** Set the start time and the end time for the scheduled coordinator jobs as prompted and click  in the upper right corner to save the job.
- Step 6** Click  in the upper right corner of the editor, select  from the displayed menu, set the start time of the bundle task, click **+Add parameter** to add parameters, and close the dialog box to save the settings.



 **NOTE**

Because the time zone is changed, the difference between the time and the local time may be several hours.

- Step 7** Click  in the upper right corner of the editor. In the dialog box that is displayed, click **Submit** to submit the job.


----End

21.4.5 Querying Job Execution Results

Scenario

After the jobs are submitted, you can view the execution status of a specific job on Hue.

Procedure

- Step 1** Access the Hue web UI. For details, see [Accessing the Hue Web UI](#).
- Step 2** Click . On the displayed page, you can view information about the Workflow, Schedule, and Bundle tasks.

View the jobs in the current cluster.

 **NOTE**

The number on **Job Browser** indicates the total number of jobs in the cluster.

Job Browser displays the following job information:

Table 21-7 MRS job attributes

Attribute	Description
Name	Job name
User	User who starts a job
Type	Job type
Status	Job status, including Succeeded , Running , and Failed .
Progress	Job running progress
Group	Group to which a job belongs
Start	Start time of a job
Duration	Job running duration
Id	Job ID, which is generated by the system automatically.

 **NOTE**

If the MRS cluster has Spark, the **Spark-JDBCServer** job is started by default to execute tasks.

----End

21.5 Oozie Log Overview

Log Description

Log path: The default storage paths of Oozie log files are as follows:

- Run log: **/var/log/Bigdata/oozie**
- Audit log: **/var/log/Bigdata/audit/oozie**

Log archiving rule: Oozie logs are classified into run logs, script logs, and audit logs. The maximum size of a run log file is 20 MB, and a maximum of 20 run log files can be reserved. The maximum size of an audit log file is 20 MB, and a maximum of 20 audit log files can be reserved.

 **NOTE**

A compressed log file is generated for **oozie.log** every hour. 720 compressed files (log files of one month) are retained by default.

Table 21-8 Oozie log list

Log Type	Log File Name	Description
Run log	jetty.log	Oozie built-in jetty server log file, which is used to process the request and response information of OozieServlet
	jetty.out	Oozie process startup log file
	oozie_db_temp.log	Oozie database connection log
	oozie-instrumentation.log	Oozie dashboard log file, which records the Oozie running status and configuration information of each component
	oozie-jpa.log	openJPa run log file
	oozie.log	Oozie run log file
	oozie-<SSH_USER>-<DATE>-<PID>-gc.log	Log file that records the garbage collection of the Oozie service
	oozie-ops.log	Oozie operation log file
	check-serviceDetail.log	Oozie health check logs
	oozie-error.log	Oozie running error logs
	threadDump-<DATE>.log	Log file that records stack information when the service process exits normally
Script logs	postinstallDetail.log	Work log file generated after the installation and before the startup
	prestartDetail.log	Pre-startup log file
	startDetail.log	Service startup log file
	stopDetail.log	Service stop log file
	upload-sharelib.log	Operation logs uploaded by sharelib
Audit log	oozie-audit.log	Audit log

Log Level

Table 21-9 describes the log levels provided by Oozie.

The priorities of log levels are 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.

Table 21-9 Log levels

Level	Description
ERROR	Logs of this level record abnormal information about events that cause process exceptions.
WARN	Logs of this level record exception information about the current event processing.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record system information and information about database underlying data transmission.

To modify log levels, perform the following operations:

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie** > **Configurations**.
- Step 3** Select **All Configurations**.
- Step 4** On the menu bar on the left, select the log menu of the target role.
- Step 5** Select a desired log level.
- Step 6** Click **Save**, and then click **OK**. The settings take effect after the processing is complete.

----End

Log Formats

The following table lists the Oozie log formats.

Table 21-10 Log formats

Log Type	Format	Example
Run log	<yyyy-MM-dd HH:mm:ss,SSS><Log level><Location where the log event occurs><Log level><Message in the log>	2015-05-29 21:01:45,268 INFO StatusTransitService\$StatusTransitRun- nable:539 - USER[-] GROUP[-] Released lock for [org.apache.oozie.service.StatusTransitSe rvice]
Script logs	<yyyy-MM-dd HH:mm:ss,SSS><Host name > <Log level > <Message in the log>	2015-06-01 17:18:03 001 suse11-192-168-0-111 oozie INFO Running oozie service check script

Log Type	Format	Example
Audit log	<code><yyyy-MM-dd HH:mm:ss,SSS><Log Level>< Thread name // Message in the log / Location where the log event occurs</code>	2015-06-01 22:38:41,323 INFO http- bio-21003-exec-8 IP [192.168.0.111] USER [null], GROUP [null], APP [null], JOBID [null], OPERATION [null], PARAMETER [null], RESULT [SUCCESS], HTTPCODE [200], ERRORCODE [null], ERRORMESSAGE [null] org.apache.oozie.util.XLog.log(XLog.java: 539)

21.6 Common Issues About Oozie

21.6.1 Oozie Scheduled Tasks Are Not Executed on Time

Question

Why are not Coordinator scheduled jobs executed on time on the Hue or Oozie client?

Answer

Use UTC time. For example, set `start=2016-12-20T09:00Z` in `job.properties` file.

21.6.2 Why Update of the share lib Directory of Oozie on HDFS Does Not Take Effect?

Symptom

A new JAR package is uploaded to the `/user/oozie/share/lib` directory on HDFS. However, an error indicating that the class cannot be found is reported during task execution.

Solution

Run the following command on the client to refresh the directory:

```
oozie admin -oozie https://xxx.xxx.xxx.xxx:21003/oozie -sharelibupdate
```

21.6.3 Common Oozie Troubleshooting Methods

1. Check the job logs on Yarn. Run the command executed through Hive SQL using beeline to ensure that Hive is running properly.
2. If error information such as "classnotfoundException" is displayed, check whether the JAR package of the faulty class exists in the `/user/oozie/share/lib` directory of each component. If no, add the JAR package and go to

Why Update of the share lib Directory of Oozie on HDFS Does Not Take Effect? If the faulty class still cannot be found after the **share lib** directory is updated, check whether **sharelibDirNew** is **/user/oozie/share/lib** in the output of the command for updating the directory.

3. If "NoSuchMethodError" is displayed, check whether the JAR packages of each component in the **/user/oozie/share/lib** directory have multiple versions. Note that the JAR packages uploaded by the service cannot conflict with each other. You can check whether a JAR package conflict occurs based on the loaded JAR packages in Oozie run logs on Yarn.
4. If the self-developed code is abnormal, run the Oozie sample to check whether Oozie is running properly.
5. Contact technical support personnel. By using this method, you must collect run logs of Oozie on Yarn, Oozie logs, and component run logs. For example, if an exception occurs when Hive runs on Oozie, you need to collect Hive logs.

22 Using Presto

22.1 Accessing the Presto Web UI

You can view the Presto statistics on the graphical Presto web UI. You are advised to use Google Chrome to access the Presto web UI because it cannot be accessed using Internet Explorer.

Prerequisites

- Presto has been installed in a cluster.
- The cluster client has been installed in a directory, for example, **/opt/client**. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite.

Accessing the Presto Web UI

- Method 1 (for MRS 3.x or later)
 - a. Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the target cluster* > **Services**.
 - b. Select **Presto**. In the **Basic Information** area, click **Coordinator(Coordinator)** next to **Coordinator WebUI**. The Coordinator web UI is displayed.

Figure 22-1 Coordinator WebUI

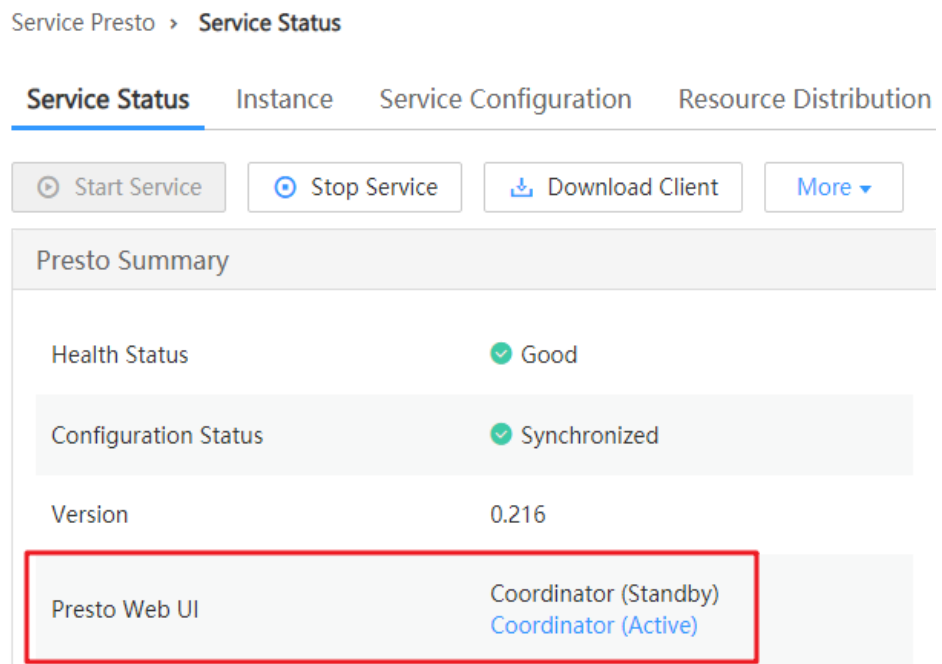
Basic Information

Running Status: ● Normal
Configuration Status: ⊕ Synchronized
Version: 333

Coordinator WebUI: [Coordinator\(Coordinator\)](#)
[Coordinator\(Coordinator\)](#)

- Method 2 (for versions earlier than MRS 3.x)
 - a. Log in to MRS Manager and choose **Services**.
 - b. Select **Presto**. In the **Presto Summary** area, click **Coordinator (Active)** next to **Presto Web UI**. The Presto web UI is displayed.

Figure 22-2 Presto WebUI



NOTE

When accessing the Presto web UI for the first time, you must add the address to the trusted site list.

22.2 Using a Client to Execute Query Statements

You can perform an interactive query on an MRS cluster client. For clusters with Kerberos authentication enabled, users who submit topologies must belong to the **presto** group.

The Presto component of MRS 3.x does not support Kerberos authentication.

Prerequisites

- The password of user **admin** has been obtained. The password of user **admin** is specified by the user during MRS cluster creation.
- The client has been updated.
- The Presto client has been manually installed for MRS 3.x clusters.

Procedure

- Step 1** For clusters with Kerberos authentication enabled, log in to MRS Manager and create a role with the **Hive Admin Privilege** permission.
- Step 2** Create a user to the **Presto** and **Hive** groups, bind the role created in **Step 1** to the user, and download the user authentication file..
- Step 3** Upload the downloaded **user.keytab** and **krb5.conf** files to the node where the MRS client resides.

 **NOTE**

For clusters with Kerberos authentication enabled, **Step 2** to **Step 3** must be performed. For normal clusters, start from **Step 4**.

- Step 4** Prepare a client based on service conditions and log in to the node where the client is installed.

For example, if you have updated the client on the Master2 node, log in to the Master2 node to use the client. For details, see [Using an MRS Client](#).

- Step 5** Run the following command to switch the user:

```
sudo su - omm
```

- Step 6** Run the following command to switch to the client directory, for example, **/opt/client**.

```
cd /opt/client
```

- Step 7** Run the following command to configure environment variables:

```
source bigdata_env
```

- Step 8** Connect to the Presto Server. The following provides two client connection methods based on the client type.

- Using the client provided by MRS
 - For clusters with Kerberos authentication disabled, run the following command to connect to the Presto Server of the cluster:
presto_cli.sh
 - For clusters with Kerberos authentication disabled, run the following command to connect to the Presto Server of other clusters. In the command, **ip** indicates the floating IP address of the cluster Presto Server, which can be obtained by searching for **PRESTO_COORDINATOR_FLOAT_IP** in the Presto configuration items. **port** indicates the Presto Server port number and is set to **7520** by default.
presto_cli.sh --server http://ip:port
 - For clusters with Kerberos authentication enabled, 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's principal* --krb5-keytab-path *user.keytab file path* --user *presto username*
 - For clusters with Kerberos authentication enabled, run the following command to connect to the Presto Server of other clusters. In the

command, **ip** indicates the floating IP address of the cluster Presto Server, which can be obtained by searching for **PRESTO_COORDINATOR_FLOAT_IP** in the Presto configuration items. **port** indicates the Presto Server port number and is set to **7521** by default.

```
presto_cli.sh --krb5-config-path krb5.conf file path --krb5-principal User's principal --krb5-keytab-path user.keytab file path --server https://ip:port --krb5-remote-service-name Presto Server name
```

- Using the native client
The native client of Presto is **Presto/presto/bin/presto** in the client directory.

Step 9 Run a query statement, for example, **show catalogs;**

 **NOTE**

For clusters with Kerberos authentication enabled, when querying **Hive Catalog** data, the user who runs the Presto client must have the permission to access Hive tables and run the **grant all on table [table_name] to group hive;** command in Hive beeline to grant permissions to the Hive group.

Step 10 After the query is complete, run the following command to exit the client:

```
quit;  
----End
```

22.3 Presto FAQ

22.3.1 How Do I Configure Multiple Hive Connections for Presto?

Question

How do I configure multiple Hive connections for Presto?

Procedure

Step 1 Copy and distribute the **core-site.xml** and **hdfs-site.xml** files of the target Hive cluster to the Presto cluster and place them in a path (for example, **/home/omm**) on which user **omm** has the read permission. Change the file owner to **omm:wheel** and file permission to **750**.

Step 2 Go to the Presto configuration page.

- For versions later than MRS 1.8.10 and version MRS 2.x: Click the cluster name on the MRS console, choose **Components > Presto > Service Configuration**, and select **All** from the **Basic** drop-down list.

 **NOTE**

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later: Log in to FusionInsight Manager and choose **Cluster > Name of the desired cluster > Services > Presto**. Click **Configurations** then **All Configurations**.

Step 3 Search for **connector-customize** in the search box.

Step 4 Add a connector named **myhive**.

Add the following configuration to **connector-customize**:

```
myhive.connector.name=hive-hadoop2
```

```
myhive.hive.metastore.uri=thrift://{Metastore server IP address of the target Hive cluster}:{Metastore server port of the target Hive cluster} (Use commas (,) to separate multiple Metastore server addresses, for example, thrift://192.0.2.1:9083,thrift://192.0.2.2:9083.)
```

```
myhive.hive.config.resources=/etc/hadoop/conf/core-site.xml,/etc/hadoop/conf/hdfs-site.xml (absolute path of the configuration files in 1)
```

For details about other configurations, see <https://trino.io/docs/333/connector/hive.html#hive-configuration-properties>.

When you add a configuration, add the connector name as the prefix to the configuration name. In this example, the connector name is **myhive**.

Step 5 Save the configuration and restart the Presto service.

 **NOTE**

Use the Presto client, JDBC, or UI tool to connect to the Presto service. Run the **show schemas from myhive;** command. If the database in the target cluster is displayed, the configuration is successful.

----End

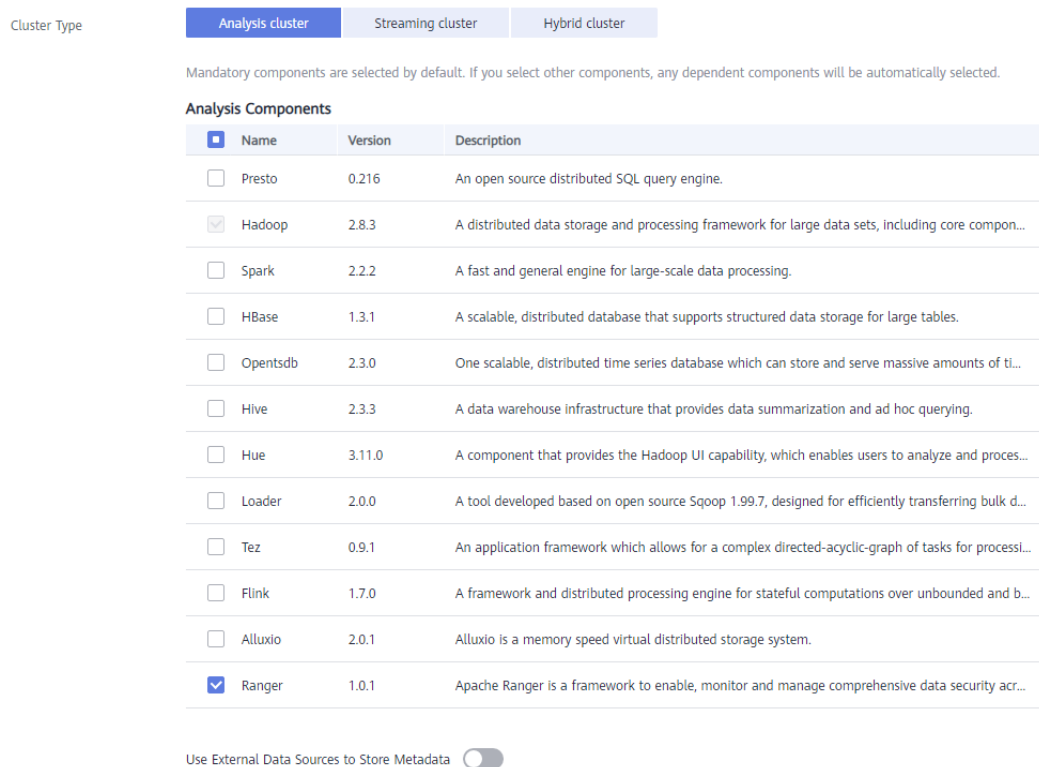
23 Using Ranger (MRS 1.9.2)

23.1 Creating a Ranger Cluster

Step 1 Create a cluster by referring to **Configuring a Cluster > Custom Creation of a Cluster** in the *User Guide*. Select the Ranger component during cluster creation.

Currently, only normal MRS 1.9.2 clusters support Ranger. Security clusters with Kerberos authentication enabled do not support Ranger.

Figure 23-1 Selecting the Ranger component



Step 2 Enable or disable **Use External Data Sources to Store Metadata**.

- Enabled: An external MySQL database is used to store the user, group, and policy data of Ranger.
- Disabled: The user, group, and policy data of Ranger is stored in the local database of the current cluster by default.

Step 3 If **Use External Data Sources to Store Metadata** is enabled, set **Data Connection Type** to **RDS MySQL database**. Select an existing data connection instance or click **Create Data Connection** to create a data connection.

Figure 23-2 Using the RDS MySQL database

Use External Data Sources to Store Metadata

Name	Data Connection Type	Data Connection Instance
Ranger	RDS MySQL database	<input type="text"/>

[Create Data Connection](#)

The user configured in the data connection used by Ranger must be the root user or authorized user.

NOTE

If the selected data connection is an **RDS MySQL database**, ensure that the database user is a **root** user. If the database user is not a **root** user, log in to the database as user **root** and run the following SQL statement to grant permissions to the database user. In the command, **`\${db_name}`** and **`\${db_user}`** indicate the database name and username entered during data connection creation.

```
grant select on mysql.user to `${db_user}`;
grant all privileges on `${db_name}`.* to `${db_user}`@'%' with grant option;
grant reload on *.* to `${db_user}`@'%' with grant option;
flush privileges;
```

Step 4 Configure other parameters by referring to **Configuring a Cluster > Custom Creation of a Cluster** in the *User Guide*.

NOTE

- After the cluster is created, Ranger does not control users' permissions to access Hive and HBase.
- When Ranger is used to manage component permissions, for example, manage Hive table permissions, if a user submits a Hive job (operation on Hive data tables) on the interface or client, a message may be displayed indicating that the user does not have the permissions. In this case, you need to configure the database or table permissions for the user who submits the job in Ranger. For details, see the step for adding a policy in [Configuring Hive/Impala Access Permissions in Ranger](#) or [Configuring HBase Access Permissions in Ranger](#).

----End

23.2 Accessing the Ranger Web UI and Synchronizing Unix Users to the Ranger Web UI

You can manage Ranger on the Ranger web UI.

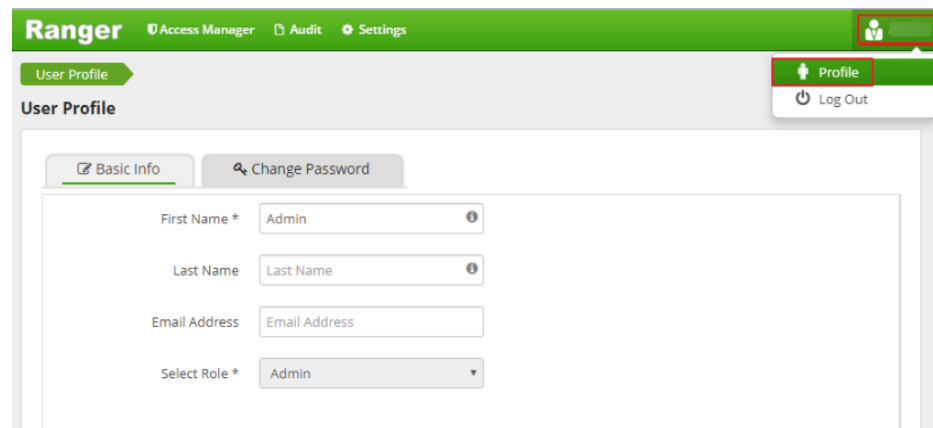
Accessing the Ranger Admin Web UI

- Step 1** On the MRS management console, click the cluster name to go to the cluster details page.
- Step 2** Click the **Components** tab.
- Step 3** Select **Ranger**. In **Ranger Summary**, click **RangerAdmin** corresponding to **Ranger Web UI**.
- Step 4** On the Ranger web UI login page, the default username for MRS 1.9.2 is **admin** and the password is **admin@12345**. The default username for MRS 1.9.3 is **admin** and the password is **ranger@A1!**.

After logging in to the Ranger Web UI for the first time, change the password and keep it secure.

- Step 5** Click the username in the upper right corner, choose **Profile** from the drop-down list, and click **Change Password** to change the password.

Figure 23-3 Changing the Ranger web UI login password



- Step 6** After changing the password, click the username in the upper right corner, choose **Log Out** from the drop-down list, and use the new password to log in to the web UI again.

----End

Using Ranger UserSync to Synchronize Unix OS Users on Cluster Nodes

Ranger UserSync is an important component of Ranger. It can synchronize Unix system users or LDAP users to the Ranger web UI. Currently, MRS can synchronize only Unix users on the node where the Ranger UserSync process resides.

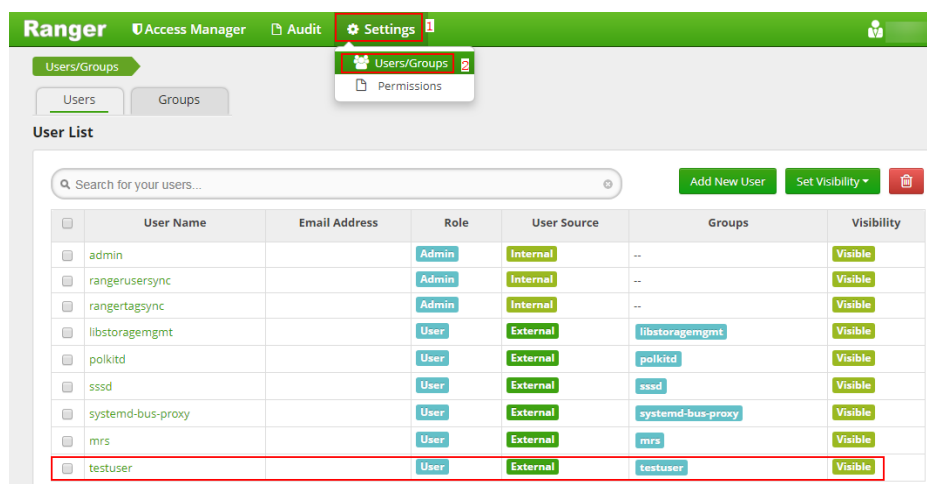
- Step 1** Log in to the node where the Ranger UserSync process is located.
- Step 2** Run the **useradd** command to add a system user, for example, **testuser**.

Figure 23-4 Adding the testuser system user

```
[root@node-master1aHRf ~]# useradd testuser
[root@node-master1aHRf ~]# passwd testuser
Changing password for user testuser.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
```

Step 3 After the user is added, wait for about 1 minute and log in to the Ranger web UI. Then, you can see that the user is synchronized.

Figure 23-5 User synchronization completed



----End


23.3 Configuring Hive/Impala Access Permissions in Ranger

After an MRS cluster with Ranger installed is created, Hive and Impala access control is not integrated into Ranger. This section describes how to integrate Hive into Ranger. Impala follows the same procedure.

Step 1 Log in to the Ranger web UI.

Step 2 In the **Service Manager** area, click **+** next to **HIVE** to add a Hive service.

Figure 23-6 Adding a Hive service



Step 3 Set the parameters for adding a Hive service according to [Table 23-1](#). Use the default values for the parameters that are not listed in the table.

Table 23-1 Parameter description

Parameter	Description	Example Value
Service Name	Name of the service to be created. The value is fixed to hivedev .	hivedev
Username	You can set this parameter to any value.	admin
Password	You can set this parameter to any value.	-
jdbc.driverClassName	Driver class for connecting to Hive. The value is fixed to org.apache.hive.jdbc.HiveDriver .	org.apache.hive.jdbc.HiveDriver
jdbc.url	URL for connecting to Hive. The format is ZooKeeper mode: jdbc:hive2://<host>:2181;/serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2 <host> indicates a ZooKeeper address. To obtain the ZooKeeper address, log in to MRS Manager, choose Services > ZooKeeper > Instance , and view the management IP address of the ZooKeeper instance.	jdbc:hive2://xx.xx.xx.xx:2181,xx.xx.xx.xx:2181,xx.xx.xx.xx:2181;/serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2

Figure 23-7 Creating hivedev

The screenshot shows the 'Create Service' form in the Service Manager. The form is divided into two main sections: 'Service Details' and 'Config Properties'.
Service Details:
- Service Name *: hivedev
- Description: (empty text area)
- Active Status: Enabled Disabled
- Select Tag Service: (dropdown menu)
Config Properties:
- Username *: admin
- Password *:
- jdbc.driverClassName *: org.apache.hive.jdbc.HiveDriver
- jdbc.url *: jdbc:hive2://192.168.0.7:2181,1:
- Common Name for Certificate: (empty text area)
- Add New Configurations: A table with columns 'Name' and 'Value'. Below the table is a '+' button.
- Test Connection: (button)
At the bottom of the form are 'Add' and 'Cancel' buttons.

Step 4 Click **Add** to add the service.

Step 5 Start the Ranger Hive plugin to authorize Ranger to manage Hive.

1. On the MRS management console, click the cluster name to go to the cluster details page.
2. Click the **Components** tab.
3. Choose **Hive > Service Configuration** and switch **Basic** to **All**.
4. Search for **hive.security.authorization** and modify the following configurations:
 - hive.security.authorization.enabled = true
 - hive.security.authorization.manager = org.apache.ranger.authorization.hive.authorizer.RangerHiveAuthorizerFactory
5. Click **Save Configuration** and select **Restart the affected services or instances** to restart the Hive service.

Step 6 Add an access control policy.

1. Log in to the Ranger web UI.
2. In the **HIVE** area, click the added service **hivedev**.
3. Click **Add New Policy** to add an access control policy.

4. Set the parameters according to [Table 23-2](#). Use the default values for the parameters that are not listed in the table.

Table 23-2 Parameter description

Parameter	Description	Example Value
Policy Name	Policy name	Policy001
database	Name of the database that the policy allows to access	test
table	Name of the table corresponding to the database that the policy allows to access	table1
Hive Column	Column name of the table corresponding to the database that the policy allows to access	name
Allow Conditions	<ul style="list-style-type: none"> - Select Group: user group that the policy allows to access - Select User: user in the user group that the policy allows to access - Permissions: permissions that the policy allows the user to have 	<ul style="list-style-type: none"> - Select Group: testuser - Select User: testuser - Permissions: Create and Select

Figure 23-8 Adding an access control policy for **hivedev**

The screenshot shows the 'Create Policy' interface in Ranger. The 'Policy Details' section includes the following fields and values:

- Policy Type: Access
- Policy Name: Policy001 (enabled)
- Database: test (include)
- Table: table1 (include)
- Hive Column: name (include)
- Audit Logging: YES
- Description: (empty)
- Policy Label: Policy Label

The 'Allow Conditions' section contains a table with the following data:

Select Group	Select User	Permissions	Delegate Admin
x testuser	x testuser	Create select	<input type="checkbox"/>

5. Click **Add** to add the policy. According to the preceding policy, user **testuser** in the **testuser** user group has the **Create** and **Select** permissions on the **name** column of **table1** in the **test** database of Hive, but no permissions to access other columns.

Step 7 Log in to the Hive client by referring to [Using Hive from Scratch](#), and check whether Hive has been integrated into Ranger.

1. Run the following command to access the Hive beeline:

```
source /opt/client/bigdata_env
beeline
```

2. Run the following command to set up a connection and log in as user **testuser**:

```
!connect jdbc:hive2://
xx.xx.xx.xx:2181,xx.xx.3.81:2181,192.168.3.153:2181;/serviceDiscoveryMode
=zooKeeper;zooKeeperNamespace=hiveserver2
```

Figure 23-9 Logging in to Hive

```
[root@node-master1ahRf ~]# source /opt/client/bigdata_env
[root@node-master1ahRf ~]# beeline
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/opt/share/log4j-slf4j-impl-2.6.2/log4j-slf4j-impl-2.6.2.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: Found binding in [jar:file:/opt/share/slf4j-log4j12-1.7.5/slf4j-log4j12-1.7.5.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.apache.logging.slf4j.Log4jLoggerFactory]
Connecting to jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181;/serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2
Connected to: Apache Hive (version )
Driver: Hive JDBC (version )
Transaction isolation: TRANSACTION_REPEATABLE_READ
Beeline version  by Apache Hive
0: jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181/> !connect jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181;/serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2
Connecting to jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181;/serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2
Enter username for jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181/: testuser
Enter password for jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181/:
Connected to: Apache Hive (version )
Driver: Hive JDBC (version )
Transaction isolation: TRANSACTION_REPEATABLE_READ
0: jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181/>
```

3. Query data and check whether Ranger is integrated.

Figure 23-10 Verifying the integration of Ranger with Hive

```
0: jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181/> select * from table1;
Error: Error while compiling statement: FAILED: HiveAccessControlException Permission denied: user [testuser] does not have [SELECT] privilege on [test/table1/*] (state=42000,code=40000)
0: jdbc:hive2://192.168.3.88:2181,192.168.3.83:2181,192.168.3.244:2181/> select name from table1;
INFO : State: Compiling.
INFO : Compiling command(queryId=omn_20191204095459_5713dad6-1fff-4a98-96f8-0da351c63df9): select name from table1
INFO : Semantic Analysis Completed
INFO : Returning Hive schema: Schema(fieldSchemas:[FieldSchema(name=name, type:string, comment:null)], properties:null)
INFO : EXPLAIN output for queryId omn_20191204095459_5713dad6-1fff-4a98-96f8-0da351c63df9 : STAGE DEPENDENCIES:
Stage-0 is a root stage [FETCH]

STAGE PLANS:
Stage: Stage-0
Fetch Operator
limit: -1
Processor Tree:
TableScan
alias: table1
GatherStats: false
Select Operator
expressions: name (type: string)
outputColumnNames: _col0
ListSink

INFO : Completed compiling command(queryId=omn_20191204095459_5713dad6-1fff-4a98-96f8-0da351c63df9); Time taken: 0.12 seconds
INFO : Concurrency mode is disabled, not creating a lock manager
INFO : State: Executing.
INFO : Executing command(queryId=omn_20191204095459_5713dad6-1fff-4a98-96f8-0da351c63df9): select name from table1
INFO : Completed executing command(queryId=omn_20191204095459_5713dad6-1fff-4a98-96f8-0da351c63df9); Time taken: 0.001 seconds
INFO : OK
+-----+
| name |
+-----+
+-----+
No rows selected (0.167 seconds)
```

----End


23.4 Configuring HBase Access Permissions in Ranger

After an MRS cluster with Ranger installed is created, HBase access control is not integrated into Ranger. This section describes how to integrate HBase into Ranger.

Step 1 Log in to the Ranger web UI.

Step 2 In the **Service Manager** area, click  next to **HBASE** to add an HBase service.

Figure 23-11 Adding an HBase service

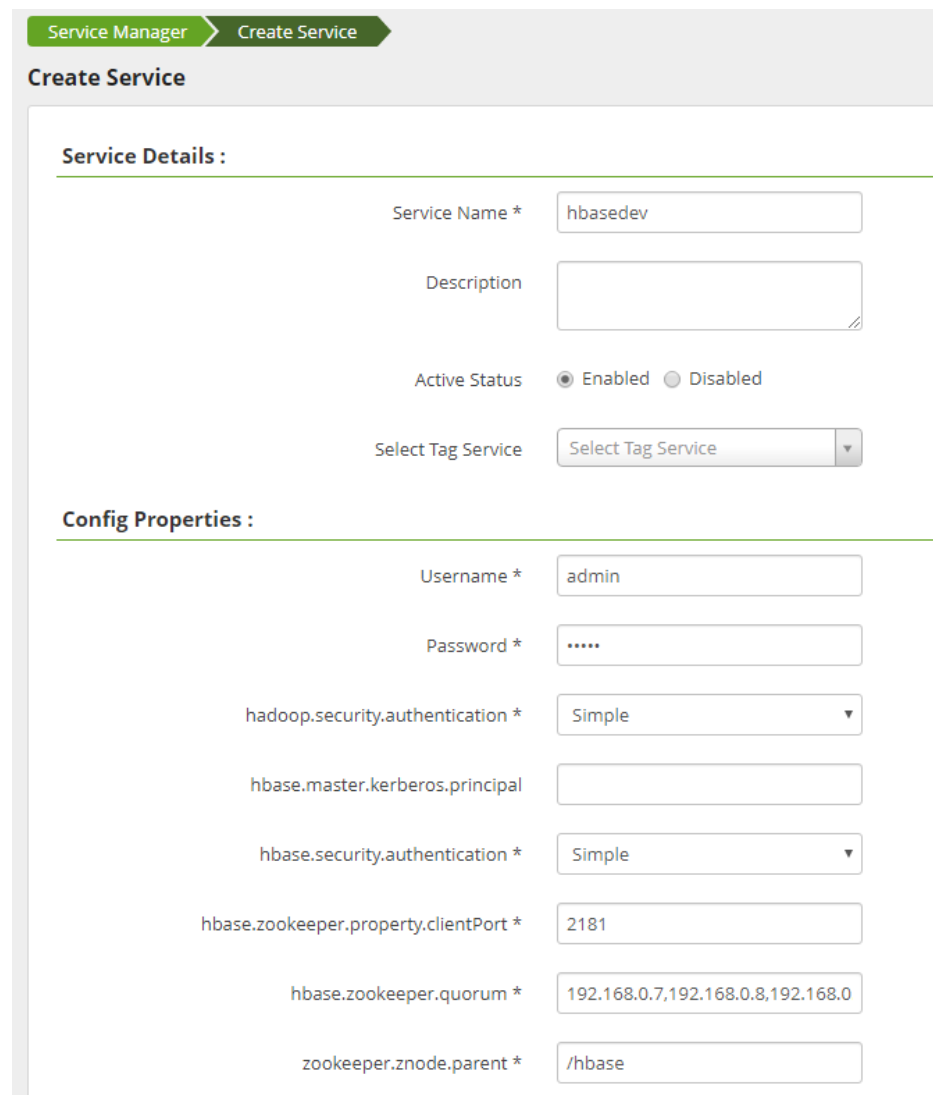


Step 3 Set the parameters for adding an HBase service according to [Table 23-3](#). Use the default values for the parameters that are not listed in the table.

Table 23-3 Parameter description

Parameter	Description	Example Value
Service Name	Name of the service to be created. The value is fixed to hbasedev .	hbasedev
Username	You can set this parameter to any value.	admin
Password	You can set this parameter to any value.	-
hadoop.security.authentication	Hadoop authentication mode. The value is fixed to Simple .	Simple
hbase.security.authentication	HBase authentication mode. The value is fixed to Simple .	Simple
hbase.zookeeper.property.clientPort	Port number of ZooKeeper in the HBase cluster.	2181
hbase.zookeeper.quorum	ZooKeeper address in the HBase cluster.	192.168.0.7,192.168.0.8,192.168.0.9
zookeeper.znode.parent	Path of the root node of HBase in ZooKeeper. The value is fixed to /hbase .	/hbase

Figure 23-12 Creating hbasedev



Service Manager > **Create Service**

Create Service

Service Details :

Service Name *

Description

Active Status Enabled Disabled

Select Tag Service

Config Properties :

Username *

Password *

hadoop.security.authentication *

hbase.master.kerberos.principal

hbase.security.authentication *

hbase.zookeeper.property.clientPort *

hbase.zookeeper.quorum *


zookeeper.znode.parent *

Step 4 Click **Add** to add the service.

Step 5 Start the Ranger HBase plugin to authorize Ranger to manage HBase.


1. On the MRS management console, click the cluster name to go to the cluster details page.
2. Click the **Components** tab.
3. Choose **HBase > Service Configuration** and switch **Basic** to **All**.
4. Search for **hbase.security.authorization** and change its value to **true** (select the first HBase parameter).

Figure 23-13 Modifying `hbase.security.authorization`




5. Search for `hbase.coprocessor.master.classes` and append `,org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor` to its original value.

Figure 23-14 `hbase.coprocessor.master.classes`



6. Search for `hbase.coprocessor.region.classes` and append `,org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor` to its original value.

Figure 23-15 `hbase.coprocessor.region.classes`



7. Click **Save Configuration** and select **Restart the affected services or instances** to restart the HMaster and RegionServer instances.

Step 6 Create a policy under **HBase Service hbasedev**.

1. Log in to the Ranger web UI.
2. In the **HBASE** area, click the added service **hbasedev**.
3. Click **Add New Policy** to add an access control policy.
4. Set the parameters according to **Table 23-4**. Use the default values for the parameters that are not listed in the table.

Table 23-4 Parameter description

Parameter	Description	Example Value
Policy Name	Policy name	Policy002

Parameter	Description	Example Value
HBase Table	Name of the HBase table that the policy allows to access	test1
HBase Column-family	Column family of the HBase table that the policy allows to access	cf1
HBase Column	Column name of the table corresponding to the HBase table that the policy allows to access	name
Allow Conditions	<ul style="list-style-type: none"> - Select Group: user group that the policy allows to access - Select User: user in the user group that the policy allows to access - Permissions: permissions that the policy allows the user to have 	<ul style="list-style-type: none"> - Select Group: testuser - Select User: testuser - Permissions: Create and Select

Figure 23-16 Adding an access control policy for hbasedev

Service Manager > hbasedev Policies > Create Policy

Create Policy

Policy Details :

Policy Type **Access**

Policy Name * **enabled**

HBase Table * **include**

HBase Column-family * **include**

HBase Column * **include**

Audit Logging **YES**

Description

Policy Label

Allow Conditions :

Select Group	Select User	Permissions	Delegate Admin
<input type="text" value="testuser"/>	<input type="text" value="testuser"/>	Read Write	<input type="checkbox"/>

5. Click **Add** to add the policy. According to the preceding policy, user **testuser** in the **testuser** user group has the **Create** and **Select** permissions on the **cf1:name** column in the **test1** table of the **default** namespace in HBase, but no permissions to access other columns.

Step 7 Update and log in to the HBase client by referring to [Using HBase from Scratch](#), and check whether HBase has been integrated into Ranger.

1. Run the following command to access the HBase shell:

```
source /opt/client/bigdata_env  
hbase shell
```

Figure 23-17 Accessing the HBase shell

```
[root@node-master1aHRf conf]# source /opt/client/bigdata_env  
[root@node-master1aHRf conf]# hbase shell  
INFO: Watching file:/opt/client/HBase/hbase/conf/log4j.properties for changes with interval : 60000
```

2. Add data and check whether Ranger is integrated.
 - a. Add data to the **cf1:name** column in the **test1** table.
put 'test1','001','cf1:name','tom'
 - b. Add data to the **cf1:age** column in the **test1** table. If the user has no permission to access this column, the data fails to be added.
put 'test1','001','cf1:age',10

Figure 23-18 Verifying the integration of Ranger with HBase

```
hbase(main):037:0> put 'test1','001','cf1:name','Tom'  
0 row(s) in 0.1120 seconds  
  
hbase(main):038:0> scan 'test1'  
ROW COLUMN+CELL  
001 column=cf1:name, timestamp=1575426581007, value=Tom  
1 row(s) in 0.0170 seconds  
  
hbase(main):039:0> put 'test1','001','cf1:age',10  
2019-12-04 10:30:15,637 WARN [hconnection-0x52ff99cd-shared--pool1-t26] client.AsyncProcess: #3, table=test1, attempt=1/7 failed=ops, last exception: org.apache.hadoop.hbase.security.AccessDeniedException: org.apache.hadoop.hbase.security.AccessDeniedException: Insufficient permissions for user 'testuser',action: put, tableName=test1, family:cf1, column: age  
at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.requirePermission(RangerAuthorizationCoprocessor.java:582)
```

----End

24 Using Ranger (MRS 3.x)

24.1 Logging In to the Ranger Web UI

Ranger provides a centralized permission management framework to implement fine-grained permission access control on components, such as HDFS, HBase, Hive, and YARN, and provides a web UI for Ranger administrators to perform operations.

Ranger User Type

Ranger users are classified into **admin**, **user**, and **auditor**. Different users have different permissions to view and operate the Ranger management interface.

- **Admin:** A Ranger security administrator who can view all page content, manage permission management plug-ins and access control policies, view audit information, and set user types.
- **Auditor:** A Ranger audit administrator who can view the permission management plug-ins and access control policies.
- **User:** A common user who can be assigned with specific permissions by the Ranger administrator.

Logging In to the Ranger Web UI

Security mode (Kerberos authentication is enabled for clusters)

Step 1 Log in to FusionInsight Manager as user **admin**. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Services > Ranger**. The Ranger service overview page is displayed.

Step 2 Click **RangerAdmin** in the **Basic Information** area. The Ranger web UI is displayed.

- The **admin** user in Ranger belongs to the **User** type and can only view the **Access Manager** as well as **Security Zone** pages.
- To view all management pages, switch to user **rangeradmin** or other users who have the Ranger administrator permissions.

- a. On the Ranger WebUI, click the user name in the upper right corner and choose **Log Out** to log out of the Ranger WebUI.



- b. Log in to the system as user **rangeradmin** (default password: **Rangeradmin@123**) or another user who has the Ranger administrator permissions.

----End

Normal mode (Kerberos authentication is disabled for clusters)

Step 1 Log in to FusionInsight Manager as user **admin**. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster > Services > Ranger**. The Ranger service overview page is displayed.

Step 2 Click **RangerAdmin** in the **Basic Information** area. The Ranger web UI is displayed.

The **admin** user in Ranger belongs to the **Admin** type and can view all management pages of Ranger without switching to user **rangeradmin**.

NOTE

When a user logs in to the Ranger WebUI as user **rangeradmin** in normal mode, error 401 is reported.

----End

On the homepage of Ranger web UI, you can view the permission management plug-ins of the services integrated in Ranger. The plug-ins can be used to set more fine-grained permissions. For details about functions of main operations you can perform on the page, see [Table 24-1](#).

Table 24-1 Functions of each operation portal on the Ranger page

Portal	Function
Access Manager	You can view the permission management plug-ins of each service integrated in Ranger. The plug-ins can be used to set more fine-grained permissions. For details, see Configuring Component Permission Policies .
Audit	You can view the audit logs related to Ranger running and permission control. For details, see Viewing Ranger Audit Information .
Security Zone	Ranger administrators can divide resources of each component into multiple security zones where different Ranger administrators set security policies for specified resources of services to facilitate management. For details, see Configuring a Security Zone .

Portal	Function
Settings	You can view Ranger permission settings, such as users, user groups, and roles. For details, see Viewing Ranger Permission Information .

24.2 Enabling Ranger Authentication

Scenario

This section guides you how to enable Ranger authentication. Ranger authentication is enabled by default in security mode and disabled by default in normal mode.

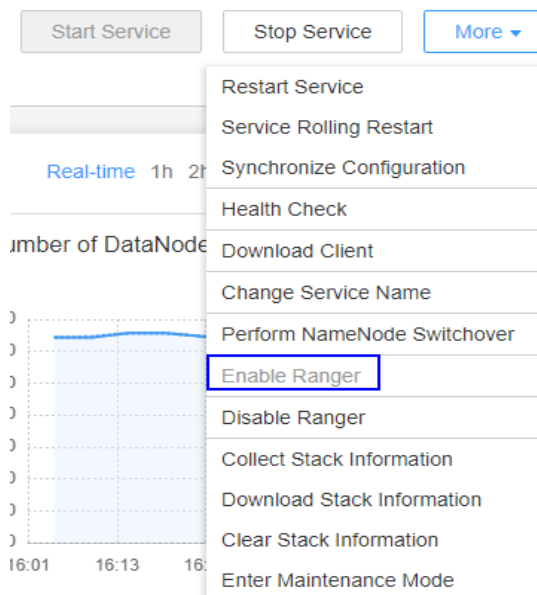
Procedure

- Step 1** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > **Services** > *Name of the service for which Ranger authentication is enabled*.
- Step 2** In the upper right corner of the **Dashboard** page, click **More** and select **Enable Ranger**. In the displayed dialog box, enter the password and click **OK**. After the operation is successful, click **Finish**.

NOTE

- If **Enable Ranger** is dimmed, Ranger authentication is enabled. See [Figure 24-1](#).
- 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.

Figure 24-1 Enabling Ranger Authentication



Step 3 Perform a rolling service restart or restart the service.

----End

24.3 Configuring Component Permission Policies

In the newly installed MRS cluster, Ranger is installed by default, with the Ranger authentication model enabled. The Ranger administrator can set fine-grained security policies for accessing component resources through the component permission plug-ins.

Currently, the following components in a cluster in security mode support Ranger: HDFS, Yarn, HBase, Hive, Spark2x, Kafka, Storm..

Configuring User Permission Policies Using Ranger

Step 1 Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).

Step 2 In the **Service Manager** area on the Ranger homepage, click the permission plug-in name of a component. The page for security access policy list of the component is displayed.

 **NOTE**

In the policy list of each component, many items are generated by default to ensure the permissions of some default users or user groups (such as the **supergroup** user group). Do not delete these items. Otherwise, the permissions of the default users or user groups are affected.

Step 3 Click **Add New Policy** and configure resource access policies for related users or user groups based on the service scenario plan.

The following policies are examples for different components:

- [Adding a Ranger Access Permission Policy for HDFS](#)
- [Adding a Ranger Access Permission Policy for HBase](#)
- [Adding a Ranger Access Permission Policy for Hive](#)
- [Adding a Ranger Access Permission Policy for Yarn](#)
- [Adding a Ranger Access Permission Policy for Spark2x](#)
- [Adding a Ranger Access Permission Policy for Kafka](#)
- [Adding a Ranger Access Permission Policy for Storm](#)

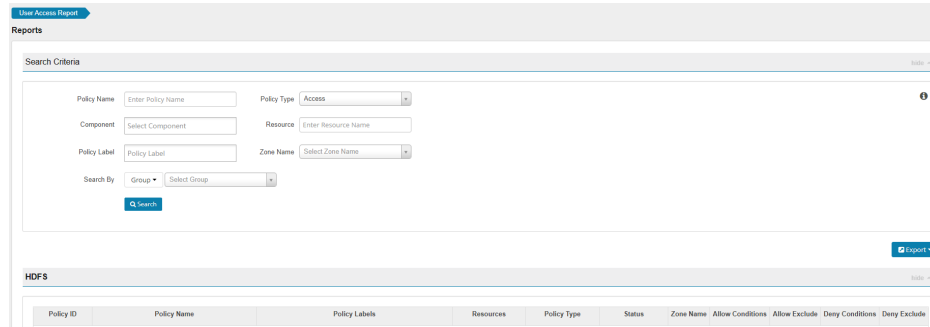
After the policies are added, wait for about 30 seconds for them to take effect.

 **NOTE**

Each time a component is started, the system checks whether the default Ranger service of the component exists. If the service does not exist, the system creates the Ranger service and adds a default policy for it. If a service is deleted by mistake, you can restart or restart the corresponding component service in rolling mode to restore the service. If the default policy is deleted by mistake, you can manually delete the service and then restart the component service.

Step 4 Choose **Access Manager > Reports** to view all security access policies of each component.

If there are many system policies, filter and search for policies by the policy name, policy type, component, resource, policy label, security zone, user, or user group. Alternatively, click **Export** to export related policies.



NOTE

- Generally, only one policy can be configured for a fixed resource object. If multiple policies are configured for the same resource object, the policies cannot be saved.
- For details about the priorities of different policies, see [Condition Priorities of the Ranger Permission Policy](#).

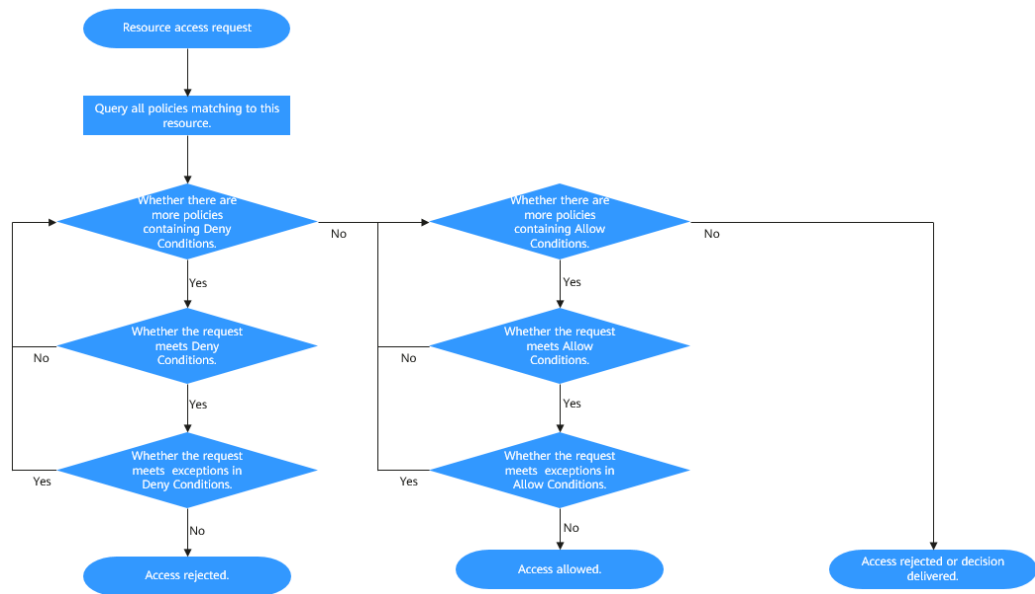
----End

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.

24.4 Viewing Ranger Audit Information

Ranger administrators can view audit logs about Ranger running and permission control audit logs after Ranger is used by components for authentication.

Viewing Ranger Audit Information

- Step 1** Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).
- Step 2** Choose **Audit** to view the audit information. For details about each tab page, see [Table 24-2](#). If there are a large number of audit records, you can filter them in the search box by keyword.

Table 24-2 Audit information

Tab	Description
Access	Audit information about users' access to components after Ranger authentication.
Admin	Audit information about operations on Ranger, such as creating, updating, and deleting security access policies, creating and deleting component permission policies, and creating, updating, and deleting roles.
Login Sessions	Session audit information about users who log in to Ranger.
Plugins	Component permission policy information in Ranger.

Tab	Description
Plugin Status	Audit information about synchronization of the permission policy of each component node.
User Sync	Audit information about synchronization between Ranger and LDAP users.

----End

24.5 Configuring a Security Zone

Security zone can be configured using Ranger. Ranger administrators can divide resources of each component into multiple security zones where Ranger administrators set security policies for specified resources in the zones to facilitate management. Policies defined in a security zone apply only to resources in the zone. After service resources are allocated to the security zone, the access permission policies for the resources in the non-security zone do not take effect. The administrator of a security zone can set policies only in the security zone that the administrator belongs to.

Adding a Security Zone

Step 1 Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).


Step 2 Click **Security Zone**. On the zone list page, click  to add a zone.

Table 24-3 Parameters for configuring a security zone

Parameter	Description	Example Value
Zone Name	Security zone	test
Zone Description	Description of the security zone	-
Admin Users/ Admin Usergroups	Management users and user groups in a security zone. You can add and modify permission policies for related resources in the security zone. At least one user or user group must be configured.	zone_admin
Auditor Users/ Auditor Usergroups	Audit users or user groups to be added. You can view the resource permission policies in the security zone. At least one user or user group must be configured.	zone_user
Select Tag Services	Tag information of a service	-

Parameter	Description	Example Value
Select Resource Services	Services and resources in a security zone. After selecting a service, you need to add specific resource objects in the Resource column, such as the file directories of the HDFS server, Yarn queues, Hive databases and tables, and HBase tables and columns.	/ testzone

For example, to create a security zone for the `/testzone` directory in HDFS, the configuration is as follows:

Zone Details :

Zone Name *

Zone Description

Zone Administration :

Admin Users

Admin Usergroups

Auditor Users

Auditor Usergroups

Services :

Select Tag Services

Select Resource Services *

Service Name	Service Type	Resource
hacluster	HDFS	<input type="text" value="path: /testzone"/> <input type="button" value="✓"/> <input type="button" value="✗"/> <input type="button" value="+"/>

Step 3 Click **Save** and wait until the security zone is added successfully.


The Ranger administrator can view all security zones on the **Security Zone** page and click **Edit** to modify the attributes of a security zone. If resources do not need to be managed in a security zone, the Ranger administrator can click **Delete** to delete the security zone.

----End

Configuring Permission Policies in a Security Zone

Step 1 Log in to the Ranger management page as the Ranger administrator of a security zone.

Step 2 Select a security zone from the **Security Zone** drop-down list in the upper right corner of the Ranger home page to switch to the permission view of the security zone.



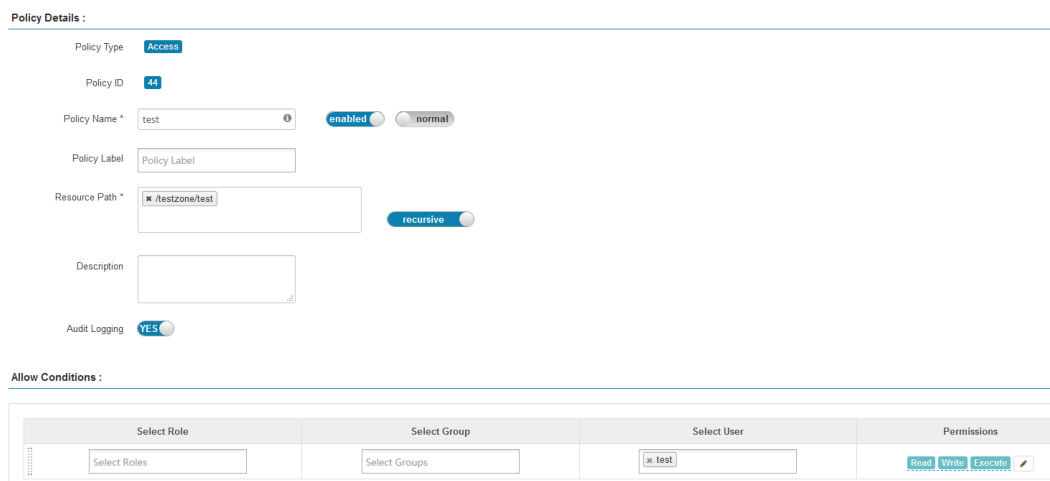
Step 3 Click the permission plug-in name of a component. The page for security access policy list of the component is displayed.

NOTE

In the policy list of each component, the default items generated by the system are automatically inherited to the security zone to ensure the permissions of some default users or user groups in the cluster.

Step 4 Click **Add New Policy** and configure resource access policies for related users or user groups based on the service scenario plan.

In this example, a policy that allows user test to access the **/testzone/test** directory is configured in the security zone.



The following access policies are examples for different components:

- [Adding a Ranger Access Permission Policy for HDFS](#)
- [Adding a Ranger Access Permission Policy for HBase](#)
- [Adding a Ranger Access Permission Policy for Hive](#)
- [Adding a Ranger Access Permission Policy for Yarn](#)
- [Adding a Ranger Access Permission Policy for Spark2x](#)
- [Adding a Ranger Access Permission Policy for Kafka](#)

- **Adding a Ranger Access Permission Policy for Storm**

After the policies are added, wait for about 30 seconds for them to take effect.

 **NOTE**

- Policies defined in a security zone apply only to resources in the zone. After service resources are allocated to the security zone, the access permission policies for the resources in the non-security zone do not take effect.
- To configure access policies for resources outside the current security zone, click **Security Zone** in the upper right corner of the Ranger homepage to exit the current security zone.

----End

24.6 Viewing Ranger Permission Information

You can view Ranger permission settings, such as users, user groups, and roles.

Viewing Ranger Permission Information

Step 1 Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).

Step 2 Choose **Settings > Users/Groups/Roles** to view information about users, user groups, or roles in the system.

- **Users:** displays all user information synchronized from LDAP or OS to Ranger.
- **Groups:** displays information about all user groups and role information synchronized from LDAP or OS to Ranger.
- **Roles:** displays information about roles created in Ranger.

 **NOTE**

- The users, roles, user groups created on FusionInsight Manager are automatically synchronized to Ranger periodically. The default period is 300,000 milliseconds (5 minutes). After roles and user groups in FusionInsight Manager are synchronized to Ranger, they become user groups. Only roles and user groups that are associated with users can be automatically synchronized to Ranger.
- The role created on the Ranger page is a set of users or user groups, which is used to flexibly set the permission access policies of components. The role is different from that on FusionInsight Manager.

----End

Adjusting Ranger User Types

Step 1 Log in to the Ranger management page.

To change the Ranger user type, you must log in as an **admin** user. For details about the user types, see [Ranger User Type](#).

Step 2 Choose **Settings > Users/Groups/Roles**. In the list of users, click the name of the user whose type you want to change.

Step 3 Set **Select Role** to the type to be modified.

Step 4 Click **Save**.

----End

Creating a Ranger Role

Ranger administrators can flexibly configure permission access policies for components based on users, user groups, or roles. User and user group information is automatically synchronized from LDAP, and roles can be manually added.

Step 1 Log in to the Ranger management page.

Step 2 Choose **Settings > Users/Groups/Roles > Roles > Add New Role**.

Step 3 Enter the role name and description as prompted.

Step 4 Add users, user groups, and sub-roles to the role.

- In the **Users** area, select a created user in the system and click **Add Users**.
- In the **Groups** area, select a created user group and click **Add Group**.
- In the **Roles** area, select a created role in the system and click **Add Role**.

Users:

User Name	Is Role Admin	Action
test01	<input type="checkbox"/>	

Select User

Groups:

Group Name	Is Role Admin	Action
hadoop	<input type="checkbox"/>	

Select Group

Roles:

Role Name	Is Role Admin	Action
admin	<input type="checkbox"/>	

Select Role

Step 5 Click **Save**. The role is added.

NOTE

Added roles cannot be deleted but can be modified.

----End

24.7 Adding a Ranger Access Permission Policy for HDFS

Scenario

Ranger administrators can use Ranger to configure the read, write, and execution permissions on HDFS directories or files for HDFS users.

Prerequisites



- The Ranger service has been installed and is running properly.
- You have created users, user groups, or roles for which you want to configure permissions.

Procedure

- Step 1** Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).
- Step 2** On the homepage, click the component plug-in name in the **HDFS** area, for example, **hacluster**.
- Step 3** Click **Add New Policy** to add an HDFS permission control policy.
- Step 4** Configure the parameters listed in the table below based on the service demands.

Table 24-4 HDFS permission parameters

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10,192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.

Parameter	Description
Resource Path	<p>Resource path, which is the HDFS path folder or file to which the current policy applies. You can enter multiple values and use the wildcard (*), for example, <code>/test/*</code>.</p> <p>To enable a subdirectory to inherit the permission of its upper-level directory, enable the recursion function.</p> <p>If recursion is enabled for the parent directory and a policy is configured for the subdirectory, the subdirectory has the policies of both the parent directory and the subdirectory. If the policy of the parent directory conflicts with that of the subdirectory, the policy of the subdirectory prevails.</p> <ul style="list-style-type: none"> • non-recursive: recursion disabled • recursive: recursion enabled
Description	Policy description.
Audit Logging	Whether to audit the policy.
Allow Conditions	<p>Permission and exception conditions allowed by a policy. The priority of an exception condition is higher than that of a normal condition.</p> <p>In the Select Role, Select Group, and Select User columns, select the role, user group, or user to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, and click Add Permissions to add the corresponding permission.</p> <ul style="list-style-type: none"> • Read: permission to read data • Write: permission to write data • Execute: execution permission • Select/Deselect All: Select or deselect all. <p>If users or user groups in the current condition need to manage this policy, select Delegate Admin. These users or user groups will become the agent administrators. The agent administrators can update and delete this policy and create sub-policies based on the original policy.</p> <p>To add multiple permission control rules, click . To delete a permission control rule, click .</p> <p>Exclude from Allow Conditions: exception rules excluded from the allowed conditions</p>
Deny All Other Accesses	<p>Whether to reject all other access requests.</p> <ul style="list-style-type: none"> • True: All other access requests are rejected. • False: Deny Conditions can be configured.

Parameter	Description
Deny Conditions	Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is the same as that of Allow Conditions . The priority of the rejection condition is higher than that of the allowed conditions configured in Allow Conditions . Exclude from Deny Conditions: exception rules excluded from the denied conditions

For example, to add the write permission for the `/user/test` directory of user `testuser`, the configuration is as follows:

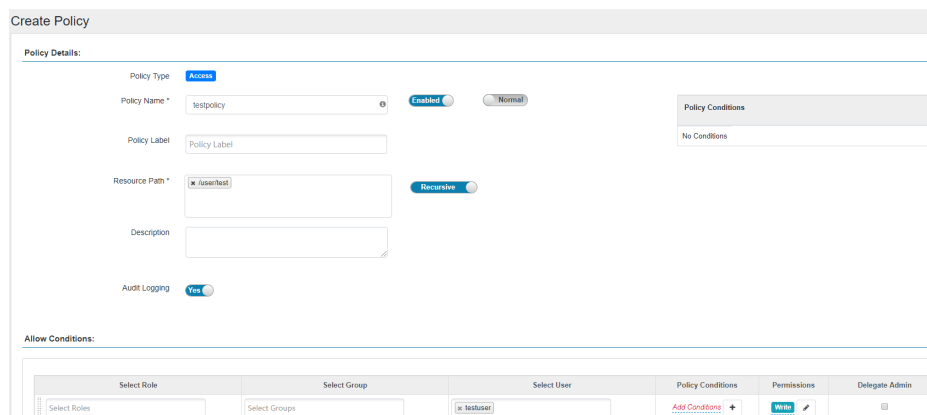





Table 24-5 Setting permissions

Task	Role Authorization
Setting the HDFS administrator permission	<ol style="list-style-type: none"> 1. On the homepage, click the component plug-in name in the HDFS area, for example, hacluster. 2. Select the policy whose Policy Name is all - path and click  to edit the policy. 3. In the Allow Conditions area, select a user from the Select User drop-down list.
Setting the permission for users to check and recover HDFS	<ol style="list-style-type: none"> 1. Add a folder or a file path in Resource Path. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Read and Execute.

Task	Role Authorization
Setting the permission for users to read directories or files of other users	<ol style="list-style-type: none"> 1. Add a folder or a file path in Resource Path. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Read and Execute.
Setting the permission for users to write data to files of other users	<ol style="list-style-type: none"> 1. Add a folder or a file path in Resource Path. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Write and Execute.
Setting the permission for users to create or delete sub-files or sub-directories in the directory of other users	<ol style="list-style-type: none"> 1. Add a folder or a file path in Resource Path. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Write and Execute.
Setting the permission for users to execute directories or files of other users	<ol style="list-style-type: none"> 1. Add a folder or a file path in Resource Path. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Execute.
Setting the permission for allowing subdirectories to inherit all permissions of their parent directories	<ol style="list-style-type: none"> 1. Add a folder or a file path in Resource Path. 2. Enable the recursion function. Recursive indicates that recursion is enabled.

Step 5 (Optional) Add the validity period of the policy. Click **Add Validity period** in the upper right corner of the page, set **Start Time** and **End Time**, and select **Time Zone**. Click **Save**. To add multiple policy validity periods, click . To delete a policy validity period, click .

Step 6 Click **Add** to view the basic information about the policy in the policy list. After the policy takes effect, check whether the related permissions are normal.

To disable a policy, click  to edit the policy and set the policy to **Disabled**.

If a policy is no longer used, click  to delete it.

----End

24.8 Adding a Ranger Access Permission Policy for HBase

Scenario

Ranger administrators can use Ranger to configure permissions on HBase tables, column families, and columns for HBase users.

Prerequisites

- The Ranger service has been installed and is running properly.
- You have created users, user groups, or roles for which you want to configure permissions.

Procedure

- Step 1** Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).
- Step 2** On the home page, click the component plug-in name in the **HBASE** area, for example, **HBase**.
- Step 3** Click **Add New Policy** to add an HBase permission control policy.
- Step 4** Configure the parameters listed in the table below based on the service demands.

Table 24-6 HBase permission parameters

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10,192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.

Parameter	Description
HBase Table	<p>Name of a table to which the policy applies.</p> <p>The value can contain wildcard (*). For example, table1:* indicates all tables in table1.</p> <p>The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.</p> <p>NOTE The value of hbase.rpc.protection of the HBase service plug-in on Ranger must be the same as that of hbase.rpc.protection on the HBase server. For details, see When an HBase Policy Is Added or Modified on Ranger, Wildcard Characters Cannot Be Used to Search for Existing HBase Tables.</p>
HBase Column-family	<p>Name of the column families to which the policy applies.</p> <p>The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.</p>
HBase Column	<p>Name of the column to which the policy applies.</p> <p>The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.</p>
Description	Policy description.
Audit Logging	Whether to audit the policy.




Parameter	Description
Allow Conditions	<p>Policy allowed condition. You can configure permissions and exceptions allowed by the policy.</p> <p>In the Select Role, Select Group, and Select User columns, select the role, user group, or user to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, and click Add Permissions to add the corresponding permission.</p> <ul style="list-style-type: none"> • Read: permission to read data • Write: permission to write data • Create: permission to create data • Admin: permission to manage data • Select/Deselect All: Select or deselect all. <p>If users or user groups in the current condition need to manage this policy, select Delegate Admin. These users or user groups will become the agent administrators. The agent administrators can update and delete this policy and create sub-policies based on the original policy.</p> <p>To add multiple permission control rules, click . To delete a permission control rule, click .</p> <p>Exclude from Allow Conditions: policy exception conditions</p>
Deny All Other Accesses	<p>Whether to reject all other access requests.</p> <ul style="list-style-type: none"> • True: All other access requests are rejected. • False: Deny Conditions can be configured.
Deny Conditions	<p>Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is similar to that of Allow Conditions.</p> <p>The priority of Deny Conditions is higher than that of allowed conditions configured in Allow Conditions.</p> <p>Exclude from Deny Conditions: exception rules excluded from the denied conditions</p>



Table 24-7 Setting permissions

Task	Role Authorization
Setting the HBase administrator permission	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the HBase area, for example, HBase. 2. Select the policy whose Policy Name is all - table, column-family, column and click  to edit the policy. 3. In the Allow Conditions area, select a user from the Select User drop-down list.
Setting the permission for users to create tables	<ol style="list-style-type: none"> 1. In HBase Table, specify a table name. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Create. 4. This user has the following permissions: create table drop table truncate table alter table enable table flush table flush region compact disable enable desc
Setting the permission for users to write data to tables	<ol style="list-style-type: none"> 1. In HBase Table, specify a table name. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Write. 4. The user has the put, delete, append, incr and bulkload operation permissions.
Setting the permission for users to read data from tables	<ol style="list-style-type: none"> 1. In HBase Table, specify a table name. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Read. 4. This user has the get and scan permissions.


Task	Role Authorization
Setting the permission for users to manage namespaces or tables	<ol style="list-style-type: none"> 1. In HBase Table, specify a table name. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Admin. 4. The user has the rsgroup, peer, assign and balance operation permissions.
Setting the permission for reading data from or writing data to columns	<ol style="list-style-type: none"> 1. In HBase Table, specify a table name. 2. In HBase Column-family, specify the column family name. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Read and Write.


 **NOTE**

If a user performs the **desc** operation in **hbase shell**, the user must be granted the read permission on the **hbase:quota** table.

Step 5 (Optional) Add the validity period of the policy. Click **Add Validity period** in the upper right corner of the page, set **Start Time** and **End Time**, and select **Time Zone**. Click **Save**. To add multiple policy validity periods, click . To delete a policy validity period, click .

Step 6 Click **Add** to view the basic information about the policy in the policy list. After the policy takes effect, check whether the related permissions are normal.

To disable a policy, click  to edit the policy and set the policy to **Disabled**.

If a policy is no longer used, click  to delete it.

----End

24.9 Adding a Ranger Access Permission Policy for Hive

Scenario

Ranger administrators can use Ranger to set permissions for Hive users. The default administrator account of Hive is **hive** and the initial password is **Hive@123**.

Prerequisites

- The Ranger service has been installed and is running properly.

- You have created users, user groups, or roles for which you want to configure permissions.
- The users must be added to the **hive** group.

Procedure

- Step 1** Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).
- Step 2** On the home page, click the component plug-in name in the **HADOOP SQL** area, for example, **Hive**.
- Step 3** On the **Access** tab page, click **Add New Policy** to add a Hive permission control policy.
- Step 4** Configure the parameters listed in the table below based on the service demands.

Table 24-8 Hive permission parameters

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10 , 192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
database	Name of the Hive database to which the policy applies. The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.
table	Name of the Hive table to which the policy applies. To add a UDF-based policy, switch to UDF and enter the UDF name. The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.
Hive Column	Name of the column to which the policy applies. The value * indicates all columns. The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.
Description	Policy description.
Audit Logging	Whether to audit the policy.


Parameter	Description
Allow Conditions	<p>Policy allowed condition. You can configure permissions and exceptions allowed by the policy.</p> <p>In the Select Role, Select Group, and Select User columns, select the role, user group, or user to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, and click Add Permissions to add the corresponding permission.</p> <ul style="list-style-type: none"> ● select: permission to query data ● update: permission to update data ● Create: permission to create data ● Drop: permission to drop data ● Alter: permission to alter data ● Index: permission to index data ● All: all permissions ● Read: permission to read data ● Write: permission to write data ● Temporary UDF Admin: temporary UDF management permission ● Select/Deselect All: Select or deselect all. <p>To add multiple permission control rules, click .</p> <p>If users or user groups in the current condition need to manage this policy, select Delegate Admin. These users will become the agent administrators. The agent administrators can update and delete this policy and create sub-policies based on the original policy.</p>
Deny Conditions	<p>Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is similar to that of Allow Conditions.</p>

Table 24-9 Setting permissions

Task	Role Authorization
<p>role admin operation</p>	<ol style="list-style-type: none"> 1. On the home page, click Settings and choose Roles. 2. Click the role with Role Name set to admin. In the Users area, click Select User and select a username. 3. Click Add Users, select Is Role Admin in the row where the username is located, and click Save. <p>NOTE Only user rangeradmin has the permission to access the Settings option on the Ranger page. After being bound to the Hive administrator role, perform the following operations during each maintenance operation:</p> <ol style="list-style-type: none"> 1. Log in to the node where the Hive client is installed as the client installation user. 2. Run the following command to configure environment variables: For example, if the Hive client installation directory is /opt/hiveclient, run source /opt/hiveclient/bigdata_env. 3. Run the following command to authenticate the user: kinit Hive service user 4. Run the following command to log in to the client tool: beeline 5. Run the following command to update the administrator permissions: set role admin;
<p>Creating a database table</p>	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter or select the corresponding database on the right side of database and enter or select * on the right side of column. (To create a table, enter or select the corresponding table on the right side of table.) 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Create.
<p>Deleting a table</p>	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter or select the corresponding database on the right side of database and enter and select * on the right side of column. (To delete a table, enter or select the corresponding table on the right side of table.) 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Drop.


Task	Role Authorization
Query operation (select , desc , and show)	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter or select the corresponding database on the right side of database and enter or select * (* indicates all columns) on the right side of column. (To create a table, enter or select the corresponding table on the right side of table.) 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select select.
Alter operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter and select the corresponding database on the right side of database and enter or select * on the right side of column. (For tables, enter or select the corresponding table on the right side of table.) 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Alter.
LOAD operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. On the right side of database, enter or select the corresponding database. On the right side of table, enter or select the corresponding table. On the right side of column, enter a column and select *. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select update.
INSERT and DELETE operations	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. On the right side of database, enter or select the corresponding database. On the right side of table, enter or select the corresponding table. On the right side of column, enter a column and select *. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select update. 5. Configure the submit permission on the Yarn task queue. For details about how to configure the permission, see Adding a Ranger Access Permission Policy for Yarn.

Task	Role Authorization
GRANT/REVOKE operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. On the right side of database, enter or select the corresponding database. On the right side of table, enter or select the corresponding table. On the right side of column, enter a column and select *. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Select Delegate Admin.
ADD JAR operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Click database, and select global from the drop-down list. On the right of global, enter related information or select *. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Temporary UDF Admin.
UDF operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter or select the corresponding database on the right of database, and enter the corresponding udf function name on the right of udf. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select required permissions for the user (udf supports the Create, select, and Drop permissions).
VIEW operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. On the right side of database, enter or select the corresponding database. On the right side of table, enter or select the corresponding table to be viewed. On the right side of column, enter a column and select *. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select permissions for the user as required.
dfs command operation	<p>The dfs operation can be performed only after you have run the set role admin command.</p>
Operations on other user database tables	<ol style="list-style-type: none"> 1. Perform the preceding operations to add the corresponding permissions. 2. Grant the read, write, and execution permissions on the HDFS paths of other user database tables to the user. For details, see Adding a Ranger Access Permission Policy for HDFS.

 NOTE

- If you have specified an HDFS path when running commands, you need to be granted with the read, write, and execution permissions on the HDFS paths. For details, see [Adding a Ranger Access Permission Policy for HDFS](#). You do not need to configure the Ranger policy of HDFS. You can use the Hive permission plug-in to add permissions to the role and assign the role to the corresponding user. If the HDFS Ranger policy can match the file or directory permission of the Hive database table, the HDFS Ranger policy is preferentially used.
- The URL policy in the Ranger policy is involved in the scenario where the Hive table is stored on OBS. Set the URL to the complete path of the object on OBS. The Read and Write permissions are used together with the URL. URL policies are not involved in other scenarios.
- The global policy in the Ranger policy is used only with the **Temporary UDF Admin** permission to control the upload of UDF packages.
- The **hiveservice** policy in the Ranger policy is used only with the **Service Admin** permission to control the permission to run the **kill query <queryId>** command to end the task that is being executed.
- The **lock**, **index**, **refresh**, and **replAdmin** permissions are not supported.
- Run the **show grant** command to view the table permission. The **grantor** column of the table **owner** is displayed as user **hive**. If the Ranger page is used or the **grant** command is used to grant permissions in the background, the **grantor** column is displayed as the corresponding user. To view the result of using the Hive permission plug-in, set **hive-ext.ranger.previous.privileges.enable** to **true** and run the **show grant** command.

Step 5 Click **Add** to view the basic information about the policy in the policy list. After the policy takes effect, check whether the related permissions are normal.

To disable a policy, click  to edit the policy and set the policy to **Disabled**.

If a policy is no longer used, click  to delete it.

----End


Hive Data Masking

Ranger supports data masking for Hive data. It can process the returned result of the **select** operation you performed to mask sensitive information.

Step 1 Log in to the Ranger web UI. Click **Hive** in the **HADOOP SQL** area on the homepage.




Step 2 On the **Masking** tab page, click **Add New Policy** to add a Hive permission control policy.



Step 3 Configure the parameters listed in the table below based on the service demands.

Table 24-10 Hive data masking parameters

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10 , 192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
Hive Database	Name of the Hive database to which the current policy applies.
Hive Table	Name of the Hive table to which the current policy applies.
Hive Column	Column name.
Description	Policy description.
Audit Logging	Whether to audit the policy.

Parameter	Description
Mask Conditions	<p>In the Select Role, Select Group, and Select User columns, select the object to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, then click Add Permissions, and select select. Click Select Masking Option and select a data masking policy.</p> <ul style="list-style-type: none"> • Redact: Use x to mask all letters and n to mask all digits. • Partial mask: show last 4: Only the last four characters are displayed, and the rest characters are displayed using x. • Partial mask: show first 4: Only the first four characters are displayed, and the rest characters are displayed using x. • Hash: Replace the original value with the hash value. The Hive built-in function mask_hash is used. This is valid only for fields of the string, character, and varchar types. NULL is returned for fields of other types. • Nullify: Replace the original value with the NULL value. • Unmasked (retain original value): Keep the original value. • Date: show only year: Only the year part of the date string is displayed, and the default month and date start from January and Monday (01/01). • Custom: You customize policies using any valid return data type which is the same as the data type in the masked column. <p>To add a multi-column masking policy, click .</p>

Step 4 Click **Add** to view the basic information about the policy in the policy list.

Step 5 After you perform the **select** operation on a table configured with a data masking policy on the Hive client, the system processes and displays the data.

 **NOTE**


To process data, you must have the permission to submit tasks to the Yarn queue.

----End


Hive Row-Level Data Filtering

Ranger allows you to filter data at the row level when you perform the **select** operation on Hive data tables.

Step 1 Log in to the Ranger web UI. Click **Hive** in the **HADOOP SQL** area on the homepage.




Step 2 On the **Row Level Filter** tab page, click **Add New Policy** to add a row data filtering policy.



Step 3 Configure the parameters listed in the table below based on the service demands.

Table 24-11 Parameters for filtering Hive row data

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10 , 192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
Hive Database	Name of the Hive database to which the current policy applies.
Hive Table	Name of the Hive table to which the current policy applies.
Description	Policy description.
Audit Logging	Whether to audit the policy.
Row Filter Conditions	<p>In the Select Role, Select Group, and Select User columns, select the object to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, then click Add Permissions, and select Select. Click Row Level Filter and enter data filtering rules.</p> <p>For example, if you want to filter the data in the zhangsan row in the name column of table A, the filtering rule is name <>'zhangsan'. For more information, see the official Ranger document.</p> <p>To add more rules, click .</p>

Step 4 Click **Add** to view the basic information about the policy in the policy list.

Step 5 After you perform the **select** operation on a table configured with a data masking policy on the Hive client, the system processes and displays the data.

 NOTE

To process data, you must have the permission to submit tasks to the Yarn queue.

----End

24.10 Adding a Ranger Access Permission Policy for Yarn

Scenario

Ranger administrators can use Ranger to configure YARN administrator permissions for YARN users, allowing them to manage YARN queue resources.

Prerequisites

- The Ranger service has been installed and is running properly.
- You have created users, user groups, or roles for which you want to configure permissions.

Procedure

- Step 1** Log in to FusionInsight Manager and choose **Cluster > Services > Yarn**.
- Step 2** On the page that is displayed, click the **Configuration** tab then the **All Configurations** sub-tab. On this sub-tab page, search for the **yarn.acl.enable** parameter, and change its value to **true**. If the value is **true**, no further action is required.
- Step 3** Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).
- Step 4** On the home page, click the component plug-in name in the **YARN** area, for example, **Yarn**.
- Step 5** Click **Add New Policy** to add a Yarn permission control policy.
- Step 6** Configure the parameters listed in the table below based on the service demands.

Table 24-12 Yarn permission parameters

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10,192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.







Parameter	Description
Queue	<p>Queue name. The wildcard (*) is supported.</p> <p>To enable a sub-queue to inherit the permission of its upper-level queue, enable the recursion function.</p> <ul style="list-style-type: none"> • Non-recursive: recursion disabled • Recursive: recursion enabled
Description	Policy description.
Audit Logging	Whether to audit the policy.
Allow Conditions	<p>Policy allowed condition. You can configure permissions and exceptions allowed by the policy.</p> <p>In the Select Role, Select Group, and Select User columns, select the role, user group, or user to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, and click Add Permissions to add the corresponding permission.</p> <ul style="list-style-type: none"> • submit-app: permission to submit queue tasks • admin-queue: permission to manage queue tasks • Select/Deselect All: Select or deselect all. <p>If users or user groups in the current condition need to manage this policy, select Delegate Admin. These users will become the agent administrators. The agent administrators can update and delete this policy and create sub-policies based on the original policy.</p> <p>To add multiple permission control rules, click . To delete a permission control rule, click .</p> <p>Exclude from Allow Conditions: policy exception conditions</p>
Deny All Other Accesses	<p>Whether to reject all other access requests.</p> <ul style="list-style-type: none"> • True: All other access requests are rejected. • False: Deny Conditions can be configured.
Deny Conditions	<p>Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is similar to that of Allow Conditions. The priority of Deny Conditions is higher than that of allowed conditions configured in Allow Conditions.</p> <p>Exclude from Deny Conditions: exception rules excluded from the denied conditions</p>


Table 24-13 Setting permissions

Task	Role Authorization
Setting the Yarn administrator permission	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the YARN area, for example, Yarn. 2. Select the policy whose Policy Name is all - queue and click  to edit the policy. 3. In the Allow Conditions area, select a user from the Select User drop-down list.
Setting the permission for a user to submit tasks in a specified Yarn queue	<ol style="list-style-type: none"> 1. In Queue, specify a queue name. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select submit-app.
Setting the permission for a user to manage tasks in a specified Yarn queue	<ol style="list-style-type: none"> 1. In Queue, specify a queue name. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select admin-queue.

Step 7 (Optional) Add the validity period of the policy. Click **Add Validity period** in the upper right corner of the page, set **Start Time** and **End Time**, and select **Time Zone**. Click **Save**. To add multiple policy validity periods, click . To delete a policy validity period, click .

Step 8 Click **Add** to view the basic information about the policy in the policy list. After the policy takes effect, check whether the related permissions are normal.

To disable a policy, click  to edit the policy and set the policy to **Disabled**.

If a policy is no longer used, click  to delete it.

----End

 **NOTE**

The permissions on Ranger Yarn are independent of each other. There is inclusion relationship among the permissions. Currently, the following permissions are supported:

- **submit-app**: permission to submit queue tasks
- **admin-queue**: permission to manage queue tasks

Although the **admin-queue** has the permission to submit tasks, it does not have the inclusion relationship with the **submit-app** permission.

24.11 Adding a Ranger Access Permission Policy for Spark2x

Scenario

Ranger administrators can use Ranger to set permissions for Spark2x users.

NOTE

1. After Ranger authentication is enabled or disabled on Spark2x, you need to restart Spark2x.
2. Download the client again or manually update the client configuration file *Client installation directory/Spark2x/spark/conf/spark-defaults.conf*.
Enable Ranger: `spark.ranger.plugin.authorization.enable=true`
Disable Ranger: `spark.ranger.plugin.authorization.enable=false`
3. In Spark2x, spark-beeline (applications connected to JDBCServer) supports the Ranger IP address filtering policy (**Policy Conditions** in the Ranger permission policy), while spark-submit and spark-sql do not.

Prerequisites

- The Ranger service has been installed and is running properly.
- The Ranger authentication function of the Hive service has been enabled. After the Hive service is restarted, the Spark2x service is restarted.
- You have created users, user groups, or roles for which you want to configure permissions.
- The created user has been added to the **hive** user group.

Procedure

Step 1 Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).

Step 2 On the home page, click the component plug-in name in the **HADOOP SQL** area, for example, **Hive**.



Step 3 On the **Access** tab page, click **Add New Policy** to add a Spark2x permission control policy.

The screenshot shows the 'Hive Policies' page in the Ranger web UI. It features a search bar and a table listing existing policies. The table has columns for Policy ID, Policy Name, Policy Labels, Default Policy, Status, Audit Logging, Roles, Groups, Users, and Action.

Policy ID	Policy Name	Policy Labels	Default Policy	Status	Audit Logging	Roles	Groups	Users	Action
9	all-database	--	True	Enabled	Enabled	admin	public, supergroup	rangeradmin, hive, OWNER	
10	all-flowservice	--	True	Enabled	Enabled	--	supergroup	rangeradmin, hive	

Step 4 Configure the parameters listed in the table below based on the service demands.

Table 24-14 Spark2x permission parameters

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10 , 192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
database	Name of the Spark2x database to which the policy applies. The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.
table	Name of the Spark2x table to which the policy applies. To add a UDF-based policy, switch to UDF and enter the UDF name. The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.
column	Name of the column to which the policy applies. The value * indicates all columns. The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.
Description	Policy description.
Audit Logging	Whether to audit the policy.


Parameter	Description
Allow Conditions	<p>Policy allowed condition. You can configure permissions and exceptions allowed by the policy.</p> <p>In the Select Role, Select Group, and Select User columns, select the role, user group, or user to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, and click Add Permissions to add the corresponding permission.</p> <ul style="list-style-type: none"> ● select: permission to query data ● update: permission to update data ● Create: permission to create data ● Drop: permission to drop data ● Alter: permission to alter data ● Index: permission to index data ● All: all permissions ● Read: permission to read data ● Write: permission to write data ● Temporary UDF Admin: temporary UDF management permission ● Select/Deselect All: Select or deselect all. <p>To add multiple permission control rules, click .</p> <p>If users or user groups in the current condition need to manage this policy, select Delegate Admin. These users will become the agent administrators. The agent administrators can update and delete this policy and create sub-policies based on the original policy.</p>
Deny Conditions	<p>Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is similar to that of Allow Conditions.</p>

Table 24-15 Setting permissions

Task	Operation
<p>role admin operation</p>	<ol style="list-style-type: none"> 1. On the home page, click Settings and choose Roles > Add New Role. 2. Set Role Name to admin. In the Users area, click Select User and select a username. 3. Click Add Users, select Is Role Admin in the row where the username is located, and click Save. <p>NOTE After being bound to the Hive administrator role, perform the following operations during each maintenance operation:</p> <ol style="list-style-type: none"> 1. Log in to the node where the Hive client is installed as the client installation user. 2. Run the following command to configure environment variables: For example, if the Spark2x client installation directory is /opt/client, run source /opt/client/bigdata_env. 3. Run the following command to perform user authentication: kinit Spark2xService user 4. Run the following command to log in to the client tool: spark-beeline 5. Run the following command to update the administrator permissions: set role admin;
<p>Creating a database table</p>	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter and select the corresponding database on the right of database. (If you want to create a database, enter the name of the database to be created or enter * to indicate a database with any name, and then select the name.) Enter and select the corresponding table name on the right of table and column. Wildcard characters (*) are supported. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Create.

Task	Operation
Deleting a table	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter and select the corresponding database on the right of database. (If you want to delete a database, enter the name of the database to be created or enter * to indicate a database with any name, and then select the name.) Enter and select the corresponding table name on the right of table and column. Wildcard characters (*) are supported. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Drop. <p>NOTE For CarbonData tables, only the owner of the corresponding database or table can perform the drop operation.</p>
ALTER operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter and select the corresponding database on the right of database, enter and select the corresponding table on the right of table, and enter and select the corresponding column name on the right of column. Wildcard characters (*) are supported. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Alter.
LOAD operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter and select the corresponding database on the right of database, enter and select the corresponding table on the right of table, and enter and select the corresponding column name on the right of column. Wildcard characters (*) are supported. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select update.

Task	Operation
INSERT operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter and select the corresponding database on the right of database, enter and select the corresponding table on the right of table, and enter and select the corresponding column name on the right of column. Wildcard characters (*) are supported. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select update. 5. The user also needs to have the submit-app permission of the Yarn task queue. By default, the Hadoop user group has the submit-app permission of all Yarn task queues. For details about how to load a network instance to a cloud connection, see Adding a Ranger Access Permission Policy for Yarn.
GRANT operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Enter and select the corresponding database on the right of database, enter and select the corresponding table on the right of table, and enter and select the corresponding column name on the right of column. Wildcard characters (*) are supported. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Select Delegate Admin.
ADD JAR operation	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. Click database, and select global from the drop-down list. On the right of global, enter related information and select *. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Temporary UDF Admin.


Task	Operation
VIEW and INDEX permissions	<ol style="list-style-type: none"> 1. Enter the policy name in Policy Name. 2. On the right side of database, enter the database name and select the corresponding database. (If you want to delete a database, enter the database name and select *.) On the right side of table, enter a table name and select the view and index names. On the right side of column, enter a Hive column name, and select *. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select permissions for the user as required.
Operations on other user database tables	<ol style="list-style-type: none"> 1. Perform the preceding operations to add the corresponding permissions. 2. Grant the read, write, and execution permissions on the HDFS paths of other user database tables to the current user. For details, see Adding a Ranger Access Permission Policy for HDFS.

 **NOTE**

After Spark SQL access policy is added on Ranger, you need to add the corresponding path access policies in the HDFS access policy. Otherwise, data files cannot be accessed. For details, see [Adding a Ranger Access Permission Policy for HDFS](#).

- The global policy in the Ranger policy is only used to associate with the **Temporary UDF Admin** permission to control the upload of UDF packages.
- When Ranger is used to control Spark SQL permissions, the **empower** syntax is not supported.

Step 5 Click **Add** to view the basic information about the policy in the policy list. After the policy takes effect, check whether the related permissions are normal.

To disable a policy, click  to edit the policy and set the policy to **Disabled**.

If a policy is no longer used, click  to delete it.

----End

Data Masking of the Spark2x Table

Ranger supports data masking for Spark2x data. It can process the returned result of the **select** operation you performed to mask sensitive information.


Step 1 Log in to the Ranger WebUI and click the component plug-in name, for example, **Hive**, in the **HADOOP SQL** area on the home page.

Step 2 On the **Masking** tab page, click **Add New Policy** to add a Spark2x permission control policy.

Step 3 Configure the parameters listed in the table below based on the service demands.

Table 24-16 Spark2x data masking parameters

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10 , 192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
Hive Database	Name of the Spark2x database to which the current policy applies.
Hive Table	Name of the Spark2x table to which the current policy applies.
Hive Column	Name of the Spark2x column to which the current policy applies.
Description	Policy description.
Audit Logging	Whether to audit the policy.

Parameter	Description
Mask Conditions	<p>In the Select Group and Select User columns, select the user group or user to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, then click Add Permissions, and select select.</p> <p>Click Select Masking Option and select a data masking policy.</p> <ul style="list-style-type: none"> ● Redact: Use x to mask all letters and n to mask all digits. ● Partial mask: show last 4: Only the last four characters are displayed. ● Partial mask: show first 4: Only the first four characters are displayed. ● Hash: Perform hash calculation for data. ● Nullify: Replace the original value with the NULL value. ● Unmasked(retain original value): The original data is displayed. ● Date: show only year: Only the year information is displayed. ● Custom: You can use any valid Hive UDF (returns the same data type as the data type in the masked column) to customize the policy. <p>To add a multi-column masking policy, click .</p>
Deny Conditions	<p>Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is similar to that of Allow Conditions.</p>


----End

Spark2x Row-Level Data Filtering

Ranger allows you to filter data at the row level when you perform the **select** operation on Spark2x data tables.

- Step 1** Log in to the Ranger WebUI and click the component plug-in name, for example, **Hive**, in the **HADOOP SQL** area on the home page.
- Step 2** On the **Row Level Filter** tab page, click **Add New Policy** to add a row data filtering policy.
- Step 3** Configure the parameters listed in the table below based on the service demands.

Table 24-17 Parameters for filtering Spark2x row data

Parameter	Description
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10,192.168.1.20 , or 192.168.1.* .
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
Hive Database	Name of the Spark2x database to which the current policy applies.
Hive Table	Name of the Spark2x table to which the current policy applies.
Description	Policy description.
Audit Logging	Whether to audit the policy.
Row Filter Conditions	<p>In the Select Role, Select Group, and Select User columns, select the object to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, then click Add Permissions, and select select. Click Row Level Filter and enter data filtering rules.</p> <p>For example, if you want to filter the data in the zhangsan row in the name column of table A, the filtering rule is name <>'zhangsan'. For more information, see the official Ranger document.</p> <p>To add more rules, click .</p>

Step 4 Click **Add** to view the basic information about the policy in the policy list.

Step 5 After you perform the **select** operation on a table configured with a data masking policy on the Spark2x client, the system processes and displays the data.

----End

24.12 Adding a Ranger Access Permission Policy for Kafka

Scenario

Ranger administrators can use Ranger to configure the read, write, and management permissions of the Kafka topic and the management permission of the cluster for the Kafka user. This section describes how to add the production permission of the **test** topic for the **test** user.

Prerequisites


- The Ranger service has been installed and is running properly.
- You have created users, user groups, or roles for which you want to configure permissions.

Procedure

- Step 1** Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).
- Step 2** On the home page, click the component plug-in name in the **KAFKA** area, for example, **Kafka**.
- Step 3** Click **Add New Policy** to add a Kafka permission control policy.
- Step 4** Configure the following parameters based on the service demands.

Table 24-18 Kafka permission parameters

Parameter	Description
Policy Type	Access type.
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10 , 192.168.1.20 , or 192.168.1.* .
Policy Name	Policy name, which can be customized and must be unique in the service.
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
topic	Name of the topic applicable to the current policy. You can enter multiple values. The value can contain wildcards, such as test , test* , and * . The Include policy applies to the current input object, and the Exclude policy applies to objects other than the current input object.
Description	Policy description.
Audit Logging	Whether to audit the policy.

Parameter	Description
Allow Conditions	<p>Permission and exception conditions allowed by a policy. The priority of an exception condition is higher than that of a normal condition.</p> <p>In the Select Role, Select Group, and Select User columns, select the role, user group, or user to which you want to assign permissions.</p> <p>Click Add Conditions, add the IP address range to which the policy applies, and click Add Permissions to add corresponding permissions.</p> <ul style="list-style-type: none"> ● Publish: production permission ● Consume: consumption permission ● Describe: query permission ● Create: topic creation permission ● Delete: topic deletion permission ● Describe Configs: configuration query permission ● Alter: permission to change the number of partitions of a topic. ● Alter Configs: configuration modification permission ● Select/Deselect All: Select or deselect all. <p>To add multiple permission control rules, click .</p> <p>If users or user groups in the current condition need to manage this policy, select Delegate Admin. These users will become the agent administrators. The agent administrators can update and delete this policy and create sub-policies based on the original policy.</p>
Deny Conditions	<p>Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is the same as that of Allow Conditions. The priority of the rejection condition is higher than that of the allowed conditions configured in Allow Conditions.</p>

For example, to add the production permission for the **test** topic of user **testuser**, configure the following information:

Figure 24-2 Kafka permission parameters







Table 24-19 Setting permissions




Scenario	Role Authorization
Setting the Kafka administrator permissions	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the KAFKA area, for example, Kafka. 2. Select the policy whose Policy Name is all - topic and click  to edit the policy. 3. In the Allow Conditions area, select a user from the Select User drop-down list. 4. Click Add Permissions and select Select/Deselect All.
Setting the permission for a user to create a topic	<ol style="list-style-type: none"> 1. Specify a topic name in topic. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Create. <p>NOTE Currently, the Kafka kernel supports the --zookeeper and --bootstrap-server methods to create topics. The --zookeeper method will be deleted from the community in later versions. Therefore, you are advised to use the --bootstrap-server method to create topics.</p> <p>Note: Currently, Kafka supports only the authentication of topic creation in --bootstrap-server mode and does not support that in --zookeeper mode.</p>


Scenario	Role Authorization
Setting the permission for a user to delete a topic	<ol style="list-style-type: none"> 1. Specify a topic name in topic. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Delete. <p>NOTE Currently, the Kafka kernel supports the --zookeeper and --bootstrap-server methods to delete topics. The --zookeeper method will be deleted from the community in later versions. Therefore, you are advised to use the --bootstrap-server method to delete topics.</p> <p>Note: Currently, Kafka supports only the authentication of topic deletion in --bootstrap-server mode and does not support that in --zookeeper mode.</p>
Setting the permission for a user to query a topic	<ol style="list-style-type: none"> 1. Specify a topic name in topic. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Describe and Describe Configs. <p>NOTE Currently, the Kafka kernel supports the --zookeeper and --bootstrap-server methods to query topics. The --zookeeper method will be deleted from the community in later versions. Therefore, you are advised to use the --bootstrap-server method to query topics.</p> <p>Note: Currently, Kafka supports only the authentication of topic query in --bootstrap-server mode and does not support that in --zookeeper mode.</p>
Setting the production permission of a user on a topic	<ol style="list-style-type: none"> 1. Specify a topic name in topic. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Publish.
Setting the consumption permission of a user on a topic	<ol style="list-style-type: none"> 1. Specify a topic name in topic. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Consume. <p>NOTE During topic consumption, offset management is involved. Therefore, the Consume permission of ConsumerGroup must be enabled at the same time. For details, see Setting a User's Permission to Submit ConsumerGroup Offsets.</p>
Setting the permission for a user to expand a topic (by adding partitions)	<ol style="list-style-type: none"> 1. Specify a topic name in topic. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Alter.



Scenario	Role Authorization
Setting the permission for a user to modify the topic configuration	Currently, the Kafka kernel does not support to modify topic parameters based on --bootstrap-server . Therefore, Ranger does not support authentication for this behavior.
Setting all the management permissions of a user on a cluster	<ol style="list-style-type: none"> 1. Enter a cluster name and select the cluster on the right side of cluster. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Kafka Admin.
Setting the permission for a user to create a cluster	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the KAFKA area, for example, Kafka. 2. Select the policy whose Policy Name is all - cluster and click  to edit the policy. 3. Enter a cluster name and select the cluster on the right side of cluster. 4. In the Allow Conditions area, select a user from the Select User drop-down list. 5. Click Add Permissions and select Create. <p>NOTE The authentication of the Create operation of a cluster involves the following two scenarios:</p> <ol style="list-style-type: none"> 1. After the auto.create.topics.enable parameter is enabled in the cluster, the client sends data to a topic that has not been created in the service. In this case, the system checks whether the user has the Create permission of the cluster. 2. If a user creates a large number of topics and is granted the Cluster Create permission, the user can create any topic in the cluster.
Setting the permission for a user to modify the cluster configuration	<ol style="list-style-type: none"> 1. Enter a cluster name and select the cluster on the right side of cluster. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Alter Configs. <p>NOTE The configuration modification permission allows you to modify the Broker and Broker Logger configurations. After the configuration modification permission is granted to a user, the user can query configuration details even if the user does not have the query permission. (The configuration modification permission includes the configuration query permission.)</p>

Scenario	Role Authorization
<p>Setting the permission for a user to query the cluster configuration</p>	<ol style="list-style-type: none"> 1. Enter a cluster name and select the cluster on the right side of cluster. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Describe and Describe Configs. <p>NOTE You can only query Broker and Broker Logger information in the cluster, excluding topics.</p>
<p>Setting the Idempotent Write permission in a cluster for a user</p>	<ol style="list-style-type: none"> 1. Enter a cluster name and select the cluster on the right side of cluster. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Idempotent Write. <p>NOTE This permission authenticates the Idempotent Produce behavior of the user's client.</p>
<p>Setting the permission to migrate partitions in a cluster for a user</p>	<ol style="list-style-type: none"> 1. Enter a cluster name and select the cluster on the right side of cluster. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Alter. <p>NOTE The Alter permission of a cluster can be used to control permissions in the following scenarios:</p> <ol style="list-style-type: none"> 1. In the Partition Reassign scenario, migrate the storage directory of replicas. 2. Elect a leader replica in each partition of the cluster. 3. Add or delete ACLs. <p>Operations in scenarios Step 4.1 and Step 4.2 are between a controller and broker and between brokers in the cluster. When a cluster is created, this permission is granted to the built-in Kafka user by default. It is meaningless for a common user to be granted with this permission.</p> <p>Scenario Step 4.3 involves the ACL management. ACLs are designed for authentication. Currently, Kafka authentication is hosted to Ranger. Therefore, this scenario is not involved (the configuration does not take effect).</p>


Scenario	Role Authorization
<p>Setting the Cluster Action permission in a cluster for a user</p>	<ol style="list-style-type: none"> 1. Enter a cluster name and select the cluster on the right side of cluster. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Cluster Action. <p>NOTE This permission controls the synchronization between the leader and follower replicas in the cluster and the communication between nodes. It has been granted to the built-in Kafka user during cluster creation. It is meaningless for a common user to grant this permission.</p>
<p>Setting the TransactionalId permission for a user</p>	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the KAFKA area, for example, Kafka. 2. Select the policy whose Policy Name is all - transactionalid and click  to edit the policy. <ol style="list-style-type: none"> 1. Set transactionalid to a transaction ID. 2. In the Allow Conditions area, select a user from the Select User drop-down list. 3. Click Add Permissions and select Publish and Describe. <p>NOTE The Publish permission is used to authenticate client requests for which the transaction feature is enabled, for example, starting and ending a transaction, submitting an offset, and generating transactional data. The Describe permission is used to authenticate the requests from the client and coordinator that have enabled the transaction feature. If the transaction feature is enabled, you are advised to grant both the Publish and Describe permissions to users.</p>


Scenario	Role Authorization
<p>Setting the DelegationToken permission for a user</p>	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the KAFKA area, for example, Kafka. 2. Select the policy whose Policy Name is all - delegationtoken and click  to edit the policy. 3. Set delegationtoken to a delegation token. 4. In the Allow Conditions area, select a user from the Select User drop-down list. 5. Click Add Permissions and select Describe. <p>NOTE Currently, Ranger only controls the query permission of DelegationToken, but does not control its create, renew, and expire permissions.</p>
<p>Setting the permission for a user to query ConsumerGroup Offsets</p>	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the KAFKA area, for example, Kafka. 2. Select the policy whose Policy Name is all - consumergroup and click  to edit the policy. 3. In consumergroup, configure the consumer group to be managed. 4. In the Allow Conditions area, select a user from the Select User drop-down list. 5. Click Add Permissions and select Describe.
<p>Set the user's submission permission on ConsumerGroup Offsets.</p>	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the KAFKA area, for example, Kafka. 2. Select the policy whose Policy Name is all - consumergroup and click  to edit the policy. 3. In consumergroup, configure the consumer group to be managed. 4. In the Allow Conditions area, select a user from the Select User drop-down list. 5. Click Add Permissions and select Consume. <p>NOTE After a user is granted with the Consume permission of ConsumerGroup, the user is also granted with the Describe permission.</p>

Scenario	Role Authorization
Setting the permission for a user to delete ConsumerGroup Offsets	<ol style="list-style-type: none"> 1. On the home page, click the component plug-in name in the KAFKA area, for example, Kafka. 2. Select the policy whose Policy Name is all - consumergroup and click  to edit the policy. 3. In consumergroup, configure the consumer group to be managed. 4. In the Allow Conditions area, select a user from the Select User drop-down list. 5. Click Add Permissions and select Delete. <p>NOTE When a user is granted with the Delete permission of ConsumerGroup, the user is also granted with the Describe permission.</p>

Step 5 (Optional) Add the validity period of the policy. Click **Add Validity period** in the upper right corner of the page, set **Start Time** and **End Time**, and select **Time Zone**. Click **Save**. To add multiple policy validity periods, click . To delete a policy validity period, click .

Step 6 Click **Add** to view the basic information about the policy in the policy list. After the policy takes effect, check whether the related permissions are normal.

To disable a policy, click  to edit the policy and set the policy to **Disabled**.

If a policy is no longer used, click  to delete it.

----End

24.13 Adding a Ranger Access Permission Policy for Storm

Scenario

Ranger administrators can use Ranger to set permissions for Storm users.

Prerequisites

- The Ranger service has been installed and is running properly.
- You have created users, user groups, or roles for which you want to configure permissions.
- The Ranger authentication function has been enabled on the page. The option in the following figure controls whether to enable the Ranger plug-in


for permission control. If the function is enabled, the Ranger authentication is used. Otherwise, the authentication mechanism of the component is used.



Procedure

- Step 1** Log in to the Ranger web UI as the Ranger administrator **rangeradmin**. For details, see [Logging In to the Ranger Web UI](#).
- Step 2** On the homepage, click **Storm** in the **STORM** area.
- Step 3** Click **Add New Policy** to add a Storm permission control policy.
- Step 4** Configure the parameters listed in the table below based on the service demands.


Table 24-20 Storm permission parameters


Parameter	Description
Policy Conditions	IP address filtering policy, which can be customized. You can enter one or more IP addresses or IP address segments. The IP address can contain the wildcard character (*), for example, 192.168.1.10 , 192.168.1.20 , or 192.168.1.* .
Policy Name	Policy name, which can be customized and must be unique in the service. The include policy applies to the current input object, and the exclude policy applies to objects other than the current input object.
Policy Label	A label specified for the current policy. You can search for reports and filter policies based on labels.
Storm Topology	Name of the topology to which the current policy applies. One or more values can be entered.
Description	Policy description.
Audit Logging	Whether to audit the policy.

Parameter	Description
Allow Conditions	<p>Policy allowed condition. You can configure permissions and exceptions allowed by the policy. In the Select Role, Select Group, and Select User columns, select the role, user group, or user to which the permission is to be granted, click Add Conditions, add the IP address range to which the policy applies, and click Add Permissions to add the corresponding permissions.</p> <ul style="list-style-type: none"> ● Submit Topology: Submit a topology. <p>NOTE The Submit Topology permission takes effect only when Storm Topology is set to *.</p> <ul style="list-style-type: none"> ● File Upload: Upload a file. ● File Download: Download a file. ● Kill Topology: Delete a topology. ● Rebalance: Perform the rebalance operation. ● Activate: Activate the topology permission. ● Deactivate: Deactivate the topology permission. ● Get Topology Conf: Obtain topology configurations. ● Get Topology: Obtain a topology. ● Get User Topology: Obtain user's topology. ● Get Topology Info: Obtain topology information. ● Upload New Credential: Upload a new credential. ● Select/Deselect All: Select or deselect all. <p>To add multiple permission control rules, click .</p> <p>If users or user groups in the current condition need to manage this policy, select Delegate Admin. These users will become the agent administrators. The agent administrators can update and delete this policy and create sub-policies based on the original policy.</p>
Deny Conditions	<p>Policy rejection condition, which is used to configure the permissions and exceptions to be denied in the policy. The configuration method is similar to that of Allow Conditions.</p>

Step 5 (Optional) Add the validity period of the policy. Click **Add Validity period** in the upper right corner of the page, set **Start Time** and **End Time**, and select **Time Zone**. Click **Save**. To add multiple policy validity periods, click . To delete a policy validity period, click .

Step 6 Click **Add** to view the basic information about the policy in the policy list. After the policy takes effect, check whether the related permissions are normal.

To disable a policy, click  to edit the policy and set the policy to **Disabled**.

If a policy is no longer used, click  to delete it.

----End

24.14 Ranger Log Overview

Log Description

Log path: The default storage path of Ranger logs is `/var/log/Bigdata/ranger/Role name`.

- RangerAdmin: `/var/log/Bigdata/ranger/rangeradmin` (run logs)
- TagSync: `/var/log/Bigdata/ranger/tagsync` (run logs)
- UserSync: `/var/log/Bigdata/ranger/usersync` (run logs)

Log archive rule: The automatic compression and archive function is enabled for Ranger logs. By default, when the size of a log file exceeds 20 MB, the log file is automatically compressed. The naming rule of the compressed log file is as follows: `<Original log file name>-<yyyy-mm-dd_hh-mm-ss>.[ID].log.zip`. A maximum of 20 compressed file are retained.

Table 24-21 Ranger log list

Type	Name	Description
RangerAdmin run log file	access_log.<DATE>.log	Tomcat access log
	catalina.out	Tomcat service run log
	gc-worker.log	RangerAdmin garbage collection (GC) log
	postinstallDetail.log	Work log generated after an instance is started before installation
	prestartDetail.log	Log that records preparations before instance startup

Type	Name	Description
	ranger-admin- <hostname>.log	RangerAdmin run log
	ranger_admin_sql- <hostname>.log	RangerAdmin log used to retrieve DBService
	startDetail.log	Instance startup log
TagSync run log	cleanupDetail.log	Instance clearing log
	gc-worker.log	GC log file of an instance
	postinstallDetail.log	Work log generated after an instance is started before installation
	prestartDetail.log	Log that records preparations before instance startup
	ranger-tagsync- <hostname>.log	TagSync run log
	startDetail.log	Instance startup log
	tagsync.out	TagSync run log
UserSync run log	auth.log	UnixAuth service run log
	cleanupDetail.log	Instance clearing log
	gc-worker.log	GC log file of an instance
	postinstallDetail.log	Work log generated after an instance is started before installation
	prestartDetail.log	Log that records preparations before instance startup
	ranger-usersync- <hostname>.log	UserSync run log
	startDetail.log	Instance startup log

Log Levels

Table 24-22 describes the log levels provided by HDFS. 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 specified level are printed. The number of printed logs decreases as the specified log level increases.

Table 24-22 Log levels

Level	Description
FATAL	Logs of this level record fatal error information about the current event processing that may result in a system crash.
ERROR	Logs of this level record error information about the current event processing, which indicates that system running is abnormal.
WARN	Logs of this level record abnormal information about the current event processing. These abnormalities will not result in system faults.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster > Services > Ranger > Configurations**.
- Step 3** Select **All Configurations**.
- Step 4** On the menu bar on the left, select the log menu of the target role.
- Step 5** Select a desired log level.
- Step 6** Click **Save**. In the displayed dialog box, click **OK** to make the configuration take effect.

 **NOTE**

The configurations take effect immediately without the need to restart the service.

----**End**

Log Formats

The following table lists the Ranger log formats.

Table 24-23 Log formats

Type	Format	Example Value
Run log	<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>	2020-04-29 20:09:28,543 INFO http-bio-21401- exec-56 Request comes from API call, skip cas filter. CasAuthenticationFilter- Wrapper.java:25

24.15 Common Issues About Ranger

24.15.1 Why Ranger Startup Fails During the Cluster Installation?

Problem

During cluster installation, Ranger fails to be started, and the error message "ERROR: cannot drop sequence X_POLICY_REF_ACCESS_TYPE_SEQ " is displayed in the task list of the Manager process. How do I resolve this problem and properly install Ranger?

Answer

This issue may occur when two RangerAmdin instances are installed. If the instance installation fails, manually restart one RangerAdmin instance and then restart the other instance.

24.15.2 How Do I Determine Whether the Ranger Authentication Is Used for a Service?

Question

How do I determine whether the Ranger authentication is enabled for a service that supports the authentication?

Answer

Log in to FusionInsight Manager and choose **Cluster** > **Services** > *Name of the desired service*. On the service details page, click **More** and check whether the **Enable Ranger** option is available.

- If yes, the Ranger authentication plug-in is not enabled for the service. You can click **Enable Ranger** to enable the function.
- If no, the Ranger authentication plug-in has been enabled for the service. You can configure the permission policy for accessing the service resources on the Ranger management page.

 NOTE

If this option does not exist, the current service does not support the Ranger authentication plug-in and Ranger authentication is disabled.

24.15.3 Why Cannot a New User Log In to Ranger After Changing the Password?

Question

When a new user logs in to Ranger, why is the 401 error reported after the password is changed?

Answer

The UserSync synchronizes user data at an interval of 5 minutes by default. Therefore, a new user created on Manager cannot log in to the Ranger before the user data is successfully synchronized because the Ranger database does not have the user information. The user can log in to the Ranger only after the specified interval ends.

In non-security mode, the Ranger does not synchronize user data from Manager. Therefore, only the **admin** user can log in to the Ranger page.

24.15.4 When an HBase Policy Is Added or Modified on Ranger, Wildcard Characters Cannot Be Used to Search for Existing HBase Tables

Question


When a Ranger access permission policy is added for HBase and wildcard characters are used to search for an existing HBase table in the policy, the table cannot be found. The following error is reported in **/var/log/Bigdata/ranger/rangeradmin/ranger-admin-*log**:

```
Caused by: javax.security.sasl.SaslException: No common protection layer between client and server
at com.sun.security.sasl.gsskerb.GssKrb5Client.doFinalHandshake(GssKrb5Client.java:253)
at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:186)
at
org.apache.hadoop.hbase.security.AbstractHBaseSaslRpcClient.evaluateChallenge(AbstractHBaseSaslRpcClient.java:142)
at org.apache.hadoop.hbase.security.NettyHBaseSaslRpcClientHandler
$2.run(NettyHBaseSaslRpcClientHandler.java:142)
at org.apache.hadoop.hbase.security.NettyHBaseSaslRpcClientHandler
$2.run(NettyHBaseSaslRpcClientHandler.java:138)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:422)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1761)
at
org.apache.hadoop.hbase.security.NettyHBaseSaslRpcClientHandler.channelRead0(NettyHBaseSaslRpcClientHandler.java:138)
at
org.apache.hadoop.hbase.security.NettyHBaseSaslRpcClientHandler.channelRead0(NettyHBaseSaslRpcClientHandler.java:42)
at
org.apache.hadoop.hbase.thirdparty.io.netty.channel.SimpleChannelInboundHandler.channelRead(SimpleChannelInboundHandler.java:105)
```

```
at  
org.apache.hbase.thirdparty.io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractCh  
annelHandlerContext.java:362)
```

Answer

The value of **hbase.rpc.protection** of the HBase service plug-in on Ranger must be the same as that of **hbase.rpc.protection** on the HBase server.

- Step 1** Log in to the Ranger management page. For details, see [Logging In to the Ranger Web UI](#).
- Step 2** In the **HBASE** area on the home page, click the component plug-in name, for example, the  button of HBase.
- Step 3** Search for the configuration item **hbase.rpc.protection** and change its value to the value of **hbase.rpc.protection** on the HBase server.
- Step 4** Click **Save**.

----End

25 Using Spark

25.1 Precautions

This section applies to versions earlier than MRS 3.x.

25.2 Getting Started with Spark

This section describes how to use Spark to submit a SparkPi job. SparkPi, a typical Spark job, is used to calculate the value of Pi (π).

Procedure

Step 1 Prepare the SparkPi program.

Multiple open source Spark sample programs are provided, including SparkPi. Click <https://archive.apache.org/dist/spark/spark-2.1.0/spark-2.1.0-bin-hadoop2.7.tgz> to download the Spark program package.

Decompress the software package to obtain the `spark-examples_2.11-2.1.0.jar` file, the sample program package, in the `spark-2.1.0-bin-hadoop2.7/examples/jars` directory. The `spark-examples_2.11-2.1.0.jar` sample program package contains the SparkPi program.

Step 2 Upload data to OBS.

1. Log in to OBS Console.
2. Choose **Parallel File System** > **Create Parallel File System** to create a file system named **sparkpi**.
sparkpi is only an example. The file system name must be globally unique. Otherwise, the parallel file system fails to be created. Use the default values for other parameters.
3. Click the file system name **sparkpi** and click **Files**.
4. Click **Create Folder** to create the **program** folder..
5. Go to the **program** folder, click **Upload Object**, select the program package downloaded in **Step 1** from the local PC, and set **Storage Class** to **Standard**.

Step 3 Log in to the MRS console. In the left navigation pane, choose **Clusters > Active Clusters**, and click a cluster name.

Step 4 Submit the SparkPi job.

On the details page of the MRS cluster, click the **Jobs** tab. On this tab page, click **Create**. The **Create Job** dialog box is displayed.

- Set **Type** to **SparkSubmit**.
- Set **Name** to **sparkPi**.
- Set **Program Path** to the path where programs are stored on OBS, for example, **obs://sparkpi/program/spark-examples_2.11-2.1.0.jar**.
- In **Program Parameter**, select **--class** for **Parameter** and set **Value** to **org.apache.spark.examples.SparkPi**.
- Set **Parameters** to **10**.
- Leave **Service Parameter** blank.

A job can be submitted only when the cluster is in the **Running** state.

After a job is submitted successfully, it is in the **Accepted** state by default. You do not need to manually execute the job.

Step 5 View the job execution result.

1. Go to the **Jobs** tab page and view job execution status.

The job execution takes a while. After the jobs are complete, refresh the job list.

Once a job has succeeded or failed, you cannot execute it again. However, you can add or copy a job, and set job parameters to submit a job again.

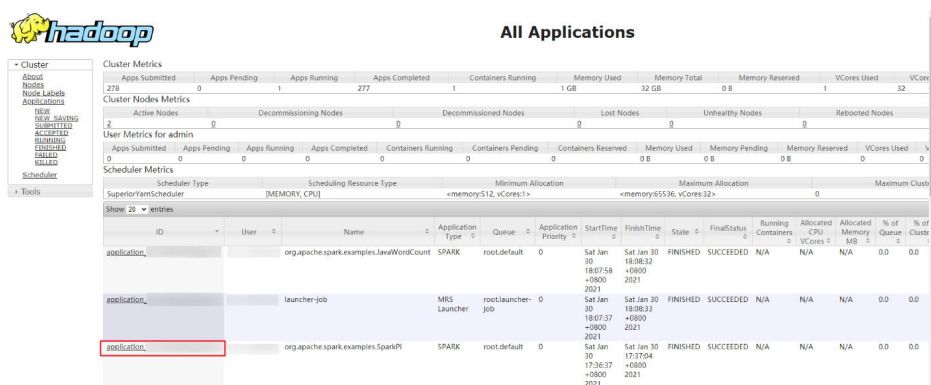
2. Go to the native Yarn page and view the job output information.

a. On the **Jobs** tab page, locate the row that contains the target job and click **View Details** in the **Operation** column to obtain the actual job ID.

b. Log in to Manager and choose **Services > Yarn > ResourceManager WebUI > ResourceManager (Active)**. The Yarn page is displayed.

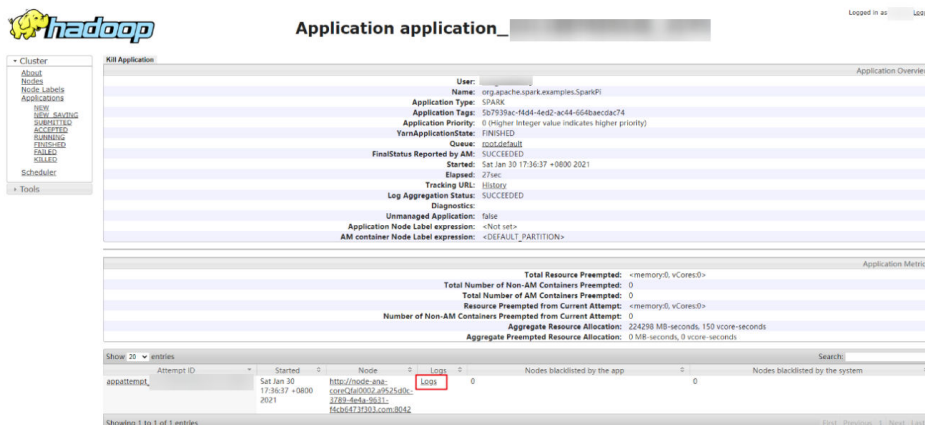
c. Click the ID corresponding to the actual job ID.

Figure 25-1 Yarn web UI




d. Click **Logs** in the job log area.

Figure 25-2 sparkPi job logs




- e. Click [here](#) to obtain more detailed logs.

Figure 25-3 More detailed logs of sparkPi jobs



- f. Obtain the job execution result.

Figure 25-4 sparkPi job execution result



----End

25.3 Getting Started with Spark SQL

Spark provides the Spark SQL language that is similar to SQL to perform operations on structured data. This section describes how to use Spark SQL from scratch. Create a table named `src_data`, write a data record in each row of the table, and store the data in the `mrs_20160907` cluster. Then use SQL statements to query data in the table, and delete the table at last.

Prerequisites

You have obtained the AK/SK for writing data from an OBS data source to a Spark SQL table. To obtain it, perform as follows:

1. Log in to the management console.
2. Click the username and select **My Credentials** from the drop-down list.

3. On the displayed **My Credentials** page, click **Access Keys**.
4. Click **Create Access Key** to switch to the **Create Access Key** dialog box.
5. Enter the password and , and click **OK** to download the access key. Keep the access key secure.

Procedure

Step 1 Prepare data sources for Spark SQL analysis.

The sample text file is as follows:

```
abcd3ghji  
efgh658ko  
1234jjyu9  
7h8kodfg1  
kk99icxz3
```

Step 2 Upload data to OBS.

1. Log in to OBS Console.
2. Choose **Parallel File System > Create Parallel File System** to create a file system named **sparksql**.
sparksql is only an example. The file system name must be globally unique. Otherwise, the parallel file system fails to be created.
3. Click the name of the **sparksql** file system and click **Files**.
4. Click **Create Folder** to create the **input** folder.
5. Go to the **input** folder, choose **Upload File > add file**, select the local TXT file, and click **Upload**.

Step 3 Log in to the MRS console. In the left navigation pane, choose **Clusters > Active Clusters**, and click a cluster name.

Step 4 Import the text file from OBS to HDFS.

1. Click the **Files** tab.
2. On the **HDFS File List** tab page, click **Create Folder**, and create a folder named **userinput**.
3. Go to the **userinput** folder, and click **Import Data**.
4. Select the OBS and HDFS paths and click **OK**.

OBS Path: `obs://sparksql/input/sparksql-test.txt`

HDFS Path: `/user/userinput`

Step 5 Submit the SQL statement.

1. On the details page of the MRS cluster, click the **Jobs** tab.
A job can be submitted only when the **mrs_20160907** cluster is in the **Running** state.
2. Enter the Spark SQL statement for table creation.
When entering Spark SQL statements, ensure that the statement characters are not more than 10,000.

Syntax:

```
CREATE [EXTERNAL] TABLE [IF NOT EXISTS] table_name [(col_name  
data_type [COMMENT col_comment], ...)] [COMMENT table_comment]
```

```
[PARTITIONED BY (col_name data_type [COMMENT col_comment], ...)]  
[CLUSTERED BY (col_name, col_name, ...) [SORTED BY (col_name [ASC]  
DESC], ...)] INTO num_buckets BUCKETS] [ROW FORMAT row_format]  
[STORED AS file_format] [LOCATION hdfs_path];
```

You can use the following two methods to create a table example:

- Method 1: Create table **src_data** and write data in every row.
 - The data source is stored in the folder of HDFS: **create external table src_data(line string) row format delimited fields terminated by '\n' stored as textfile location '/user/userinput'**;
 - The data source is stored in the **/sparksql/input** folder of OBS: **create external table src_data(line string) row format delimited fields terminated by '\\n' stored as textfile location 'obs://AK:SK@sparksql/input'**;
For details about how to obtain the AK/SK, see [Prerequisites](#).
- Method 2: Create table **src_data1** and load data to the table in batches.
create table src_data1 (line string) row format delimited fields terminated by ',' ;
load data inpath '/user/userinput/sparksql-test.txt' into table src_data1;

 NOTE

When method 2 is used, the data from OBS cannot be loaded to the created tables directly.

3. Enter the Spark SQL statement for table query.

Syntax:

```
SELECT col_name FROM table_name;
```

Example of querying all data in the **src_data** table:

```
select * from src_data;
```

4. Enter the Spark SQL statement for table deletion.

Syntax:

```
DROP TABLE [IF EXISTS] table_name;
```

Example of deleting the **src_data** table:

```
drop table src_data;
```

5. Click **Check** to check the statement correctness.
6. Click **OK**.

After the Spark SQL statements are submitted, the statement execution results are displayed in the result column.

Step 6 Delete the cluster.

----End

25.4 Using the Spark Client

After an MRS cluster is created, you can create and submit jobs on the client. The client can be installed on nodes inside or outside the cluster.

- Node in the cluster: After an MRS cluster is created, the client has been installed on the master and core nodes in the cluster by default.
- Node outside the cluster: You can install the client on a node outside the cluster.

Using the Spark Client

Step 1 Log in to the node where the client is installed.

Step 2 Run the following command to go to the client installation directory:

```
cd /opt/client
```

Step 3 Run the following command to configure environment variables:

```
source bigdata_env
```

Step 4 If the cluster is in security mode, run the following command to authenticate the user. In normal mode, user authentication is not required.

```
kinit Component service user
```

Step 5 Run the Spark shell command. The following provides an example:

```
spark-beeline
```

```
----End
```

25.5 Accessing the Spark Web UI

The Spark web UI is used to view the running status of Spark applications. Google Chrome is recommended for better user experience.

Spark has two web UIs.

- Spark UI: used to display the status of running applications.
The UI includes the following parts: Jobs, Stages, Storage, Environment, Executors, SQL, and JDBC/ODBC Server. The Streaming application has the Streaming tab in addition to the preceding parts.
- History Server UI: used to display the status of Spark applications that are complete or incomplete.
The UI includes the application ID, application name, start time, end time, execution time, and owner information.

Spark UI

Step 1 Access the component management page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components**.

NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 2 Select **Yarn**. In the **Yarn Summary** area, click **ResourceManager** in **ResourceManager Web UI** to access the web UI.

Step 3 Locate the Spark application. Click **ApplicationMaster** in the last column of the application information. The Spark UI is displayed.

Figure 25-5 ApplicationMaster

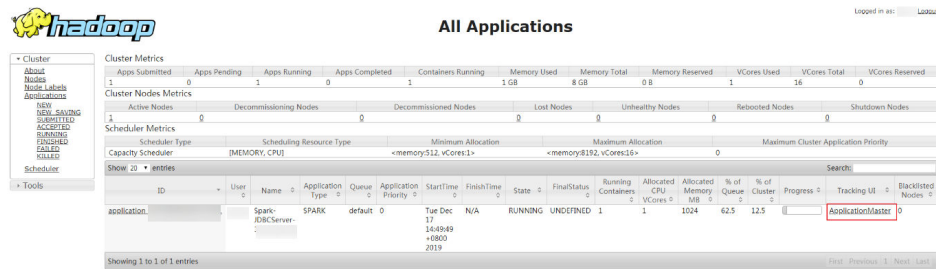


Figure 25-6 Spark UI



----End

History Server

Step 1 Access the component management page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components**.


NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 2 Select **Spark**. In the **Spark Summary** area, click **JobHistory** corresponding to **Spark Web UI** to access the web UI.

Figure 25-7 Spark History Server



----End

25.6 Interconnecting Spark with OpenTSDB

25.6.1 Creating a Table and Associating It with OpenTSDB

Function

MRS Spark can be used to access the data source of OpenTSDB, create and associate tables in the Spark, and query and insert the OpenTSDB data.

Use the **CREATE TABLE** command to create a table and associate it with an existing metric in OpenTSDB.

NOTE

If no metric exists in OpenTSDB, an error will be reported when the corresponding table is queried.

Syntax

```
CREATE TABLE [IF NOT EXISTS] OPENTSDB_TABLE_NAME USING OPENTSDB OPTIONS (
'metric' = 'METRIC_NAME',
'tags' = 'TAG1,TAG2'
);
```

Keyword

Parameter	Description
metric	Indicates the name of the metric in OpenTSDB corresponding to the table to be created.
tags	Indicates the tags corresponding to the metric. The tags are used for classification, filtering, and quick retrieval. You can set 1 to 8 tags, which are separated by commas (.). The parameter value includes values of all tagKs in the corresponding metric.

Precautions

When creating a table, you do not need to specify the **timestamp** and **value** fields. The system automatically builds the following fields based on the specified tags. The fields **TAG1** and **TAG2** are specified by tags.

- TAG1 String
- TAG2 String
- timestamp Timestamp
- value double

Example

Create table **opentsdb_table** and associate it with metric **city.temp** of the OpenTSDB component.

```
CREATE table opentsdb_table using opentsdb OPTIONS ('metric='city.temp', 'tags'='city,location');
```

25.6.2 Inserting Data to the OpenTSDB Table

Function

Run the **INSERT INTO** statement to insert the data in the table to the associated OpenTSDB metric.

Syntax

```
INSERT INTO TABLE_NAME SELECT * FROM SRC_TABLE;  
INSERT INTO TABLE_NAME VALUES(XXX);
```

Keyword

Parameter	Description
TABLE_NAME	Indicates the name of the associated OpenTSDB table.
SRC_TABLE	Indicates the name of the table from which data is obtained. This parameter can be set to a name of a common table.

Precautions

- The inserted data cannot be **null**. If the inserted data is the same as the original data or only the **value** is different, the inserted data overwrites the original data.
- **INSERT OVERWRITE** is not supported.
- You are advised not to concurrently insert data into a table. If you concurrently insert data into a table, there is a possibility that conflicts occur, leading to data insertion failures.
- The **TIMESTAMP** format supports only yyyy-MM-dd hh:mm:ss.

Example

Insert data into table **opentsdb_table**.

```
insert into opentsdb_table values('city1','city2','2018-05-03 00:00:00',21);
```

25.6.3 Querying an OpenTSDB Table

This **SELECT** command is used to query data in an OpenTSDB table.

Syntax

```
SELECT * FROM table_name WHERE tagk=tagv LIMIT number;
```


Keyword

Parameter	Description
LIMIT	Used to limit the query results.
number	Only the INT type is supported.

Precautions

- The to-be-queried table must exist. Otherwise, an error is reported.
- The value of **tagv** must exist. Otherwise, an error occurs.

Example

Query data in the **opentsdb_table** table.

```
SELECT * FROM opentsdb_table LIMIT 100;
SELECT * FROM opentsdb_table WHERE city='city1';
```

25.6.4 Modifying the Default Configuration Data

By default, OpenTSDB connects to the local TSD process of the node where the Spark executor resides. In MRS, use the default configuration.

Table 25-1 OpenTSDB data source configuration

Parameter	Description	Example Value
spark.sql.datasource.opentsdb.host	Indicates the IP address of the connected TSD process.	Null (default value) xx.xx.xx.xx indicates the IP address. Separate multiple IP addresses with commas (,).
spark.sql.datasource.opentsdb.port	Indicates the port number of the TSD process.	4242 (default value)

spark.sql.datasource.opentsdb.randomSeed	Indicates whether to use the random seed when the spark.sql.datasource.opentsdb.host is set to multiple addresses. If this parameter is set to false , all executors on the same node are connected to the same host. In this way, spark.blacklist.enabled=true can be used to implement task fault tolerance.	false (default value)
--	---	-----------------------

Example

Run the **set** statement in **spark-sql** and **spark-beeline**, and then run other SQL statements.

```
set spark.sql.datasource.opentsdb.host = 192.168.2.143,192.168.2.158;  
SELECT * FROM opentsdb_table ;
```

26 Using Spark2x

26.1 Precautions

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

26.2 Basic Operation

26.2.1 Getting Started

This section describes how to use Spark2x to submit Spark applications, including Spark Core and Spark SQL. Spark Core is the kernel module of Spark. It executes tasks and is used to compile Spark applications. Spark SQL is a module that executes SQL statements.

Scenario Description

Develop a Spark application to perform the following operations on logs about netizens' dwell time for online shopping on a weekend.

- Collect statistics on female netizens who dwell on online shopping for more than 2 hours on the weekend.
- The first column in the log file records names, the second column records genders, and the third column records the dwell durations in the unit of minute. Three columns are separated by comma (,).

log1.txt: logs collected on Saturday

```
LiuYang,female,20  
YuanJing,male,10  
GuoYijun,male,5  
CaiXuyu,female,50  
Liyuan,male,20  
FangBo,female,50  
LiuYang,female,20  
YuanJing,male,10  
GuoYijun,male,50  
CaiXuyu,female,50  
FangBo,female,60
```

log2.txt: logs collected on Sunday

```
LiuYang,female,20
YuanJing,male,10
CaiXuyu,female,50
FangBo,female,50
GuoYijun,male,5
CaiXuyu,female,50
Liyuan,male,20
CaiXuyu,female,50
FangBo,female,50
LiuYang,female,20
YuanJing,male,10
FangBo,female,50
GuoYijun,male,50
CaiXuyu,female,50
FangBo,female,60
```

Prerequisites

- On Manager, you have created a user and granted the HDFS, Yarn, Kafka, and Hive permissions to the user.
- You have installed and configured tools such as IntelliJ IDEA and JDK based on the development language.
- You have installed the Spark2x client and configured the client network connection.
- For Spark SQL programs, you have started Spark SQL or Beeline on the client to enter SQL statements.

Procedure

Step 1 Obtain the sample project and import it to IDEA. Import the JAR package on which the sample project depends. Use IDEA to configure and generate JAR packages.

Step 2 Prepare the data required by the sample project.

Save the original log files in the scenario description to the HDFS system.

1. Create two text files (**input_data1.txt** and **input_data2.txt**) on the local host and copy the content in the **log1.txt** and **log2.txt** files to the **input_data1.txt** and **input_data2.txt** files, respectively.
2. Create the **/tmp/input** directory in HDFS, and upload **input_data1.txt** and **input_data2.txt** to the **/tmp/input** directory:

Step 3 Upload the generated JAR package to the Spark2x running environment (Spark2x client), for example, **/opt/female**.

Step 4 Go the client directory, configure the environment variables, and log in to the system. When you use a client to connect to a specific instance in a scenario where multiple Spark2x instances are installed or Spark and Spark2x instances are installed, run the following commands to load the environment variables of the instance.

```
source bigdata_env
```

```
source Spark2x/component_env
```

```
kinit <Service user for authentication>
```

Step 5 Run the following script in the **bin** directory to submit the Spark application:

```
spark-submit --class com..bigdata.spark.examples.FemaleInfoCollection --master
yarn-client /opt/female/FemaleInfoCollection.jar <inputPath>
```

NOTE

- **FemaleInfoCollection.jar** is the JAR package generated in **Step 1**.
- **<inputPath>** is the directory created in **Step 2.2**.

Step 6 (Optional) After calling the **spark-sql** or **spark-beeline** script in the **bin** directory, directly enter SQL statements to perform operations such as query.

For example, create a table, insert a piece of data, and then query the table.

```
spark-sql> CREATE TABLE TEST(NAME STRING, AGE INT);
Time taken: 0.348 seconds
spark-sql>INSERT INTO TEST VALUES('Jack', 20);
Time taken: 1.13 seconds
spark-sql> SELECT * FROM TEST;
Jack      20
Time taken: 0.18 seconds, Fetched 1 row(s)
```

Step 7 View the running result of the Spark application.

- View the running result data in a specified file.

The storage path and format of the result data are specified by the Spark application.
- Check the running status on the web page.
 - Log in to Manager. Select **Spark2x** from the **Service** drop-down list.
 - Go to the Spark2x overview page and click an instance in the Spark web UI, for example, **JobHistory2x(host2)**.
 - The History Server UI is displayed.

The History Server UI is used to display the status of Spark applications that are complete or incomplete.


Figure 26-1 History Server UI

Version	App ID	App Name	Started	Completed	Duration	Spark User	Last Updated	Event Log
	application_...	Spark Pi	2021-06-22 16:39:06	2021-06-22 16:39:56	50 s		2021-06-22 16:39:56	Download
	application_...	Spark Pi	2021-06-22 16:39:11	2021-06-22 16:39:55	45 s		2021-06-22 16:39:55	Download
	application_...	Spark Pi	2021-06-22 16:39:10	2021-06-22 16:39:55	44 s		2021-06-22 16:39:55	Download
	application_...	Spark Pi	2021-06-22 16:39:10	2021-06-22 16:39:46	35 s		2021-06-22 16:39:46	Download
	application_...	Spark Pi	2021-06-22 16:39:06	2021-06-22 16:39:44	38 s		2021-06-22 16:39:44	Download
	application_...	Spark Pi	2021-06-22 16:39:05	2021-06-22 16:39:26	21 s		2021-06-22 16:39:26	Download
	application_...	Spark Pi	2021-06-22 16:38:13	2021-06-22 16:39:05	52 s		2021-06-22 16:39:05	Download
	application_...	Spark Pi	2021-06-22 16:38:13	2021-06-22 16:38:57	45 s		2021-06-22 16:38:58	Download
	application_...	Spark Pi	2021-06-22 16:38:12	2021-06-22 16:38:57	45 s		2021-06-22 16:38:57	Download
	application_...	Spark Pi	2021-06-22 16:38:12	2021-06-22 16:38:54	42 s		2021-06-22 16:38:54	Download
	application_...	Spark Pi	2021-06-22 16:38:09	2021-06-22 16:38:47	38 s		2021-06-22 16:38:47	Download
	application_...	Spark Pi	2021-06-22 16:38:05	2021-06-22 16:38:46	41 s		2021-06-22 16:38:46	Download
	application_...	Spark Pi	2021-06-22 16:38:06	2021-06-22 16:38:27	21 s		2021-06-22 16:38:27	Download
	application_...	Spark Pi	2021-06-22 16:38:55	2021-06-22 16:39:06	1.2 min		2021-06-22 16:39:06	Download

- Select an application ID and click this page to go to the Spark UI of the application.

Spark UI: used to display the status of running applications.

Figure 26-2 Spark UI



- View Spark logs to learn application runtime conditions.
View [Spark2x Logs](#) to learn application running status, and adjust applications based on log information.

----End

26.2.2 Configuring Parameters Rapidly

Overview

This section describes how to quickly configure common parameters and lists parameters that are not recommended to be modified when Spark2x is used.

Common parameters to be configured

Some parameters have been adapted during cluster installation. However, the following parameters need to be adjusted based on application scenarios. Unless otherwise specified, the following parameters are configured in the **spark-defaults.conf** file on the Spark2x client.

Table 26-1 Common parameters to be configured

Configuration Item	Description	Default Value
spark.sql.parquet.compression.codec	Used to set the compression format of a non-partitioned Parquet table. Set the queue in the spark-defaults.conf configuration file on the JDBCServer server.	snappy
spark.dynamicAllocation.enabled	Indicates whether to use dynamic resource scheduling, which is used to adjust the number of executors registered with the application according to scale. Currently, this parameter is valid only in Yarn mode. The default value for JDBCServer is true , and that for the client is false .	false

Configuration Item	Description	Default Value
spark.executor.memory	Indicates the memory size used by each executor process. Its character sting is in the same format as the JVM memory (example: 512 MB or 2 GB).	4G
spark.sql.autoBroadcastJoinThreshold	Indicates the maximum value for the broadcast configuration when two tables are joined. <ul style="list-style-type: none"> When the size of a field in a table involved in an SQL statement is less than the value of this parameter, the system broadcasts the SQL statement. If the value is set to -1, broadcast is not performed. 	10485760
spark.yarn.queue	Specifies the Yarn queue where JDBCServer resides. Set the queue in the spark-defaults.conf configuration file on the JDBCServer server.	default
spark.driver.memory	In a large cluster, you are advised to configure the memory used by the 32 GB to 64 GB driver process, that is, the SparkContext initialization process (for example, 512 MB and 2 GB).	4G
spark.yarn.security.credential.hbase.enabled	Indicates whether to enable the function of obtaining HBase tokens. If the Spark on HBase function is required and a security cluster is configured, set this parameter to true . Otherwise, set this parameter to false .	false
spark.serializer	Used to serialize the objects that are sent over the network or need to be cached. The default value of Java serialization applies to any Serializable Java object, but the running speed is slow. Therefore, you are advised to use org.apache.spark.serializer.KryoSerializer and configure Kryo serialization. It can be any subclass of org.apache.spark.serializer.Serializer .	org.apache.spark.serializer.JavaSerializer

Configuration Item	Description	Default Value
spark.executor.cores	Indicates the number of kernels used by each executor. Set this parameter in standalone mode and Mesos coarse-grained mode. When there are sufficient kernels, the application is allowed to execute multiple executable programs on the same worker. Otherwise, each application can run only one executable program on each worker.	1
spark.shuffle.service.enabled	Indicates a long-term auxiliary service in NodeManager for improving shuffle computing performance.	false
spark.sql.adaptive.enabled	Indicates whether to enable the adaptive execution framework.	false
spark.executor.memoryOverhead	Indicates the heap memory to be allocated to each executor, in MB. This is the memory that occupies the overhead of the VM, similar to the internal string and other built-in overhead. The value increases with the executor size (usually 6% to 10%).	1 GB
spark.streaming.kafka.direct.lifo	Indicates whether to enable the LIFO function of Kafka.	false

Parameters Not Recommended to Be Modified

The following parameters have been adapted during cluster installation. You are not advised to modify them.

Table 26-2 Parameters not recommended to be modified

Configuration Item	Description	Default Value or Configuration Example
spark.password.factory	Selects the password parsing mode.	org.apache.spark.om.util.FIPasswordFactory
spark.ssl.ui.protocol	Sets the SSL protocol of the UI.	TLSv1.2

Configuration Item	Description	Default Value or Configuration Example
spark.yarn.archive	Archives Spark JAR files, which are distributed to Yarn cache. If this parameter is set, the value will replace <code><code>spark.yarn.jars </code></code> and be archived in the containers of all applications. The archive should contain the JAR files in its root directory. Archives can also be hosted on HDFS to speed up file distribution.	hdfs://hacluster/user/spark2x/jars/8.1.0.1/spark-archive-2x.zip NOTE The version 8.1.0.1 is used as an example. Replace it with the actual version number.
spark.yarn.am.extraJavaOptions	Indicates a string of extra JVM options to pass to the YARN ApplicationMaster in client mode. Use spark.driver.extraJavaOptions in cluster mode.	-Dlog4j.configuration=./__spark_conf__/_hadoop_conf__/log4j-executor.properties -Djava.security.auth.login.config=./__spark_conf__/_hadoop_conf__/jaas-zk.conf - Dzookeeper.server.principal=zookeeper/hadoop.<system domain name> - Djava.security.krb5.conf=./__spark_conf__/_hadoop_conf__/kdc.conf - Djdk.tls.ephemeralDHKeySize=2048
spark.shuffle.servicev2.port	Indicates the port for the shuffle service to monitor requests for obtaining data.	27338
spark.ssl.historyServer.enabled	Sets whether the history server uses SSL.	true
spark.files.override	When the target file exists and its content does not match that of the source file, whether to overwrite the file added through SparkContext.addFile() .	false

Configuration Item	Description	Default Value or Configuration Example
spark.yarn.cluster.driver.extraClassPath	Indicates the extraClassPath of the driver in Yarn-cluster mode. Set the parameter to the path and parameters of the server.	<code>\${BIGDATA_HOME}/common/runtime/security</code>
spark.driver.extraClassPath	Indicates the extra class path entries attached to the class path of the driver.	<code>\${BIGDATA_HOME}/common/runtime/security</code>
spark.yarn.dist.innerfiles	Sets the files that need to be uploaded to HDFS from Spark in Yarn mode.	<code>/Spark_path/spark/conf/s3p.file,/Spark_path/spark/conf/locals3.jceks</code> <i>Spark_path</i> is the installation path of the Spark client.
spark.sql.bigdata.register.dialect	Registers the SQL parser.	<code>org.apache.spark.sql.hbase.HBaseSQLParser</code>
spark.shuffle.manager	Indicates the data processing mode. There are two implementation modes: sort and hash. The sort shuffle has a higher memory utilization. It is the default option in Spark 1.2 and later versions.	<code>SORT</code>
spark.deploy.zookeeper.url	Indicates the address of ZooKeeper. Multiple addresses are separated by commas (,).	For example: <code>host1:2181,host2:2181,host3:2181</code>
spark.broadcast.factory	Indicates the broadcast mode.	<code>org.apache.spark.broadcast.TorrentBroadcastFactory</code>
spark.sql.session.state.builder	Session state constructor.	<code>org.apache.spark.sql.hive.FIHiveACLSessionStateBuilder</code>

Configuration Item	Description	Default Value or Configuration Example
spark.executor.extraLibraryPath	Sets the special library path used when the executor JVM is started.	<code>\${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-Hadoop-*/hadoop/lib/native</code>
spark.ui.customErrorPage	Indicates whether to display the custom error information page when an error occurs on the page.	true
spark.httpdProxy.enable	Indicates whether to use the httpd proxy.	true
spark.ssl.ui.enabledAlgorithms	Sets the SSL algorithm of UI.	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,TLS_DHE_DSS_WITH_AES_256_GCM_SHA384,TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,TLS_DHE_DSS_WITH_AES_128_GCM_SHA256
spark.ui.logout.enabled	Sets the logout button for the web UI of the Spark component.	true
spark.security.hideInfo.enabled	Indicates whether to hide sensitive information on the UI.	true
spark.yarn.cluster.driver.extraLibraryPath	Indicates the extraLibraryPath of the driver in Yarn-cluster mode. Set this parameter to the path and parameters of the server.	<code>\${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-Hadoop-*/hadoop/lib/native</code>
spark.driver.extraLibraryPath	Sets a special library path for starting the driver JVM.	<code>\${DATA_NODE_INSTALL_HOME}/hadoop/lib/native</code>

Configuration Item	Description	Default Value or Configuration Example
spark.ui.killedEnabled	Allows stages and jobs to be stopped on the web UI.	true
spark.yarn.access.hadoopFileSystems	Spark can access multiple NameService instances. If there are multiple NameService instances, set this parameter to all the NameService instances and separate them with commas (,).	hdfs://hacluster,hdfs://hacluster
spark.yarn.cluster.driver.extraJavaOptions	Indicates extra JVM option passed to the executor, for example, GC setting and logging. Do not set Spark attributes or heap size using this option. Instead, set Spark attributes using the SparkConf object or the spark-defaults.conf file specified when the spark-submit script is called. Set heap size using spark.executor.memory .	-Xloggc:<LOG_DIR>/gc.log - XX:+PrintGCDetails -XX:-OmitStackTraces - FastThrow -XX:+PrintGCDateStamps - XX:+PrintGCDateStamps - XX:+UseGCLogFileRotation - XX:NumberOfGCLogFiles=20 - XX:GCLogFileSize=10M - Dlog4j.configuration=../_spark_conf_/__hadoop_conf_/log4j-executor.properties -Djava.security.auth.login.config=../_spark_conf_/__hadoop_conf_/jaas-zk.conf - Dzookeeper.server.principal=zookeeper/hadoop.<system domain name> - Djava.security.krb5.conf=../_spark_conf_/__hadoop_conf_/kdc.conf - Djetty.version=x.y.z - Dorg.xerial.snappy.tmpdir=\${BIGDATA_HOME}/tmp/spark2x_app - Dcarbon.properties.filepath=../_spark_conf_/__hadoop_conf_/carbon.properties - Djdk.tls.ephemeralDHKeySize=2048

Configuration Item	Description	Default Value or Configuration Example
spark.driver.extraJavaOptions	Indicates a series of extra JVM options passed to the driver,	<pre>-Xloggc:\${SPARK_LOG_DIR}/indexserver-omm-%p-gc.log -XX:+PrintGCDetails -XX:-OmitStackTracelnFastThrow -XX:+PrintGCTimeStamps -XX:+PrintGCDateStamps -XX:MaxDirectMemorySize=512M -XX:MaxMetaspaceSize=512M -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=20 -XX:GCLogFileSize=10M -XX:OnOutOfMemoryError='kill -9 %p' -Djetty.version=x.y.z -Dorg.xerial.snappy.tmpdir=\${BIGDATA_HOME}/tmp/spark2x/JDBCServer/snappy_tmp -Djava.io.tmpdir=\${BIGDATA_HOME}/tmp/spark2x/JDBCServer/io_tmp -Dcarbon.properties.filepath=\${SPARK_CONF_DIR}/carbon.properties -Djdk.tls.ephemeralDHKeySize=2048 -Dspark.ssl.keyStore=\${SPARK_CONF_DIR}/child.keystore #{java_stack_prefer}</pre>
spark.eventLog.overwrite	Indicates whether to overwrite any existing file.	false
spark.eventLog.dir	Indicates the directory for logging Spark events if spark.eventLog.enabled is set to true . In this directory, Spark creates a subdirectory for each application and logs events of the application in the subdirectory. You can also set a unified address similar to the HDFS directory so that the History Server can read historical files.	hdfs://hacluster/spark2xJobHistory2x

Configuration Item	Description	Default Value or Configuration Example
spark.random.port.min	Sets the minimum random port.	22600
spark.authenticate	Indicates whether Spark authenticates its internal connections. If the application is not running on Yarn, see spark.authenticate.secret .	true
spark.random.port.max	Sets the maximum random port.	22899
spark.eventLog.enabled	Indicates whether to log Spark events, which are used to reconstruct the web UI after the application execution is complete.	true

Configuration Item	Description	Default Value or Configuration Example
spark.executor.extraJavaOptions	Indicates extra JVM option passed to the executor, for example, GC setting and logging. Do not set Spark attributes or heap size using this option.	<pre>-Xloggc:<LOG_DIR>/gc.log - XX:+PrintGCDetails -XX:-OmitStackTraceln- FastThrow -XX:+PrintGCTimeStamps - XX:+PrintGCDateStamps - XX:+UseGCLogFileRotation - XX:NumberOfGCLogFiles=20 - XX:GCLogFileSize=10M - Dlog4j.configuration=./log4j- executor.properties - Djava.security.auth.login.config=./jaas- zk.conf - Dzookeeper.server.principal=zookeeper/ hadoop.<system domain name> - Djava.security.krb5.conf=./kdc.conf - Dcarbon.properties.filepath=./ carbon.properties -Xloggc:<LOG_DIR>/gc.log - XX:+PrintGCDetails -XX:-OmitStackTraceln- FastThrow -XX:+PrintGCTimeStamps - XX:+PrintGCDateStamps - XX:+UseGCLogFileRotation - XX:NumberOfGCLogFiles=20 - XX:GCLogFileSize=10M - Dlog4j.configuration=./_spark_conf_/ _hadoop_conf_/log4j-executor.properties -Djava.security.auth.login.config=./ _spark_conf_/_hadoop_conf_/jaas- zk.conf - Dzookeeper.server.principal=zookeeper/ hadoop.<system domain name> - Djava.security.krb5.conf=./_spark_conf_/ _hadoop_conf_/kdc.conf - Dcarbon.properties.filepath=./ _spark_conf_/_hadoop_conf_/ carbon.properties - Djdk.tls.ephemeralDHKeySize=2048</pre>
spark.sql.authorization.enabled	Indicates whether to enable authentication for the Hive client.	true

26.2.3 Common Parameters

Overview

This section describes common configuration items used in Spark. Subsections are divided by feature so that you can quickly find required configuration items. If you

use MRS clusters, most parameters described in this section have been adapted and you do not need to configure them again. For details about the parameters that need to be configured based on the site requirements, see [Configuring Parameters Rapidly](#).

Configuring the Number of Stage Retries

When `FetchFailedException` occurs in a Spark task, a stage retry is triggered. To prevent infinite stage retries, the number of stage retries is limited. The number of retry times can be adjusted based on the site requirements.

Configure the following parameters in the `spark-defaults.conf` file on the Spark client.

Table 26-3 Parameter description

Parameter	Description	Default Value
<code>spark.stage.maxConsecutiveAttempts</code>	Indicates the maximum number of stage retries.	4

Configuring Whether to Use Cartesian Product

To enable the Cartesian product function, configure the following parameter in the `spark-defaults.conf` configuration file of Spark.

Table 26-4 Cartesian product parameters

Parameter	Description	Default Value
<code>spark.sql.crossJoin.enabled</code>	Indicates whether to allow implicit Cartesian product execution. <ul style="list-style-type: none"> true: Implicit Cartesian product execution is allowed. false: Implicit Cartesian product execution is not allowed. In this case, only CROSS JOIN can be explicitly included in the query. 	true

NOTE

- For JDBC applications, configure this parameter in the `spark-defaults.conf` configuration file of the server.
- For tasks submitted by the Spark client, configure this parameter in the `spark-defaults.conf` configuration file of the client.

Configuring Security Authentication for Long-Time Spark Tasks

In security mode, if the **kinit** command is used for security authentication when the Spark CLI (such as `spark-shell`, `spark-sql`, or `spark-submit`) is used, the task fails due to authentication expiration when the task is running for a long time.

Set the following parameters in the **spark-defaults.conf** configuration file on the client. After the configuration is complete, run the Spark CLI again.

 **NOTE**

If this parameter is set to **true**, ensure that the values of **keytab** and **principal** in **spark-defaults.conf** and **hive-site.xml** are the same.

Table 26-5 Parameter description

Parameter	Description	Default Value
spark.kerberos.principal	Indicates the principal user who has the Spark operation permission. Contact the MRS cluster administrator to obtain the principal user.	-
spark.kerberos.keytab	Indicates the name and path of the keytab file used to configure Spark operation permissions. Contact the MRS cluster administrator to obtain the Keytab file.	-
spark.security.bigdata.loginOnce	<p>Indicates whether the principal user logs in to the system only once. true: single login; false: multiple logins.</p> <p>The difference between a single login and multiple logins is as follows: The Spark community uses the Kerberos user to log in to the system for multiple times. However, the TGT or token may expire, causing the application to fail to run for a long time. The Kerberos login mode of DataSight is modified to allow users to log in only once, which effectively resolves the expiration problem. The restrictions are as follows: The principal and keytab configuration items of Hive must be the same as those of Spark.</p> <p>NOTE If this parameter is set to true, ensure that the values of keytab and principal in spark-defaults.conf and hive-site.xml are the same.</p>	true

Python Spark

Python Spark is the third programming language of Spark except Scala and Java. Different from Java and Scala that run on the JVM platform, Python Spark has its own Python process as well as the JVM process. The following configuration items

apply only to Python Spark scenarios. However, other configuration items can also take effect in Python Spark scenarios.

Table 26-6 Parameter description

Parameter	Description	Default Value
spark.python.profile	Indicates whether to enable profiling on the Python worker. Use sc.show_profiles() to display the analysis results or display the analysis results before the Driver exits. You can use sc.dump_profiles(path) to dump the results to a disk. If some analysis results have been manually displayed, they will not be automatically displayed before the driver exits. By default, pyspark.profiler.BasicProfiler is used. You can transfer the specified profiler during SparkContext initialization to overwrite the default profiler.	false
spark.python.worker.memory	Indicates the memory size that can be used by each Python worker process during aggregation. The value format is the same as that of the specified JVM memory, for example, 512 MB and 2 GB. If the memory used by a process during aggregation exceeds the value of this parameter, data will be written to disks.	512m
spark.python.worker.reuse	Indicates whether to reuse Python workers. If the reuse function is enabled, a fixed number of Python workers will be reused by the next batch of submitted tasks instead of forking a Python process for each task. This function is useful in large-scale broadcasting because the data does not need to be transferred from the JVM to the Python workers again for the next batch of submitted tasks.	true

Dynamic Allocation

Dynamic resource scheduling is a unique feature of the On Yarn mode. This function can be used only after Yarn External Shuffle is enabled. When Spark is used as a resident service, dynamic resource scheduling greatly improves resource utilization. For example, the JDBCServer process does not accept JDBC requests in most of the time. Therefore, releasing resources in this period greatly reduces the waste of cluster resources.

Table 26-7 Parameter description

Parameter	Description	Default Value
spark.dynamicAllocation.enabled	Indicates whether to use dynamic resource scheduling, which is used to adjust the number of executors registered with the application according to scale. Currently, this parameter is valid only in Yarn mode. To enable dynamic resource scheduling, set spark.shuffle.service.enabled to true . Related parameters are as follows: spark.dynamicAllocation.minExecutors , spark.dynamicAllocation.maxExecutors , and spark.dynamicAllocation.initialExecutors .	<ul style="list-style-type: none"> JDBCServer2x: true SparkResource2x: false
spark.dynamicAllocation.minExecutors	Indicates the minimum number of executors.	0
spark.dynamicAllocation.initialExecutors	Indicates the number of initial executors.	spark.dynamicAllocation.minExecutors
spark.dynamicAllocation.maxExecutors	Indicates the maximum number of executors.	2048
spark.dynamicAllocation.schedulerBacklogTimeout	Indicates the first timeout period for scheduling. The unit is second.	1s
spark.dynamicAllocation.sustainedSchedulerBacklogTimeout	Indicates the second and later timeout interval for scheduling.	1s
spark.dynamicAllocation.executorIdleTimeout	Indicates the idle timeout interval for common executors. The unit is second.	60

Parameter	Description	Default Value
spark.dynamicAllocation.cachedExecutorIdleTimeout	Indicates the idle timeout interval for executors with cached blocks.	<ul style="list-style-type: none"> JDBCServer2x: 2147483647s IndexServer2x: 2147483647s SparkResource2x: 120

Spark Streaming

Spark Streaming is a streaming data processing function provided by the Spark batch processing platform. It processes data input from external systems in **mini-batch** mode.

Configure the following parameters in the **spark-defaults.conf** file on the Spark client.

Table 26-8 Parameter description

Parameter	Description	Default Value
spark.streaming.receiver.writeAheadLog.enable	Indicates whether to enable the write-ahead log (WAL) function. After this function is enabled, all input data received by the receiver is saved in the WAL. WAL ensures that data can be restored if the driver program becomes faulty.	false
spark.streaming.unpersist	Determines whether to automatically remove RDDs generated and saved by Spark Streaming from the Spark memory. If this function is enabled, original data received by Spark Streaming is also automatically cleared. If this function is disabled, original data and RDDs cannot be automatically cleared. External applications can access the data in Streaming. This, however, occupies more Spark memory resources.	true

Spark Streaming Kafka

The receiver is an important component of Spark Streaming. It receives external data, encapsulates the data into blocks, and provides the blocks for Streaming to

consume. The most common data source is Kafka. Spark Streaming integrates Kafka to ensure reliability and can directly use Kafka as the RDD input.

Table 26-9 Parameter description

Parameter	Description	Default Value
spark.streaming.kafka.maxRatePerPartition	Indicates the maximum rate (number of records per second) for reading data from each Kafka partition if the Kafka direct stream API is used.	-
spark.streaming.blockInterval	Indicates the interval (ms) for accumulating data received by a Spark Streaming receiver into a data block before the data is stored in Spark. A minimum value of 50 ms is recommended.	200ms
spark.streaming.receiver.maxRate	Indicates the maximum rate (number of records per second) for each receiver to receive data. The value 0 or a negative value indicates no limit to the rate.	-
spark.streaming.receiver.writeAheadLog.enabled	Indicates whether to use ReliableKafkaReceiver. This receiver ensures the integrity of streaming data.	false

Netty/NIO and Hash/Sort Configuration

Shuffle is critical for big data processing, and the network is critical for the entire shuffle process. Currently, Spark supports two shuffle modes: hash and sort. There are two network modes: Netty and NIO.

Table 26-10 Parameter description

Parameter	Description	Default Value
spark.shuffle.manager	Indicates the data processing mode. There are two implementation modes: sort and hash. The sort shuffle has a higher memory utilization. It is the default option in Spark 1.2 and later versions.	SORT

Parameter	Description	Default Value
spark.shuffle.consolidateFiles	(Only in hash mode) To merge intermediate files created during shuffle, set this parameter to true . Decreasing the number of files to be created can improve the processing performance of the file system and reduce risks. If the ext4 or xfs file system is used, you are advised to set this parameter to true . Due to file system restrictions, this setting on ext3 may reduce the processing performance of a server with more than eight cores.	false
spark.shuffle.sort.byPassMergeThreshold	This parameter is valid only when spark.shuffle.manager is set to sort . When Map aggregation is not performed and the number of partitions for Reduce tasks is less than or equal to the value of this parameter, do not merge and sort data to prevent performance deterioration caused by unnecessary sorting.	200
spark.shuffle.io.maxRetries	(Only in Netty mode) If this parameter is set to a non-zero value, fetch failures caused by I/O-related exceptions will be automatically retried. This retry logic helps the large shuffle keep stable when long GC pauses or intermittent network disconnections occur.	12
spark.shuffle.io.numConnectionsPerPeer	(Only in Netty mode) Connections between hosts are reused to reduce the number of connections between large clusters. For a cluster with many disks but a few hosts, this function may make concurrent requests unable to occupy all disks. Therefore, you can increase the value of this parameter.	1
spark.shuffle.io.preferDirectBufs	(Only in Netty mode) The off-heap buffer is used to reduce GC during shuffle and cache block transfer. In an environment where off-heap memory is strictly limited, you can disable it to force all applications from Netty to use heap memory.	true
spark.shuffle.io.retryWait	(Only in Netty mode) Specifies the duration for waiting for fetch retry, in seconds. The maximum delay caused by retry is maxRetries x retryWait . The default value is 15 seconds.	5

Common Shuffle Configuration

Table 26-11 Parameter description

Parameter	Description	Default Value
spark.shuffle.spill	If this parameter is set to true , data is overflowed to the disk to limit the memory usage during a Reduce task.	true
spark.shuffle.spill.compress	Indicates whether to compress the data overflowed during shuffle. The algorithm specified by spark.io.compression.codec is used for data compression.	true
spark.shuffle.file.buffer	Specifies the size of the memory buffer for storing output streams of each shuffle file, in KB. These buffers can reduce the number of disk seek and system calls during the creation of intermediate shuffle file streams. You can also set this parameter by setting spark.shuffle.file.buffer.kb .	32KB
spark.shuffle.compress	Indicates whether to compress the output files of a Map task. You are advised to compress the broadcast variables. using spark.io.compression.codec .	true
spark.reducer.maxSizeInFlight	Specifies the maximum output size of the Map task that fetches data from each Reduce task, in MB. Each output requires a buffer, which is the fixed memory overhead of each Reduce task. Therefore, keep the value small unless there is a large amount of memory. You can also set this parameter by setting spark.reducer.maxMbInFlight .	48MB

Driver Configuration

Spark driver can be considered as the client of Spark applications. All code parsing is completed in this process. Therefore, the parameters of this process are especially important. The following describes how to configure parameters for Spark driver.

- **JavaOptions:** parameter following **-D** in the Java command, which can be obtained by **System.getProperty**
- **ClassPath:** path for loading the Java classes and Native library
- **Java Memory and Cores:** memory and CPU usage of the Java process
- **Spark Configuration:** Spark internal parameter, which is irrelevant to the Java process

Table 26-12 Parameter description

Parameter	Description	Default Value
spark.driver.extraJavaOptions	<p>Indicates a series of extra JVM options passed to the driver, for example, GC setting and logging.</p> <p>Note: In client mode, this configuration cannot be set directly in the application using SparkConf because the driver JVM has been started. You can use --driver-java-options or the default property file to set the parameter.</p>	For details, see Configuring Parameters Rapidly .
spark.driver.extraClassPath	<p>Indicates the extra class path entries attached to the class path of the driver.</p> <p>Note: In client mode, this configuration cannot be set directly in the application using SparkConf because the driver JVM has been started. You can use --driver-java-options or the default property file to set the parameter.</p>	For details, see Configuring Parameters Rapidly .
spark.driver.userClassPathFirst	<p>(Trial) Indicates whether to allow JAR files added by users to take precedence over Spark JAR files when classes are loaded in the driver. This feature can be used to mitigate conflicts between Spark dependencies and user dependencies. This feature is in the trial phase and is used only in cluster mode.</p>	false
spark.driver.extraLibraryPath	<p>Sets a special library path for starting the driver JVM.</p> <p>Note: In client mode, this configuration cannot be set directly in the application using SparkConf because the driver JVM has been started. You can use --driver-java-options or the default property file to set the parameter.</p>	<ul style="list-style-type: none"> JDBCServer2x: \$ {SPARK_INSTALLED_HOME}/spark/native SparkResource2x: \$ {DATA_NODE_INSTANCE_HOME}/hadoop/lib/native

Parameter	Description	Default Value
spark.driver.cores	Specifies the number of cores used by the driver process. This parameter is available only in cluster mode.	1
spark.driver.memory	Indicates the memory used by the driver process, that is, the memory used by the SparkContext initialization process (for example, 512 MB and 2 GB). Note: In client mode, this configuration cannot be set directly in the application using SparkConf because the driver JVM has been started. You can use --driver-java-options or the default property file to set the parameter.	4G
spark.driver.maxResultSize	Indicates the total size of serialization results of all partitions for each Spark action operation (for example, collect). The value must be at least 1 MB. If this parameter is set to 0 , the size is not limited. If the total amount exceeds this limit, the task will be aborted. If the value is too large, the memory of the driver may be insufficient (depending on the object memory overhead of spark.driver.memory and JVM). Set a proper limit to ensure sufficient memory for the driver.	1G
spark.driver.host	Specifies the host name or IP address for the driver to listen on, which is used for the driver to communicate with the executor.	(local hostname)
spark.driver.port	Specifies the port for the driver to listen on, which is used for the driver to communicate with the executor.	(random)

ExecutorLauncher Configuration

ExecutorLauncher exists only in Yarn-client mode. In Yarn-client mode, ExecutorLauncher and the driver are not in the same process. Therefore, you need to configure parameters for ExecutorLauncher.

Table 26-13 Parameter description

Parameter	Description	Default Value
spark.yarn.am.extraJavaOptions	Indicates a string of extra JVM options to pass to the YARN ApplicationMaster in client mode. Use spark.driver.extraJavaOptions in cluster mode.	For details, see Configuring Parameters Rapidly .
spark.yarn.am.memory	Indicates the amount of memory to use for the YARN ApplicationMaster in client mode, in the same format as JVM memory strings (for example, 512 MB or 2 GB). In cluster mode, use spark.driver.memory instead.	1G
spark.yarn.am.memoryOverhead	This parameter is the same as spark.yarn.driver.memoryOverhead . However, this parameter applies only to ApplicationMaster in client mode.	-
spark.yarn.am.cores	Indicates the number of cores to use for the YARN ApplicationMaster in client mode. Use spark.driver.cores in cluster mode.	1

Executor Configuration

An executor is a Java process. However, unlike the driver and ApplicationMaster, an executor can have multiple processes. Spark supports only same configurations. That is, the process parameters of all executors must be the same.

Table 26-14 Parameter description

Parameter	Description	Default Value
spark.executor.extraJavaOptions	Indicates extra JVM option passed to the executor, for example, GC setting and logging. Do not set Spark attributes or heap size using this option. Instead, set Spark attributes using the SparkConf object or the spark-defaults.conf file specified when the spark-submit script is called. Set heap size using spark.executor.memory .	For details, see Configuring Parameters Rapidly .

Parameter	Description	Default Value
spark.executor.extraClassPath	Indicates the extra classpath attached to the executor classpath. This parameter ensures compatibility with historical versions of Spark. Generally, you do not need to set this parameter.	-
spark.executor.extraLibraryPath	Sets the special library path used when the executor JVM is started.	For details, see Configuring Parameters Rapidly .
spark.executor.userClassPathFirst	(Trial) Same function as spark.driver.userClassPathFirst . However, this parameter applies to executor instances.	false
spark.executor.memory	Indicates the memory size used by each executor process. Its character string is in the same format as the JVM memory (example: 512 MB or 2 GB).	4G
spark.executorEnv. [EnvironmentVariableName]	Adds the environment variable specified by EnvironmentVariableName to the executor process. You can specify multiple environment variables.	-
spark.executor.logs.rolling.maxRetainedFiles	Sets the number of latest log files to be retained by the system during rolling. The old log files are deleted. This function is disabled by default.	-
spark.executor.logs.rolling.size.maxBytes	Sets the maximum size of the executor log file for rolling. This function is disabled by default. The value is in bytes. To automatically clear old logs, see spark.executor.logs.rolling.maxRetainedFiles .	-
spark.executor.logs.rolling.strategy	Sets the executor log rolling policy. Rolling is disabled by default. The value can be time (time-based rolling) or size (size-based rolling). If this parameter is set to time , the value of the spark.executor.logs.rolling.time.interval attribute is used as the log rolling interval. If this parameter is set to size , spark.executor.logs.rolling.size.maxBytes is used to set the maximum size of the file for rolling.	-

Parameter	Description	Default Value
spark.executor.logs.rolling.time.interval	Sets the time interval for executor log rolling. This function is disabled by default. The value can be daily , hourly , minutely , or any number of seconds. To automatically clear old logs, see spark.executor.logs.rolling.maxRetainedFiles .	daily

WebUI

The Web UI displays the running process and status of the Spark application.

Table 26-15 Parameter description

Parameter	Description	Default Value
spark.ui.killEnabled	Allows stages and jobs to be stopped on the web UI. NOTE For security purposes, the default value of this parameter is set to false to prevent misoperations. To enable this function, set this parameter to true in the spark-defaults.conf configuration file. Exercise caution when performing this operation.	true
spark.ui.port	Specifies the port for your application's dashboard, which displays memory and workload data.	<ul style="list-style-type: none"> • JDBC Server2x: 4040 • Spark Resource2x: 0 • Index Server2x: 22901
spark.ui.retainedJobs	Specifies the number of jobs recorded by the Spark UI and status API before GC.	1000
spark.ui.retainedStages	Specifies the number of stages recorded by the Spark UI and status API before GC.	1000

HistoryServer

A History Server reads the **EventLog** file in the file system and displays the running status of the Spark application.

Table 26-16 Parameter description

Parameter	Description	Default Value
spark.history.fs.logDirectory	Specifies the log directory of a History Server.	-
spark.history.ui.port	Specifies the port for JobHistory listening to connection.	18080
spark.history.fs.updateInterval	Specifies the update interval of the information displayed on a History Server, in seconds. Each update checks for changes made to the event logs in the persistent store.	10s
spark.history.fs.updateInterval.seconds	Specifies the interval for checking the update of each event log. This parameter has the same function as spark.history.fs.updateInterval . spark.history.fs.updateInterval is recommended.	10s
spark.history.updateInterval	This parameter has the same function as spark.history.fs.updateInterval.seconds and spark.history.fs.updateInterval . spark.history.fs.updateInterval is recommended.	10s

History Server UI Timeout and Maximum Number of Access Times

Table 26-17 Parameter description

Parameter	Description	Default Value
spark.session.maxAge	Specifies the session timeout interval, in seconds. This parameter applies only to the security mode. This parameter cannot be set in normal mode.	600
spark.connection.maxRequest	Specifies the maximum number of concurrent client access requests to JobHistory.	5000

EventLog

During the running of Spark applications, the running status is written into the file system in JSON format in real time for the History Server service to read and reproduce the application running status.

Table 26-18 Parameter description

Parameter	Description	Default Value
spark.eventLog.enabled	Indicates whether to log Spark events, which are used to reconstruct the web UI after the application execution is complete.	true
spark.eventLog.dir	Indicates the directory for logging Spark events if spark.eventLog.enabled is set to true . In this directory, Spark creates a subdirectory for each application and logs events of the application in the subdirectory. You can also set a unified address similar to the HDFS directory so that the History Server can read historical files.	hdfs://hacluster/spark2x/jobHistory2x
spark.eventLog.compress	Indicates whether to compress logged events when spark.eventLog.enabled is set to true .	false

Periodic Clearing of Event Logs

Event logs on JobHistory increases with submitted tasks. Too many event log files exist as the number of submitted tasks increases. Spark provides the function for periodically clearing event logs. You can enable this function and set the clearing interval using related parameters.

Table 26-19 Parameter description

Parameter	Description	Default Value
spark.history.fs.cleaner.enabled	Indicates whether to enable the clearing function.	true
spark.history.fs.cleaner.interval	Indicates the check interval of the clearing function.	1d
spark.history.fs.cleaner.maxAge	Indicates the maximum duration for storing logs.	4d

Kryo

Kryo is a highly efficient Java serialization framework, which is integrated into Spark by default. Almost all Spark performance tuning requires the process of converting the default serializer of Spark into a Kryo serializer. Kryo serialization

supports only serialization at the Spark data layer. To configure Kryo serialization, set **spark.serializer** to **org.apache.spark.serializer.KryoSerializer** and configure the following parameters to optimize Kryo serialization performance:

Table 26-20 Parameter description

Parameter	Description	Default Value
spark.kryo.classesToRegister	Specifies the name of the class that needs to be registered with Kryo when Kryo serialization is used. Multiple classes are separated by commas (,).	-
spark.kryo.referenceTracking	Indicates whether to trace the references to the same object when Kryo is used to serialize data. This function is applicable to the scenario where the object graph has circular references or the same object has multiple copies. Otherwise, you can disable this function to improve performance.	true
spark.kryo.registrationRequired	Indicates whether Kryo is used to register an object. When this parameter is set to true , an exception is thrown if an object that is not registered with Kryo is serialized. When it is set to false (default value), Kryo writes unregistered class names to the serialized object. This operation causes a large amount of performance overhead. Therefore, you need to enable this option before deleting a class from the registration queue.	false
spark.kryo.registrator	If Kryo serialization is used, use Kryo to register the class with the custom class. Use this property if you need to register a class in a custom way, such as specifying a custom field serializer. Otherwise, use spark.kryo.classesToRegister , which is simpler. Set this parameter to a class that extends KryoRegistrator.	-
spark.kryoserializer.buffer.max	Specifies the maximum size of the Kryo serialization buffer, in MB. The value must be greater than the object that attempts to be serialized. If the error "buffer limit exceeded" occurs in Kryo, increase the value of this parameter. You can also set this parameter by setting spark.kryoserializer.buffer.max .	64MB

Parameter	Description	Default Value
spark.kryoserializer.buffer	Specifies the initial size of the Kryo serialization buffer, in MB. Each core of each worker has a buffer. If necessary, the buffer size will be increased to the value of spark.kryoserializer.buffer.max . You can also set this parameter by setting spark.kryoserializer.buffer .	64KB

Broadcast

Broadcast is used to transmit data blocks between Spark processes. In Spark, broadcast can be used for JAR packages, files, closures, and returned results. Broadcast supports two modes: Torrent and HTTP. The Torrent mode divides data into small fragments and distributes them to clusters. Data can be obtained remotely if necessary. The HTTP mode saves files to the local disk and transfers the entire files to the remote end through HTTP if necessary. The former is more stable than the latter. Therefore, Torrent is the default broadcast mode.

Table 26-21 Parameter description

Parameter	Description	Default Value
spark.broadcast.factory	Indicates the broadcast mode.	org.apache.spark.broadcast.TorrentBroadcastFactory
spark.broadcast.blockSize	Indicates the block size of TorrentBroadcastFactory . If the value is too large, the concurrency during broadcast is reduced (the speed is slow). If the value is too small, BlockManager performance may be affected.	4096
spark.broadcast.compress	Indicates whether to compress broadcast variables before sending them. You are advised to compress the broadcast variables.	true

Storage

Spark features in-memory computing. Spark Storage is used to manage memory resources. Storage stores data blocks generated during RDD caching. The heap memory in the JVM acts as a whole. Therefore, **Storage Memory Size** is an important concept during Spark Storage management.

Table 26-22 Parameter description

Parameter	Description	Default Value
spark.storage.memoryMapThreshold	Specifies the block size. If the size of a block exceeds the value of this parameter, Spark performs memory mapping for the disk file. This prevents Spark from mapping too small blocks during memory mapping. Generally, memory mapping for blocks whose page size is close to or less than that of the operating system has high overhead.	2m

PORT

Table 26-23 Parameter description

Parameter	Description	Default Value
spark.ui.port	Specifies the port for your application's dashboard, which displays memory and workload data.	<ul style="list-style-type: none"> JDBC Server2x: 4040 SparkResource2x: 0
spark.blockManager.port	Specifies all ports listened by BlockManager. These ports are on both the driver and executor.	Range of Random Ports
spark.driver.port	Specifies the port for the driver to listen on, which is used for the driver to communicate with the executor.	Range of Random Ports

Range of Random Ports

All random ports must be within a certain range.

Table 26-24 Parameter description

Parameter	Description	Default Value
spark.random.port.min	Sets the minimum random port.	22600
spark.random.port.max	Sets the maximum random port.	22899

TIMEOUT

By default, computation tasks that can well process medium-scale data are configured in Spark. However, if the data volume is too large, the tasks may fail due to timeout. In the scenario with a large amount of data, the timeout parameter in Spark needs to be assigned a larger value.

Table 26-25 Parameter description

Parameter	Description	Default Value
spark.files.fetchTimeout	Specifies the communication timeout (in seconds) when fetching files added using SparkContext.addFile() of the driver.	60s
spark.network.timeout	Specifies the default timeout for all network interactions, in seconds. You can use this parameter to replace spark.core.connection.ack.wait.timeout , spark.akka.timeout , spark.storage.blockManagerSlaveTimeoutMs , or spark.shuffle.io.connectionTimeout .	360s
spark.core.connection.ack.wait.timeout	Specifies the timeout for a connection to wait for a response, in seconds. To avoid long-time waiting caused by GC, you can set this parameter to a larger value.	60

Encryption

Spark supports SSL for Akka and HTTP (for the broadcast and file server) protocols, but does not support SSL for the web UI and block transfer service.

SSL must be configured on each node and configured for each component involved in communication using a particular protocol.

Table 26-26 Parameter description

Parameter	Description	Default Value
spark.ssl.enabled	Indicates whether to enable SSL connections for all supported protocols. All SSL settings similar to spark.ssl.xxx indicate the global configuration of all supported protocols. To override the global configuration of a particular protocol, you must override the property in the namespace specified by the protocol. Use spark.ssl.YYY.XXX to overwrite the global configuration of the particular protocol specified by YYY . YYY can be either akka for Akka-based connections or fs for the broadcast and file server.	false
spark.ssl.enabledAlgorithms	Indicates the comma-separated list of passwords. The specified passwords must be supported by the JVM.	-
spark.ssl.keyPassword	Specifies the password of a private key in the keystore.	-
spark.ssl.keystore	Specifies the path of the keystore file. The path can be absolute or relative to the directory where the component is started.	-
spark.ssl.keystorePassword	Specifies the password of the keystore.	-
spark.ssl.protocol	Specifies the protocol name. This protocol must be supported by the JVM. The reference list of protocols is available on this page.	-
spark.ssl.trustStore	Specifies the path of the truststore file. The path can be absolute or relative to the directory where the component is started.	-
spark.ssl.trustStorePassword	Specifies the password of the truststore.	-

Security

Spark supports shared key-based authentication. You can use **spark.authenticate** to configure authentication. This parameter controls whether the Spark communication protocol uses the shared key for authentication. This authentication is a basic handshake that ensures that both sides have the same shared key and are allowed to communicate. If the shared keys are different, the communication is not allowed. You can create shared keys as follows:

- For Spark on Yarn deployments, set **spark.authenticate** to **true**. Then, shared keys are automatically generated and distributed. Each application exclusively occupies a shared key.

- For other types of Spark deployments, configure Spark parameter **spark.authenticate.secret** on each node. All masters, workers, and applications use this key.

Table 26-27 Parameter description

Parameter	Description	Default Value
spark.acls.enable	Indicates whether to enable Spark ACLs. If Spark ACLs are enabled, the system checks whether the user has the permission to access and modify jobs. Note that this requires the user to be identifiable. If the user is identified as invalid, the check will not be performed. Filters can be used to verify and set users on the UI.	true
spark.admin.acls	Specifies the comma-separated list of users/administrators that have the permissions to view and modify all Spark jobs. This list can be used if you are running on a shared cluster and working with the help of an MRS cluster administrator or developer.	admin
spark.authenticate	Indicates whether Spark authenticates its internal connections. If the application is not running on Yarn, see spark.authenticate.secret .	true
spark.authenticate.secret	Sets the key for authentication between Spark components. This parameter must be set if Spark does not run on Yarn and authentication is disabled.	-
spark.modify.acls	Specifies the comma-separated list of users who have the permission to modify Spark jobs. By default, only users who have enabled Spark jobs have the permission to modify the list (for example, delete the list).	-
spark.ui.view.acls	Specifies the comma-separated list of users who have the permission to access the Spark web UI. By default, only users who have enabled Spark jobs have the access permission.	-

Enabling the Authentication Mechanism Between Spark Processes

Spark processes support shared key-based authentication. You can configure **spark.authenticate** to control whether Spark performs authentication during communication. In this authentication mode, the two communication parties share the same key only using simple handshakes.

Configure the following parameters in the **spark-defaults.conf** file on the Spark client.

Table 26-28 Parameter description

Parameter	Description	Default Value
spark.authenticate	For Spark on Yarn deployments, set this parameter to true . Then, keys are automatically generated and distributed, and each application uses a unique key.	true

Compression

Data compression is policy that optimizes memory usage at the expense of CPU. Therefore, when the Spark memory is severely insufficient (this issue is common due to the characteristics of in-memory computing), data compression can greatly improve performance. Spark supports three types of compression algorithm: Snappy, LZ4, and LZF. Snappy is the default compression algorithm and invokes the native method to compress and decompress data. In Yarn mode, pay attention to the impact of non-heap memory on the container process.

Table 26-29 Parameter description

Parameter	Description	Default Value
spark.io.compression.codec	Indicates the codec for compressing internal data, such as RDD partitions, broadcast variables, and shuffle output. By default, Spark supports three types of compression algorithm: LZ4, LZF, and Snappy. You can specify algorithms using fully qualified class names, such as org.apache.spark.io.LZ4CompressionCodec , org.apache.spark.io.LZFCompressionCodec , and org.apache.spark.io.SnappyCompressionCodec .	lz4
spark.io.compression.lz4.block.size	Indicates the block size (bytes) used in LZ4 compression when the LZ4 compression algorithm is used. When LZ4 is used, reducing the block size also reduces the shuffle memory usage.	32768
spark.io.compression.snappy.block.size	Indicates the block size (bytes) used in Snappy compression when the Snappy compression algorithm is used. When Snappy is used, reducing the block size also reduces the shuffle memory usage.	32768
spark.shuffle.compress	Indicates whether to compress the output files of a Map task. You are advised to compress the broadcast variables. using spark.io.compression.codec .	true

Parameter	Description	Default Value
spark.shuffle.spill.compress	Indicates whether to compress the data overflowed during shuffle using spark.io.compression.codec .	true
spark.eventLog.compress	Indicates whether to compress logged events when spark.eventLog.enabled is set to true .	false
spark.broadcast.compress	Indicates whether to compress broadcast variables before sending them. You are advised to compress the broadcast variables.	true
spark.rdd.compress	Indicates whether to compress serialized RDD partitions (for example, the StorageLevel.MEMORY_ONLY_SER partition). Substantial space can be saved at the cost of some extra CPU time.	false

Reducing the Probability of Abnormal Client Application Operations When Resources Are Insufficient

When resources are insufficient, ApplicationMaster tasks must wait and will not be processed until enough resources are available for use. If the actual waiting time exceeds the configured waiting time, the ApplicationMaster tasks will be deleted. Adjust the following parameters to reduce the probability of abnormal client application operation.

Configure the following parameters in the **spark-defaults.conf** file on the client.

Table 26-30 Parameter description

Parameter	Description	Default Value
spark.yarn.applicationMaster.waitTries	Specifies the number of the times that ApplicationMaster waits for Spark master, which is also the times that ApplicationMaster waits for SparkContext initialization. Enlarge this parameter value to prevent ApplicationMaster tasks from being deleted and reduce the probability of abnormal client application operations.	10
spark.yarn.am.memory	Specifies the ApplicationMaster memory. Enlarge this parameter value to prevent ApplicationMaster tasks from being deleted by ResourceManager due to insufficient memory and reduce the probability of abnormal client application operations.	1G

26.2.4 Spark on HBase Overview and Basic Applications

Scenario

Spark on HBase allows users to query HBase tables in Spark SQL and to store data for HBase tables by using the Beeline tool. You can use HBase APIs to create, read data from, and insert data into tables.

Procedure

- Step 1** Log in to Manager and choose **Cluster** > *Name of the desired cluster* > **Cluster Properties** to check whether the cluster is in security mode.
- If yes, go to [Step 2](#).
 - If no, go to [Step 5](#).
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration** > **All Configurations** > **JDBCServer2x** > **Default**, and modify the following parameter.

Table 26-31 Parameter list 1

Parameter	Default Value	Changed To
spark.yarn.security.credentials.hbase.enabled	false	true

NOTE

To ensure that Spark2x can access HBase for a long time, do not modify the following parameters of the HBase and HDFS services:

- dfs.namenode.delegation.token.renew-interval
- dfs.namenode.delegation.token.max-lifetime
- hbase.auth.key.update.interval
- hbase.auth.token.max.lifetime (The value is fixed to **604800000** ms, that is, 7 days.)

If the preceding parameter configuration must be modified based on service requirements, ensure that the value of the HDFS parameter **dfs.namenode.delegation.token.renew-interval** is not greater than the values of the HBase parameters **hbase.auth.key.update.interval**, **hbase.auth.token.max.lifetime**, and **dfs.namenode.delegation.token.max-lifetime**.

- Step 3** Choose **SparkResource2x** > **Default** and modify the following parameters.

Table 26-32 Parameter list 2

Parameter	Default Value	Changed To
spark.yarn.security.credentials.hbase.enabled	false	true

Step 4 Restart the Spark2x service for the configuration to take effect.

 **NOTE**

To use the Spark on HBase function on the Spark2x client, you need to download and install the Spark2x client again.

Step 5 On the Spark2x client, use the spark-sql or spark-beeline connection to query tables created by Hive on HBase. You can create an HBase table by running SQL commands or create an external table to associate the HBase table. Before creating tables, ensure that HBase tables exist in HBase. The HBase table **table1** is used as an example.

1. Run the following commands to create the HBase table using the Beeline tool:

```
create table hbaseTable
(
  id string,
  name string,
  age int
)
using org.apache.spark.sql.hbase.HBaseSource
options(
  hbaseTableName "table1",
  keyCols "id",
  colsMapping "
name=cf1.cq1,
age=cf1.cq2
");
```

 **NOTE**

- **hbaseTable**: name of the created Spark table
 - **id string, name string, age int**: field name and field type of the Spark table
 - **table1**: name of the HBase table
 - *id*: row key column name of the HBase table
 - *name=cf1.cq1, age=cf1.cq2*: mapping between columns in the Spark table and columns in the HBase table. The **name** column of the Spark table maps the **cq1** column in the **cf1** column family of the HBase table, and the **age** column of the Spark table maps the **cq2** column in the **cf1** column family of the HBase table.
2. Run the following command to import data to the HBase table using a CSV file:

```
hbase org.apache.hadoop.hbase.mapreduce.ImportTsv -  
Dimporttsv.separator="," -  
Dimporttsv.columns=HBASE_ROW_KEY,cf1:cq1,cf1:cq2,cf1:cq3,cf1:cq4,cf1:cq5  
table1 /hperson
```

Where **table1** indicates the name of the HBase table, and **/hperson** indicates the path where the CSV file is stored.

- Run the following command to query data in spark-sql or spark-beeline, where *hbaseTable* is the corresponding Spark table name: The command is as follows:

```
select * from hbaseTable;
```

----End

26.2.5 Spark on HBase V2 Overview and Basic Applications

Scenario

Spark on HBase V2 allows users to query HBase tables in Spark SQL and to store data for HBase tables by using the Beeline tool. You can use HBase APIs to create, read data from, and insert data into tables.

Procedure

- Log in to Manager and choose **Cluster** > *Name of the desired cluster* > **Cluster Properties** to check whether the cluster is in security mode.
 - If yes, go to [Step 2](#).
 - If no, go to [Step 5](#).
- Choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration** > **All Configurations** > **JDBCServer2x** > **Default**, and modify the following parameter.

Table 26-33 Parameter list 1

Parameter	Default Value	Changed To
spark.yarn.security.credentials.hbase.enabled	false	true

NOTE

To ensure that Spark2x can access HBase for a long time, do not modify the following parameters of the HBase and HDFS services:

- dfs.namenode.delegation.token.renew-interval
- dfs.namenode.delegation.token.max-lifetime
- hbase.auth.key.update.interval
- hbase.auth.token.max.lifetime (The value is fixed to **604800000** ms, that is, 7 days.)

If the preceding parameter configuration must be modified based on service requirements, ensure that the value of the HDFS parameter **dfs.namenode.delegation.token.renew-interval** is not greater than the values of the HBase parameters **hbase.auth.key.update.interval**, **hbase.auth.token.max.lifetime**, and **dfs.namenode.delegation.token.max-lifetime**.

- Choose **SparkResource2x** > **Default** and modify the following parameters.

Table 26-34 Parameter list 2

Parameter	Default Value	Changed To
spark.yarn.security.credentials.hbase.enabled	false	true

Step 4 Restart the Spark2x service for the configuration to take effect.

 **NOTE**

If you need to use the Spark on HBase function on the Spark2x client, download and install the Spark2x client again.

Step 5 On the Spark2x client, use the spark-sql or spark-beeline connection to query tables created by Hive on HBase. You can create an HBase table by running SQL commands or create an external table to associate the HBase table. For details, see the following description. The following uses the HBase table **table1** as an example.

1. Run the following commands to create a table using the spark-beeline tool:

```
create table hbaseTable1
(id string, name string, age int)
using org.apache.spark.sql.hbase.HBaseSourceV2
options(
hbaseTableName "table2",
keyCols "id",
colsMapping "name=cf1.cq1,age=cf1.cq2");
```

 **NOTE**

- **hbaseTable1**: name of the created Spark table
 - **id string, name string, age int**: field name and field type of the Spark table
 - **table2**: name of the HBase table
 - **id**: row key column name of the HBase table
 - **name=cf1.cq1, age=cf1.cq2**: mapping between columns in the Spark table and columns in the HBase table. The **name** column of the Spark table maps the **cq1** column in the **cf1** column family of the HBase table, and the **age** column of the Spark table maps the **cq2** column in the **cf1** column family of the HBase table.
2. Run the following command to import data to the HBase table using a CSV file:
- ```
hbase org.apache.hadoop.hbase.mapreduce.ImportTsv -
Dimporttsv.separator="," -
Dimporttsv.columns=HBASE_ROW_KEY,cf1:cq1,cf1:cq2,cf1:cq3,cf1:cq4,cf1:cq5
table2 /hperson
```
- Where **table2** indicates the name of the HBase table, and **/hperson** indicates the path where the CSV file is stored.
3. Run the following command to query data in spark-sql or spark-beeline.

```
select * from hbaseTable1;
```

----End

## 26.2.6 SparkSQL Permission Management(Security Mode)

### 26.2.6.1 Spark SQL Permissions

#### SparkSQL Permissions

Similar to Hive, Spark SQL is a data warehouse framework built on Hadoop, providing storage of structured data like structured query language (SQL).

MRS supports users, user groups, and roles. Permission must be assigned to roles and then roles are bound to users or user groups. Users can obtain permissions only by binding a role or joining a group that is bound with a role.

#### NOTE

- If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for Spark2x](#).
- After Ranger authentication is enabled or disabled on Spark2x, you need to restart Spark2x and download the client again or update the client configuration file `spark/conf/spark-defaults.conf`.

Enable Ranger authentication: `spark.ranger.plugin.authorization.enable=true`

Disable Ranger authentication: `spark.ranger.plugin.authorization.enable=false`

#### Permission Management


Spark SQL permission management indicates the permission system for managing and controlling users' operations on databases, to ensure that different users can operate databases separately and securely. A user can operate another user's tables and databases only with the corresponding permissions. Otherwise, operations will be rejected.

Spark SQL permission management integrates the functions of Hive management. The Metastore service of Hive and the permission assignment function are required to enable Spark SQL permission management.

**Figure 26-3** shows the basic architecture of Spark SQL permission management. This architecture includes two parts: granting permissions on the page, and obtaining and judging a service.

- Granting permissions on the page: Spark SQL only supports granting permissions on the page. On FusionInsight Manager, choose **System > Permission** to add or delete a user, user group, or a role, and to grant permissions or cancel permissions.
- Obtaining and judging a service: When the DDL and DML commands are received from a client, Spark SQL will obtain the client's permissions on database information from MetaStore, and check whether the required permissions are included. If the required permissions are included, continue the execution. If the required permissions are not included, reject the user's operations. After the MetaStore permissions are checked, ACL permission also needs to be checked on HDFS.

Figure 26-3 Spark SQL permission management architecture



Additionally, Spark SQL provides column and view permissions to meet requirements of different scenarios.

- Column permission
  - Spark SQL permission control consists of metadata permission control and HDFS ACL permission control. When Hive MetaStore automatically synchronizes table permissions to the HDFS ACL, column-level permissions are not synchronized. In other words, a user with partial or all column-level permissions cannot access the entire HDFS file using the HDFS client.
    - In **spark-sql** mode, users with only column-level permissions cannot access HDFS files. Therefore, they cannot access the columns of the corresponding tables.
    - In Beeline/JDBCServer mode, permissions are assigned among users, for example, the permissions on the table created by user A are assigned to user B.
      - **hive.server2.enable.doAs=true** (configured in the **hive-site.xml** file on the Spark server)  
In this case, user B cannot query the information. You need to manually assign the read permission on the file in HDFS.
      - **hive.server2.enable.doAs=false**
        - Users A and B are connected by Beeline. User B can query the information.
        - User A creates a table using SQL statements, and user B can query the table in Beeline.

However, information query is not supported in other scenarios, for example, user A uses Beeline to create a table and user B uses SQL

to query the table, or user A uses SQL to create a table and user B uses SQL to query the table. You need to manually assign the read permission on the file in HDFS.

 **NOTE**

The **spark** user is the Spark administrator in HDFS ACL permission control. The permission control of the Beeline client user depends only on the metadata permission on Spark.

- View permission

View permission indicates the operation permission such as query and modification on the view of a table, regardless of the corresponding permission of a table. Namely, if you have the permission to query the view of a table, the permission to query the table is not mandatory. The view permission is applicable to the whole table but not to the columns.

Restrictions of view and column permissions on SparkSQL are similar. The following uses the view permission as an example:

- In spark-sql mode, if you have only the view permission but not the table permission and do not have the permission to read HDFS, you cannot access the table data stored in HDFS. That is, you cannot query the view of the table.
- In Beeline/JDBCServer mode, permissions are assigned among users, for example, the permissions on the view created by user A are assigned to user B.

- **hive.server2.enable.doAs=true** (configured in the **hive-site.xml** file on the Spark server)

In this case, user B cannot query the information. You need to manually assign the read permission on the file in HDFS.

- **hive.server2.enable.doAs=false**
  - Users A and B are connected by Beeline. User B can query the information.
  - User A creates a view using SQL statements, and user B can query the view in Beeline.

However, information query is not supported in other scenarios. For example, user A uses Beeline to create a view but user B cannot use SQL to query the view, or user A uses SQL to create a view but user B cannot use SQL to query the view. You need to manually assign the read permission on the file in HDFS.

Permission of operations on the view of a table is as follows:

- To create a view, you must have the CREATE permission on the database and the SELECT and SELECT\_of\_GRANT permissions on the tables.
- Creating and describing a view only entail the SELECT permission on the view. Querying views and tables at the same time entails the SELECT permission on other tables. For example, to perform **select \* from v1 join t1**, you must have the SELECT permission on the **v1** view and **t1** table, even through the **v1** view depends on the **t1** table.

 NOTE

In Beeline/JDBCServer mode, to query a view, you must have the SELECT permission on the tables. In spark-sql mode, to query a view, you must have the SELECT permission on the view and tables.

- Deleting and modifying a view entail the permission of owner on the view.

## SparkSQL Permission Model

If you want to perform SQL operations using SparkSQL, you must be granted with permissions of SparkSQL databases and tables (include external tables and views). The complete permission model of SparkSQL consists of the meta data permission and HDFS file permission. Permissions required to use a database or a table is just one type of SparkSQL permission.

- Metadata permissions

Metadata permissions are controlled at the metadata layer. Similar to traditional relational databases, SparkSQL databases involve the CREATE and SELECT permissions, and tables and columns involve the SELECT, INSERT, UPDATE, and DELETE permissions. SparkSQL also supports the permissions of **OWNERSHIP** and **ADMIN**.

- Data file permissions (that is, HDFS file permissions)

SparkSQL database and table files are stored in HDFS. The created databases or tables are saved in the **/user/hive/warehouse** directory of HDFS by default. The system automatically creates subdirectories named after database names and database table names. To access a database or table, you must have the **Read**, **Write** and **Execute** permissions on the corresponding file in HDFS.

To perform various operations on SparkSQL databases or tables, you need to associate the metadata permission and HDFS file permission. For example, to query SparkSQL data tables, you need to associate the metadata permission **SELECT** and HDFS file permissions **Read** and **Execute**.

Using the management function of Manager GUI to manage the permissions of SparkSQL databases and tables, only requires the configuration of metadata permission, and the system will automatically associate and configure the HDFS file permission. In this way, operations on the interface are simplified, and the efficiency is improved.

## Usage Scenarios and Related Permissions

Creating a database with SparkSQL service requires users to join in the hive group, without granting a role. Users have all permissions on the databases or tables created by themselves in Hive or HDFS. They can create tables, select, delete, insert, or update data, and grant permissions to other users to allow them to access the tables and corresponding HDFS directories and files.

A user can access the tables or database only with permissions. Users' permissions vary depending on different SparkSQL scenarios.

**Table 26-35** SparkSQL scenarios

| Typical Scenario                             | Required Permission                                                                                                                                                                                                                                                                              |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Using SparkSQL tables, columns, or databases | Permissions required in different scenarios are as follows: <ul style="list-style-type: none"> <li>• To create a table, the CREATE permission is required.</li> <li>• To query data, the SELECT permission is required.</li> <li>• To insert data, the INSERT permission is required.</li> </ul> |
| Associating and using other components       | In some scenarios, except the SparkSQL permission, other permissions may be also required. For example:<br>Using Spark on HBase to query HBase data in SparkSQL requires HBase permissions.                                                                                                      |

In some special SparkSQL scenarios, other permissions must be configured separately.

**Table 26-36** SparkSQL scenarios and required permissions

| Scenario                                                                                                                                                                                                                                | Required Permission                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Creating SparkSQL databases, tables, and external tables, or adding partitions to created Hive tables or external tables when data files specified by Hive users are saved to other HDFS directories except <b>/user/hive/warehouse</b> | <ul style="list-style-type: none"> <li>• The directory must exist, the client user must be the owner of the directory, and the user must have the <b>Read, Write, and Execute</b> permissions on the directory. The user must have the <b>Read</b> and <b>Execute</b> permissions of all the upper-layer directories of the directory.</li> <li>• If the Spark version is later than 2, the <b>Create</b> permission of the Hive database is required if you want to create a HBase table. However, in Spark 1.5, the <b>Create</b> permissions of both the Hive database and HBase namespace are required if you want to create a HBase table.</li> </ul> |

| Scenario                                                                                                                                                                                     | Required Permission                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Importing all the files or specified files in a specified directory to the table using load                                                                                                  | <ul style="list-style-type: none"> <li>The data source is a Linux local disk, the specified directory exists, and the system user <b>omm</b> has read and execute permission of the directory and all its upper-layer directories. The specified file exists, and user <b>omm</b> has the <b>Read</b> permission on the file and has the <b>Read</b> and <b>Execute</b> permissions on all the upper-layer directories of the file.</li> <li>The data source is HDFS, the specified directory exists, and the SparkSQL user is the owner of the directory and has the <b>Read</b>, <b>Write</b>, and <b>Execute</b> permissions on the directory and its subdirectories, and has the <b>Read</b> and <b>Execute</b> permissions on all its upper-layer directories. The specified file exists, and the SparkSQL user is the owner of the file and has the <b>Read</b>, <b>Write</b>, and <b>Execute</b> permissions on the file and has the <b>Read</b> and <b>Execute</b> permissions on all its upper-layer directories.</li> </ul> |
| Creating or deleting functions or modifying any database                                                                                                                                     | The <b>ADMIN</b> permission is required.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Performing operations on all databases and tables in Hive                                                                                                                                    | The user must be added to the <b>supergroup</b> user group, and be assigned the <b>ADMIN</b> permission.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| After assigning the <b>Insert</b> permission on some DataSource tables, assigning the <b>Write</b> permission on table directories in HDFS before performing the insert or analyze operation | When the <b>Insert</b> permission is assigned to the <b>spark datasource</b> table, if the table format is text, CSV, JSON, Parquet, or ORC, the permission on the table directory is not changed. After the <b>Insert</b> permission is assigned to the DataSource table of the preceding formats, you need to assign the <b>Write</b> permission to the table directories in HDFS separately so that users can perform the insert or analyze operation on the tables.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

### 26.2.6.2 Creating a Spark SQL Role

#### Scenario

This section describes how to create and configure a SparkSQL role on Manager. The Spark SQL role can be configured with the Spark administrator permission or the permission of performing operations on the table data.

Creating a database with Hive requires users to join in the **hive** group, without granting a role. Users have all permissions on the databases or tables created by themselves in Hive or HDFS. They can create tables, select, delete, insert, or update data, and grant permissions to other users to allow them to access the tables and corresponding HDFS directories and files. The created databases or tables are saved in the **/user/hive/warehouse** directory of HDFS by default.



 NOTE

- If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. For details, see [Adding a Ranger Access Permission Policy for Spark2x](#).
- After Ranger authentication is enabled or disabled on Spark2x, you need to restart Spark2x and download the client again or update the client configuration file `spark/conf/spark-defaults.conf`.

Enable Ranger authentication: `spark.ranger.plugin.authorization.enable=true`

Disable Ranger authentication: `spark.ranger.plugin.authorization.enable=false`

## Procedure

1. Log in to Manager, and choose **System > Permission > Role**.
2. Click **Create Role** and set a role name and enter description.
3. Set **Configure Resource Permission**. For details, see [Table 26-37](#).
  - **Hive Admin Privilege**: Hive administrator permissions.
  - **Hive Read Write Privileges**: Hive data table management permission, which is the operation permission to set and manage the data of created tables.

 NOTE

- Hive role management supports Hive administrator permissions and the permissions to access tables and views, but does not support granting permissions on databases.
- The permissions of the Hive administrator do not include the permission to manage HDFS.
- If there are too many tables in the database or too many files in tables, the permission granting may last a while. For example, if a table contains 10,000 files, the permission granting lasts about 2 minutes.

**Table 26-37** Setting a role

| Task                                 | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Hive administrator permission</p> | <p>In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Hive</b> and select <b>Hive Admin Privilege</b>.</p> <p>After being bound to the Hive administrator role, perform the following operations during each maintenance operation:</p> <ol style="list-style-type: none"> <li>1. Log in to the node where the Spark2x client is installed as the client installation user.</li> <li>2. Run the following command to configure environment variables:<br/>For example, if the Spark2x client installation directory is <code>/opt/client</code>, run <b>source /opt/client/bigdata_env</b>.<br/><b>source /opt/client/Spark2x/component_env</b></li> <li>3. Run the following command to perform user authentication:<br/><b>kinit Hive service user</b></li> <li>4. Run the following command to log in to the client tool:<br/><b>/opt/client/Spark2x/spark/bin/beeline -u "jdbc:hive2://&lt;zkNode1_IP&gt;:&lt;zkNode1_Port&gt;,&lt;zkNode2_IP&gt;:&lt;zkNode2_Port&gt;,&lt;zkNode3_IP&gt;:&lt;zkNode3_Port&gt;/;serviceDiscovery-Mode=zooKeeper;zooKeeperNamespace=sparkthriftserver2x;user.principal=spark2x/hadoop.&lt;System domain name&gt;@&lt;System domain name&gt;;sasLQop=auth-conf;auth=KERBEROS;principal=spark2x/hadoop.&lt;System domain name&gt;@&lt;System domain name&gt;;"</b></li> </ol> |

| Task                                                                                     | Operation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                          | <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• <code>&lt;zkNode1_IP&gt;:&lt;zkNode1_Port&gt;</code>, <code>&lt;zkNode2_IP&gt;:&lt;zkNode2_Port&gt;</code>, <code>&lt;zkNode3_IP&gt;:&lt;zkNode3_Port&gt;</code> indicates the ZooKeeper URL, for example, 192.168.81.37:2181,192.168.195.232:2181,192.168.169.84:2181.</li> <li>• <code>sparkthriftserver</code> indicates a ZooKeeper directory, from which a random TriftServer or ProxyThriftServer is connected by the client.</li> <li>• You can log in to Manager, choose <b>System &gt; Permission &gt; Domain and Mutual Trust</b>, and view the value of <b>Local Domain</b>, which is the current system domain name. <code>spark2x/hadoop.&lt;System domain name&gt;</code> is the username. All letters in the system domain name contained in the username are lowercase letters. For example, <b>Local Domain</b> is set to <code>9427068F-6EFA-4833-B43E-60CB641E5B6C.COM</code>, and the username is <code>spark2x/hadoo.9427068f-6efa-4833-b43e-60cb641e5b6c.com</code>.</li> </ul> <p>5. Run the following command to update the administrator permissions:<br/><b>set role admin;</b></p> |
| Setting the permission to query a table of another user in the default database          | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Hive</b> &gt; <b>Hive Read Write Privileges</b>.</li> <li>2. Click the name of the specified database in the database list. Tables in the database are displayed.</li> <li>3. In the <b>Permission</b> column of the specified table, select <b>SELECT</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Setting the permission to import data to a table of another user in the default database | <ol style="list-style-type: none"> <li>1. In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> &gt; <b>Hive</b> &gt; <b>Hive Read Write Privileges</b>.</li> <li>2. Click the name of the specified database in the database list. Tables in the database are displayed.</li> <li>3. In the <b>Permission</b> column of the specified table, select <b>DELETE</b> and <b>INSERT</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

4. Click **OK**.

### 26.2.6.3 Configuring Permissions for SparkSQL Tables, Columns, and Databases

#### Scenario

You can configure related permissions if you need to access tables or databases created by other users. SparkSQL supports column-based permission control. If a user needs to access some columns in tables created by other users, the user must be granted the permission for columns. The following describes how to grant table, column, and database permissions to users by using the role management function of Manager.

#### Procedure

The operations for granting permissions on SparkSQL tables, columns, and databases are the same as those for Hive. For details, see [Permission Management](#).

#### NOTE

- Any permission for a table in the database is automatically associated with the HDFS permission for the database directory to facilitate permission management. When any permission for a table is canceled, the system does not automatically cancel the HDFS permission for the database directory to ensure performance. In this case, users can only log in to the database and view table names.
- When the query permission on a database is added to or deleted from a role, the query permission on tables in the database is automatically added to or deleted from the role. This mechanism is inherited from Hive.
- In Spark, the column name of the struct data type cannot contain special characters, that is, characters other than letters, digits, and underscores (\_). If the column name of the struct data type contains special characters, the column cannot be displayed on the FusionInsight Manager console when you grant permissions to roles on the role page.

#### Concepts

SparkSQL statements are processed in SparkSQL. [Table 26-38](#) describes the permission requirements.

**Table 26-38** Scenarios of using SparkSQL tables, columns, or databases

| Scenario              | Required Permission                                                                                                                                                                                                                |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CREATE TABLE          | <b>CREATE</b> , RWX+ownership (for creating external tables - the location)<br><b>NOTE</b><br>When creating datasource tables in a specified file path, the RWX and ownership permission on the file next to the path is required. |
| DROP TABLE            | <b>Ownership</b> (of table)                                                                                                                                                                                                        |
| DROP TABLE PROPERTIES | <b>Ownership</b>                                                                                                                                                                                                                   |
| DESCRIBE TABLE        | <b>Select</b>                                                                                                                                                                                                                      |

| Scenario                                       | Required Permission                                                                                                                                         |
|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SHOW PARTITIONS                                | <b>Select</b>                                                                                                                                               |
| ALTER TABLE LOCATION                           | <b>Ownership</b> , RWX+ownership (for new location)                                                                                                         |
| ALTER PARTITION LOCATION                       | <b>Ownership</b> , RWX+ownership (for new partition location)                                                                                               |
| ALTER TABLE ADD PARTITION                      | <b>Insert</b> , RWX and ownership (for partition location)                                                                                                  |
| ALTER TABLE DROP PARTITION                     | <b>Delete</b>                                                                                                                                               |
| ALTER TABLE(all of them except the ones above) | <b>Update, Ownership</b>                                                                                                                                    |
| TRUNCATE TABLE                                 | <b>Ownership</b>                                                                                                                                            |
| CREATE VIEW                                    | <b>Select, Grant Of Select, CREATE</b>                                                                                                                      |
| ALTER VIEW PROPERTIES                          | <b>Ownership</b>                                                                                                                                            |
| ALTER VIEW RENAME                              | <b>Ownership</b>                                                                                                                                            |
| ALTER VIEW ADD PARTS                           | <b>Ownership</b>                                                                                                                                            |
| ALTER VIEW AS                                  | <b>Ownership</b>                                                                                                                                            |
| ALTER VIEW DROPPARTS                           | <b>Ownership</b>                                                                                                                                            |
| ANALYZE TABLE                                  | <b>Search, Insert</b>                                                                                                                                       |
| SHOW COLUMNS                                   | <b>Select</b>                                                                                                                                               |
| SHOW TABLE PROPERTIES                          | <b>Select</b>                                                                                                                                               |
| CREATE TABLE AS SELECT                         | <b>Select, CREATE</b>                                                                                                                                       |
| SELECT                                         | <b>Select</b><br><b>NOTE</b><br>The same as tables, you need to have the <b>Select</b> permission on a view when performing a SELECT operation on the view. |
| INSERT                                         | <b>Insert, Delete (for overwrite)</b>                                                                                                                       |
| LOAD                                           | <b>Insert, Delete</b> , RWX+ownership(input location)                                                                                                       |
| SHOW CREATE TABLE                              | <b>Select, Grant Of Select</b>                                                                                                                              |
| CREATE FUNCTION                                | <b>ADMIN</b>                                                                                                                                                |
| DROP FUNCTION                                  | <b>ADMIN</b>                                                                                                                                                |
| DESC FUNCTION                                  | -                                                                                                                                                           |
| SHOW FUNCTIONS                                 | -                                                                                                                                                           |

| Scenario               | Required Permission |
|------------------------|---------------------|
| MSCK (metastore check) | <b>Ownership</b>    |
| ALTER DATABASE         | <b>ADMIN</b>        |
| CREATE DATABASE        | -                   |
| SHOW DATABASES         | -                   |
| EXPLAIN                | <b>Select</b>       |
| DROP DATABASE          | <b>Ownership</b>    |
| DESC DATABASE          | -                   |
| CACHE TABLE            | <b>Select</b>       |
| UNCACHE TABLE          | <b>Select</b>       |
| CLEAR CACHE TABLE      | <b>ADMIN</b>        |
| REFRESH TABLE          | <b>Select</b>       |
| ADD FILE               | <b>ADMIN</b>        |
| ADD JAR                | <b>ADMIN</b>        |
| HEALTHCHECK            | -                   |

## 26.2.6.4 Configuring Permissions for SparkSQL to Use Other Components

### Scenario

SparkSQL may need to be associated with other components. For example, Spark on HBase requires HBase permissions. The following describes how to associate SparkSQL with HBase.

### Prerequisites

- The Spark client has been installed in a directory, for example, **/opt/client**.
- You have obtained a user account with the MRS cluster administrator permissions, for example, **admin**.

### Procedure

- **Spark on HBase authorization**  
After the permissions are assigned, you can use statements that are similar to SQL statements to access HBase tables from SparkSQL. The following uses the procedure for assigning a user the permissions to query HBase tables as an example.

#### NOTE

Set `spark.yarn.security.credentials.hbase.enabled` to `true`.

- a. On Manager, create a role, for example, **hive\_hbase\_create**, and grant the permission to create HBase tables to the role.

In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **HBase** > **HBase Scope** > **global**. Select **create** of the namespace **default**, and click **OK**.

 **NOTE**

In this example, the created table is saved in the default database of Hive and has the CREATE permission of the default database. If you save the table to a Hive database other than **default**, perform the following operations:

In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **Hive** > **Hive Read Write Privileges**, select **CREATE** for the desired database, and click **OK**.

- b. On Manager, create a role, for example, **hive\_hbase\_submit**, and grant the permission to submit tasks to the Yarn queue.

In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **Yarn** > **Scheduling Queue** > **root**. Select **Submit** of **default**, and click **OK**.

- c. On Manager, create a human-machine user, for example, **hbase\_creates\_user**, add the user to the **hive** group, and bind the **hive\_hbase\_create** and **hive\_hbase\_submit** roles to create SparkSQL and HBase tables.

- d. Log in to the node where the client is installed as the client installation user.

- e. Run the following command to configure environment variables:

```
source /opt/client/bigdata_env
source /opt/client/Spark2x/component_env
```

- f. Run the following command to authenticate the user:

```
kinit hbase_creates_user
```

- g. Run the following commands to enter the shell environment on the Spark JDBCServer client:

```
/opt/client/Spark2x/spark/bin/beeline -u "jdbc:hive2://
<zkNode1_IP>:<zkNode1_Port>,<zkNode2_IP>:<zkNode2_Port>,<zkNode3_IP>:<zkNode3_Port>";serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=sparkthriftserver2x;user.principal=spark2x/hadoop.<System domain name>@<System domain name>;sasLQop=auth-conf;auth=KERBEROS;principal=spark2x/hadoop.<System domain name>@<System domain name>;"
```

- h. Run the following command to create a table in SparkSQL and HBase, for example, create the **hbaseTable** table:

```
create table hbaseTable (id string, name string, age int) using
org.apache.spark.sql.hbase.HBaseSource options (hbaseTableName
"table1", keyCols "id", colsMapping = "", name=cf1.cq1, age=cf1.cq2");
```

The created SparkSQL table and the HBase table are stored in the Hive database **default** and the HBase namespace **default**, respectively.

- i. On Manager, create a role, for example, **hive\_hbase\_select**, and grant the role the permission to query SparkSQL on HBase table **hbaseTable** and HBase table **hbaseTable**.

- In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **HBase** > **HBase Scope** > **global** > **default**. Select **read** for the **hbaseTable** table, and click **OK** to grant the table query permission to the HBase role.
- Edit the role. In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **HBase** > **HBase Scope** > **global** > **hbase**. Select **Execute** for **hbase:meta**, and click **OK**.
- Edit the role. In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **Hive** > **Hive Read Write Privileges** > **default**. Select **SELECT** for the **hbaseTable** table, and click **OK**.
- j. On Manager, create a human-machine user, for example, **hbase\_select\_user**, add the user to the **hive** group, and bind the **hive\_hbase\_select** role to the user for querying SparkSQL and HBase tables.
- k. Run the following command to configure environment variables:  
**source /opt/client/bigdata\_env**  
**source /opt/client/Spark2x/component\_env**
- l. Run the following command to authenticate users:  
**kinit hbase\_select\_user**
- m. Run the following commands to enter the shell environment on the Spark JDBCServer client:  
**/opt/client/Spark2x/spark/bin/beeline -u "jdbc:hive2://**  
**<zkNode1\_IP>:<zkNode1\_Port>,<zkNode2\_IP>:<zkNode2\_Port>,<zkNode3\_IP>:<zkNode3\_Port>;serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=sparkthriftserver2x;user.principal=spark2x/hadoop.<System domain name>@<System domain name>;sasLQop=auth-conf;auth=KERBEROS;principal=spark2x/hadoop.<System domain name>@<System domain name>";**
- n. Run the following command to use a SparkSQL statement to query HBase table data:  
**select \* from hbaseTable;**

### 26.2.6.5 Configuring the Client and Server

This section describes how to configure SparkSQL permission management functions (client configuration is similar to server configuration). To enable table permission, add following configurations on the client and server:

- **spark-defaults.conf** configuration file



**Table 26-39** Parameter description (1)

| Parameter                       | Description                                                                                                                                                                              | Default Value |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.authorization.enabled | Specifies whether to enable permission authentication of the datasource statement. It is recommended that the parameter value be set to <b>true</b> to enable permission authentication. | true          |

- **hive-site.xml** configuration file

**Table 26-40** Parameter description (2)

| Parameter                            | Description                                                                                                                                                     | Default Value                                                           |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| hive.metastore.uris                  | Specifies the MetaStore service address of the Hive component, for example, <b>thrift://10.10.169.84:21088,thrift://10.10.81.37:21088</b> .                     | -                                                                       |
| hive.metastore.sasl.enabled          | Specifies whether the MetaStore service uses SASL to improve security. The table permission function must be enabled.                                           | true                                                                    |
| hive.metastore.kerberos.principal    | Specifies the principal of the MetaStore service in the Hive component, for example, <b>hive/hadoop.&lt;system domain name&gt;@&lt;system domain name&gt;</b> . | hive-metastore / _HOST@EXAMPLE.COM                                      |
| hive.metastore.thrift.sasl.qop       | After the SparkSQL permission management function is enabled, set the parameter to <b>auth-conf</b> .                                                           | auth-conf                                                               |
| hive.metastore.token.signature       | Specifies the token identifier of the MetaStore service, which is set to <b>HiveServer2ImpersonationToken</b> .                                                 | HiveServer2ImpersonationToken                                           |
| hive.security.authentication.manager | Specifies the manager authenticated by the Hive client, which is set to <b>org.apache.hadoop.hive.ql.security.SessionStateUserGroupAuthenticator</b> .          | org.apache.hadoop.hive.ql.security.SessionStateUserMSGroupAuthenticator |
| hive.security.authorization.enabled  | Specifies whether to enable client authentication, which is set to <b>true</b> .                                                                                | true                                                                    |

| Parameter                                            | Description                                                                                              | Default Value |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------|---------------|
| hive.security.authorization.createtable.owner.grants | Specifies which permissions are granted to the owner who creates the table, which is set to <b>ALL</b> . | ALL           |

- **core-site.xml** configuration file of the MetaStore service

**Table 26-41** Parameter description (3)

| Parameter                     | Description                                                                                                         | Default Value |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------|---------------|
| hadoop.proxyuser.spark.hosts  | Specifies the hosts from which Spark users can be masqueraded, which is set to *, indicating all hosts.             | -             |
| hadoop.proxyuser.spark.groups | Specifies the user groups from which Spark users can be masqueraded, which is set to *, indicating all user groups. | -             |

## 26.2.7 Scenario-Specific Configuration

### 26.2.7.1 Configuring Multi-active Instance Mode

#### Scenarios

In this mode, multiple ThriftServers coexist in the cluster and the client can randomly connect any ThriftServer to perform service operations. When one or multiple ThriftServers stop working, a client can connect to another functional ThriftServer.

#### Configuration Description

Log in to Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, click **All Configurations**, and search for and modify the following parameters.

**Table 26-42** Parameter description

| Parameter                                       | Description                                                                                                      | Default Value |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------|---------------|
| spark.thriftserver.zookeeper.connection.timeout | Specifies the timeout interval of connection between ZooKeeper client and ThriftServer. The unit is millisecond. | 60000         |

| Parameter                                    | Description                                                                            | Default Value |
|----------------------------------------------|----------------------------------------------------------------------------------------|---------------|
| spark.thriftserver.zookeeper.session.timeout | Specifies the timeout interval of a ZooKeeper client session. The unit is millisecond. | 90000         |
| spark.thriftserver.zookeeper.retry.times     | Specifies the retry times after ZooKeeper disconnection.                               | 3             |
| spark.yarn.queue                             | Specifies the Yarn queue where the JDBCServer service resides.                         | default       |

## 26.2.7.2 Configuring the Multi-tenant Mode

### Scenarios

In multi-tenant mode, JDBCServer are bound with tenants. Each tenant corresponds to one or more JDBCServer, and a JDBCServer provides services for only one tenant. Different tenants can be configured with different Yarn queues to implement resource isolation.

### Configuration Description

Log in to Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, click **All Configurations**, and search for and modify the following parameters.

**Table 26-43** Parameter description

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                | Default Value                                                                                   |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| spark.proxyserver.hash.enabled   | Specifies whether to connect to ProxyServer using the Hash algorithm. <ul style="list-style-type: none"> <li><b>true</b> indicates using the Hash algorithm. In multi-tenant mode, this parameter must be configured to <b>true</b>.</li> <li><b>false</b> indicates using random connection. In multi-active instance mode, this parameter must be configured to <b>false</b>.</li> </ul> | true<br><b>NOTE</b><br>After this parameter is modified, you need to download the client again. |
| spark.thriftserver.proxy.enabled | Specifies whether to use the multi-tenant mode. <ul style="list-style-type: none"> <li><b>false</b>: The multi-instance mode is used.</li> <li><b>true</b>: The multi-tenant mode is used.</li> </ul>                                                                                                                                                                                      | true                                                                                            |

| Parameter                                          | Description                                                                                                                                                                                                                                                                                  | Default Value |
|----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.thriftserver.proxy.maxThriftServerPerTenancy | Specifies the maximum number of JDBCServer instances that can be started by a tenant in multi-tenant mode.                                                                                                                                                                                   | 1             |
| spark.thriftserver.proxy.maxSessionPerThriftServer | Specifies the maximum number of sessions in a single JDBCServer instance in multi-tenant mode. If the number of sessions exceeds this value and the number of JDBCServer instances does not exceed the upper limit, a new JDBCServer instance is started. Otherwise, an alarm log is output. | 50            |
| spark.thriftserver.proxy.sessionWaitTime           | Specifies the wait time before a JDBCServer instance is stopped when it has no session connections in multi-tenant mode.                                                                                                                                                                     | 180000        |
| spark.thriftserver.proxy.sessionThreshold          | In multi-tenant mode, when the session usage (formula: number of current sessions / spark.thriftserver.proxy.maxSessionPerThriftServer x number of current JDBCServer instances) of the JDBCServer instance reaches the threshold, a new JDBCServer instance is automatically added.         | 100           |
| spark.thriftserver.proxy.healthcheck.period        | Specifies the period of JDBCServer health checks conducted by the JDBCServer proxy in multi-tenant mode.                                                                                                                                                                                     | 60000         |
| spark.thriftserver.proxy.healthcheck.recheckTimes  | Specifies the number of JDBCServer health check retries conducted by the JDBCServer proxy in multi-tenant mode.                                                                                                                                                                              | 3             |
| spark.thriftserver.proxy.healthcheck.waitTime      | Specifies the wait time for JDBCServer to respond to a health check request sent by the JDBCServer proxy.                                                                                                                                                                                    | 10000         |
| spark.thriftserver.proxy.session.check.interval    | Specifies the period of JDBCServer proxy sessions in multi-tenant mode.                                                                                                                                                                                                                      | 6h            |
| spark.thriftserver.proxy.idle.session.timeout      | Specifies the idle time interval of a JDBCServer proxy session in multi-tenant mode. If no operation is performed within this period, the session is closed.                                                                                                                                 | 7d            |

| Parameter                                             | Description                                                                                                                                           | Default Value |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.thriftserver.proxy.idle.session.check.operation | Specifies whether to check that operations still exist on a JDBCServer proxy session when the session is checked for expiration in multi-tenant mode. | true          |
| spark.thriftserver.proxy.idle.operation.timeout       | Specifies the timeout interval of an operation in multi-tenant mode. An operation that times out is closed.                                           | 5d            |

### 26.2.7.3 Configuring the Switchover Between the Multi-active Instance Mode and the Multi-tenant Mode

#### Scenarios

When using a cluster, if you want to switch between multi-active instance mode and multi-tenant mode, the following configurations are required.

- Switch from multi-tenant mode to multi-active instance mode.  
Modify the following parameters of the Spark2x service:
  - spark.thriftserver.proxy.enabled=false
  - spark.scheduler.allocation.file=#{conf\_dir}/fairscheduler.xml
  - spark.proxyserver.hash.enabled=false
- Switch from multi-active instance mode to multi-tenant mode.  
Modify the following parameters of the Spark2x service:
  - spark.thriftserver.proxy.enabled=true
  - spark.scheduler.allocation.file=./\_\_spark\_conf\_\_/\_\_hadoop\_conf\_\_/fairscheduler.xml
  - spark.proxyserver.hash.enabled=true

#### Configuration Description

Log in to Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration**, click **All Configurations**, and search for and modify the following parameters.

**Table 26-44** Parameter description

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                                                                                | Default Value                                                                                   |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| spark.thriftserver.proxy.enabled | Specifies whether to use the multi-tenant mode. <ul style="list-style-type: none"> <li><b>false</b>: The multi-instance mode is used.</li> <li><b>true</b>: The multi-tenant mode is used.</li> </ul>                                                                                                                                                                                      | true                                                                                            |
| spark.scheduler.allocation.file  | Specifies the fair scheduling file path. <ul style="list-style-type: none"> <li>If the multi-active instance mode is used, the path is changed to <b>#{conf_dir}/fairscheduler.xml</b>.</li> <li>If multi-tenant mode is used, the path is changed to <b>./__spark_conf__/_hadoop_conf_/fairscheduler.xml</b>.</li> </ul>                                                                  | ./__spark_conf__/_hadoop_conf_/fairscheduler.xml                                                |
| spark.proxyserver.hash.enabled   | Specifies whether to connect to ProxyServer using the Hash algorithm. <ul style="list-style-type: none"> <li><b>true</b> indicates using the Hash algorithm. In multi-tenant mode, this parameter must be configured to <b>true</b>.</li> <li><b>false</b> indicates using random connection. In multi-active instance mode, this parameter must be configured to <b>false</b>.</li> </ul> | true<br><b>NOTE</b><br>After this parameter is modified, you need to download the client again. |

## 26.2.7.4 Configuring the Size of the Event Queue

### Scenarios

Functions such as UI, EventLog, and dynamic resource scheduling in Spark are implemented through event transfer. Events include SparkListenerJobStart and SparkListenerJobEnd, which record each important process.

Each event is saved to a queue after it occurs. When creating a SparkContext object, Driver starts a thread to obtain an event from the queue in sequence and sends the event to each Listener. Each Listener processes the event after detecting the event.

Therefore, when the queuing speed is faster than the read speed, the queue overflows. As a result, the overflow event is lost, affecting the UI, EventLog, and dynamic resource scheduling functions. Therefore, a configuration item is added for more flexible use. You can set a proper value based on the memory size of the driver.

## Configuration Description

### Navigation path for setting parameters:

Before executing an application, modify the Spark service configuration. On Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration** and click **All Configurations**. Enter a parameter name in the search box.

**Table 26-45** Parameter description

| Parameter                                       | Description                                                                                        | Default Value |
|-------------------------------------------------|----------------------------------------------------------------------------------------------------|---------------|
| spark.scheduler.listenerbus.eventqueue.capacity | Specifies the size of the event queue. Configure this parameter based on the memory of the driver. | 100000<br>0   |

### NOTE

If the following information is displayed in the Driver log, the queue overflows.

1. Common application:  
Dropping SparkListenerEvent because no remaining room in event queue.  
This likely means one of the SparkListeners is too slow and cannot keep up with the rate at which tasks are being started by the scheduler.
2. Spark Streaming application:  
Dropping StreamingListenerEvent because no remaining room in event queue.  
This likely means one of the StreamingListeners is too slow and cannot keep up with the rate at which events are being started by the scheduler.

## 26.2.7.5 Configuring Executor Off-Heap Memory

### Scenario

When the executor off-heap memory is too small, or processes with higher priority preempt resources, the physical memory usage will exceed the maximal value. To prevent the physical memory usage from exceeding, set the following parameter.

### Configuration

#### Navigation path for setting parameters:

When submitting an application, set the following parameter using **--conf** or adjust the parameter in the **spark-defaults.conf** configuration file on the client.

**Table 26-46** Parameter description

| Parameter                     | Description                                                                                                                                                                                                                                                                      | Default Value |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.executor.memoryOverhead | Indicates the off-heap memory of each executor, in MB. Increasing the value of this parameter prevents the physical memory usage from exceeding the maximal value. The value is calculated based on $\max(384, \text{Executor - Memory} \times 0.1)$ . The minimal value is 384. | 1024          |

## 26.2.7.6 Enhancing Stability in a Limited Memory Condition

### Scenario

A large amount of memory is required when Spark SQL executes a query, especially during Aggregate and Join operations. If the memory is limited, OutOfMemoryError may occur. Stability in a limited memory condition ensures queries to be run in limited memory without OutOfMemoryError.

#### NOTE

Limited memory does not mean infinitely small memory, but ensures stable queries by using disks in a scenario where memory fails to store the data amount that is several times larger than the available memory size. For example, for queries involving Join, the data of the same key used for Join needs to be stored in memory. If the data amount is too large to be stored in the available memory, OutOfMemoryError occurs.

Stability in a limited memory condition involves the following sub-functions:

1. ExternalSort  
If the memory is inadequate during sorting, partial data overflows to disks.
2. TungstenAggregate  
By default, ExternalSort is used to sort data before data aggregation. Therefore, if the memory is inadequate, the data overflows to disks during sorting. The data has been properly sorted before aggregation and only aggregation results of the current key are remained, which use a small amount of memory.
3. SortMergeJoin and SortMergeOuterJoin  
SortMergeJoin and SortMergeOuterJoin are based on the equivalence join of sorted data. By default, ExternalSort is used to sort the data before the equivalence join. Therefore, if the memory is inadequate, the data overflows to disks during sorting. The data has been properly sorted before the equivalence join and only the data of the same key are remained, which uses a small amount of memory.

### Configuration

**Navigation path for setting parameters:**



When submitting an application, set the following parameters using `--conf` or adjust the parameters in the `spark-defaults.conf` configuration file on the client.

**Table 26-47** Parameter description

| Parameter                    | Scenario | Description                                                                                                                                                                                                                                                                                                                          | Default Value |
|------------------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.tungsten.enabled   | /        | Type: Boolean <ul style="list-style-type: none"> <li>If the value is <b>true</b>, tungsten is enabled. That is, the logic plan is equivalent to the codegeneration function, and the physical plan uses the corresponding tungsten execution plan.</li> <li>If the value is <b>false</b>, tungsten is disabled.</li> </ul>           | true          |
| spark.sql.codegen.wholeStage |          | Type: Boolean <ul style="list-style-type: none"> <li>If the value is <b>true</b>, codegeneration is enabled. That is, for some specified queries, the logic plan code will be generated dynamically when running.</li> <li>If the value is <b>false</b>, codegeneration is disabled and the existing static code is used.</li> </ul> | true          |

 **NOTE**

- To enable ExternalSort, you need to set `spark.sql.planner.externalSort` to **true** and `spark.sql.unsafe.enabled` to **false** or `spark.sql.codegen.wholeStage` to **false**.
- To enable TungstenAggregate, use either of the following methods:  
Set `spark.sql.codegen.wholeStage` and `spark.sql.unsafe.enabled` to **true** in the configuration file or CLI.  
If neither `spark.sql.codegen.wholeStage` nor `spark.sql.unsafe.enabled` is **true** or either of them is **true**, TungstenAggregate is enabled as long as `spark.sql.tungsten.enabled` is set to **true**.

### 26.2.7.7 Viewing Aggregated Container Logs on the Web UI

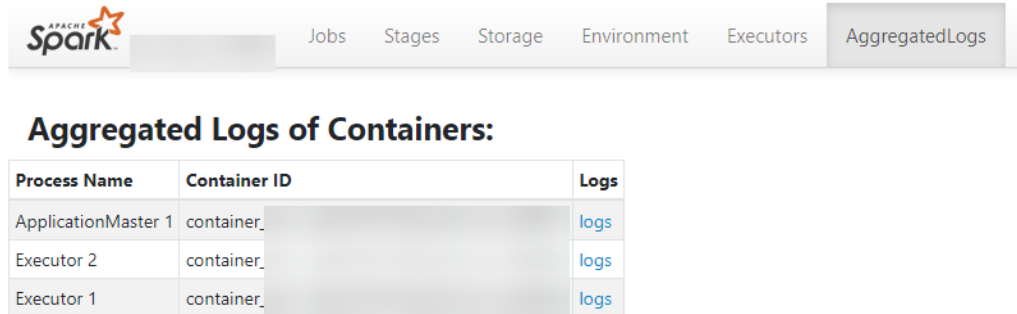
#### Scenarios

When `yarn.log-aggregation-enable` of Yarn is set to **true**, the container log aggregation function is enabled. Log aggregation indicates that after applications are run on Yarn, NodeManager aggregates all container logs of the node to HDFS and deletes local logs. For details, see [Configuring Container Log Aggregation](#).

However, all logs will be aggregated to an HDFS directory and can only be viewed by accessing an HDFS file. Open-source Spark and Yarn do not support the function of viewing aggregated logs on the web UI.

Spark supports this function. As shown in [Figure 26-4](#), the **AggregatedLogs** tab is added to the HistoryServer page. You can click **logs** to view aggregated logs.

**Figure 26-4** Aggregated log page



## Configuration Description

To display logs on the web UI, aggregated logs need to be parsed and presented. Spark parses aggregation logs using JobHistoryServer of Hadoop. Therefore, you can use the **spark.jobhistory.address** parameter to specify the URL of the JobHistoryServer page to parse and present the logs.

### Navigation path for setting parameters:

When submitting an application, set these parameters using **--conf** or adjust the following parameter in the **spark-defaults.conf** configuration file on the client.

#### NOTE

- This function depends on JobHistoryServer of Hadoop. Therefore, ensure that JobHistoryServer is running properly before using the log aggregation function.
- If the parameter value is empty, the **AggregatedLogs** tab page still exists, but you cannot view logs by clicking **logs**.
- The aggregated container logs can be viewed only when the application is running and event log files of the application exist on HDFS.
- You can click the log link on the **Executors** page to view the logs of a running task. After the task completes, the logs are aggregated to HDFS, and the log link on the **Executors** page becomes invalid. In this case, you can click **logs** on the **AggregatedLogs** page to view the aggregated logs.

**Table 26-48** Parameter description

| Parameter                | Description                                                                                                                                                                                                                                                                                                                          | Default Value |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.jobhistory.address | <p>URL of the JobHistoryServer page. The format is <i>http(s)://ip:port/jobhistory</i>. For example, <b>https://10.92.115.1:26014/jobhistory</b>.</p> <p>The default value is empty, indicating that container aggregation logs cannot be viewed on the web UI.</p> <p>Restart the service for the configuration to take effect.</p> | -             |

### 26.2.7.8 Configuring Environment Variables in Yarn-Client and Yarn-Cluster Modes

#### Scenario

Values of some configuration parameters of Spark client vary depending on its work mode (YARN-Client or YARN-Cluster). If you switch Spark client between different modes without first changing values of such configuration parameters, Spark client fails to submit jobs in the new mode.

To avoid this, configure parameters as described in [Table 26-49](#).

- In Yarn-Cluster mode, use the new parameters (path and parameters of Spark server).
- In Yarn-Client mode, uses the original parameters.  
They are `spark.driver.extraClassPath`, `spark.driver.extraJavaOptions`, and `spark.driver.extraLibraryPath`.

#### NOTE

If you choose not to add the parameters in [Table 26-49](#), Spark client can continue to operate well in either mode but the mode switch requires changes to some of its configuration parameters.

### Configuration Parameters

#### Navigation path for setting parameters:

On Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**. Click **All Configurations** and enter a parameter name in the search box.

**Table 26-49** Parameter description

| Parameter                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Default Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| spark.yarn.cluster.driver.extraClassPath   | <p>Indicates the extraClassPath of the driver in Yarn-cluster mode. Set the parameter to the path and parameters of the server.</p> <p>The original parameter <b>spark.driver.extraClassPath</b> indicates the extraClassPath of Spark client. By using different parameters to separate the settings of Spark server from the settings of Spark client, you can switch Spark client to different modes without changing parameter values.</p>                     | <p><code>\${BIGDATA_HOME}/common/runtime/security</code></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| spark.yarn.cluster.driver.extraJavaOptions | <p>Indicates the extraJavaOptions of Driver in Yarn-Cluster mode and is set to path and parameters of extraJavaOptions of Spark server.</p> <p>The original parameter <b>spark.driver.extraJavaOptions</b> indicates the path of extraJavaOptions of Spark client. By using different parameters to separate the settings of Spark server from the settings of Spark client, you can switch Spark client to different modes without changing parameter values.</p> | <p><code>-Xloggc:&lt;LOG_DIR&gt;/indexserver-%p-gc.log -XX:+PrintGCDetails -XX:-OmitStackTracelnFastThrow -XX:+PrintGCTimeStamps -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=20 -XX:GCLogFileSize=10M -Dlog4j.configuration=./__spark_conf__/__hadoop_conf__/log4j-executor.properties -Dlog4j.configuration.watch=true -Djava.security.auth.login.config=./__spark_conf__/__hadoop_conf__/jaas-zk.conf -Dzookeeper.server.principal=\${ZOOKEEPER_SERVER_PRINCIPAL} -Djava.security.krb5.conf=./__spark_conf__/__hadoop_conf__/kdc.conf -Djetty.version=x.y.z -Dorg.xerial.snappy.tmpdir=\${BIGDATA_HOME}/tmp -Dcarbon.properties.filepath=./__spark_conf__/__hadoop_conf__/carbon.properties -Djdk.tls.ephemeralDHKeySize=2048 -Dspark.ssl.keyStore=child.keystore #{java_stack_prefer}</code></p> |

### 26.2.7.9 Configuring the Default Number of Data Blocks Divided by SparkSQL

#### Scenarios

By default, SparkSQL divides data into 200 data blocks during shuffle. In data-intensive scenarios, each data block may have excessive size. If a single data block of a task is larger than 2 GB, an error similar to the following will be reported while Spark attempts to fetch the data block:

```
Adjusted frame length exceeds 2147483647: 2717729270 - discarded
```

For example, setting the number of default data blocks to 200 causes SparkSQL to encounter an error in running a TPCDS 500-GB test. To avoid this, increase the number of default blocks in data-intensive scenarios.

#### Configuration parameters

##### Navigation path for setting parameters:

On Manager, choose **Cluster > Name of the desired cluster > Service > Spark2x > Configuration** and click **All Configurations**. Enter a parameter name in the search box.

**Table 26-50** Parameter description

| Parameter                    | Description                                                    | Default Value |
|------------------------------|----------------------------------------------------------------|---------------|
| spark.sql.shuffle.partitions | Indicates the default number of blocks divided during shuffle. | 200           |

### 26.2.7.10 Configuring the Compression Format of a Parquet Table

#### Scenarios

The compression format of a Parquet table can be configured as follows:

1. If the Parquet table is a partitioned one, set the **parquet.compression** parameter of the Parquet table to specify the compression format. For example, set **tblproperties** in the table creation statement:  
**"parquet.compression"="snappy"**.
2. If the Parquet table is a non-partitioned one, set the **spark.sql.parquet.compression.codec** parameter to specify the compression format. The configuration of the **parquet.compression** parameter is invalid, because the value of the **spark.sql.parquet.compression.codec** parameter is read by the **parquet.compression** parameter. If the **spark.sql.parquet.compression.codec** parameter is not configured, the default value is **snappy** and will be read by the **parquet.compression** parameter.

Therefore, the `spark.sql.parquet.compression.codec` parameter can only be used to set the compression format of a non-partitioned Parquet table.

## Configuration parameters

### Navigation path for setting parameters:

On Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration**. Click **All Configurations** and enter a parameter name in the search box.

**Table 26-51** Parameter description

| Parameter                           | Description                                                            | Default Value |
|-------------------------------------|------------------------------------------------------------------------|---------------|
| spark.sql.parquet.compression.codec | Used to set the compression format of a non-partitioned Parquet table. | snappy        |

### 26.2.7.11 Configuring the Number of Lost Executors Displayed in WebUI

#### Scenario

In Spark WebUI, the **Executor** page can display information about Lost Executor. Executors are dynamically recycled. If the JDBCServer tasks are large, there may be too many lost executors displayed in WebUI. Therefore, the number of displayed lost executors can be configured.

#### Procedure

Configure the following parameter in the `spark-defaults.conf` file on Spark client.

**Table 26-52** Parameter description

| Parameter                      | Description                                                    | Default Value |
|--------------------------------|----------------------------------------------------------------|---------------|
| spark.ui.retainedDeadExecutors | The maximum number of Lost Executors displayed in Spark WebUI. | 100           |

### 26.2.7.12 Setting the Log Level Dynamically

#### Scenarios

In some scenarios, to locate problems or check information by changing the log level,

you can add the `-Dlog4j.configuration.watch=true` parameter to the JVM parameter of a process before the process is started. After the process is started,

you can modify the log4j configuration file corresponding to the process to change the log level.

The following processes support the dynamic setting of log levels: driver, executor, ApplicationMaster, JobHistory and JDBCServer.

Allowed log levels are as follows: FATAL, ERROR, WARN, INFO, DEBUG, TRACE, and ALL.

## Configuration Description

Add the following parameters to the JVM parameter corresponding to a process.

**Table 26-53** Parameter description

| Parameter                      | Description                                                                                                                           | Default Value                                                                   |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| - Dlog4j.configuration.watcher | Indicates a JVM parameter of a process. If this parameter is set to <b>true</b> , the dynamic configuration of log levels is enabled. | Left blank, indicating that the dynamic configuration of log levels is disabled |

**Table 26-54** lists the JVM parameters of the driver, executor, and ApplicationMaster processes. Configure the following parameters in the **spark-defaults.conf** file on the Spark client. Set the log levels of the driver, executor, and ApplicationMaster processes in the log4j configuration file specified by the - **Dlog4j.configuration** parameter.

**Table 26-54** JVM parameters of processes (1)

| Parameter                       | Description                                                   | Default Log Level |
|---------------------------------|---------------------------------------------------------------|-------------------|
| spark.driver.extraJavaOptions   | Indicates the JVM parameter of the driver process.            | INFO              |
| spark.executor.extraJavaOptions | Indicates the JVM parameter of the executor process.          | INFO              |
| spark.yarn.am.extraJavaOptions  | Indicates the JVM parameter of the ApplicationMaster process. | INFO              |

**Table 26-55** describes the JVM parameters of JobHistory Server and JDBCServer. Set the parameters in the **ENV\_VARS** configuration file. Set the log levels of JobHistory Server and JDBCServer in the **log4j.properties** configuration file.

**Table 26-55** JVM parameters of processes (2)

| Parameter         | Description                                                   | Default Log Level |
|-------------------|---------------------------------------------------------------|-------------------|
| GC_OPTS           | Indicates the JVM parameter of the JobHistory Server process. | INFO              |
| SPARK_SUBMIT_OPTS | Indicates the JVM parameter of JDBCServer.                    | INFO              |

**Example:**

To change the log level of the executor process to DEBUG dynamically, modify the **spark.executor.extraJavaOptions** JVM parameter of the executor process in the **spark-defaults.conf** file and run the following command to add the following configuration before the process is started:

```
-Dlog4j.configuration.watch=true
```

After the user application is submitted, change the log level in the Log4j configuration file (for example, **-Dlog4j.configuration=file:\${BIGDATA\_HOME}/FusionInsight\_Spark2x\_8.1.0.1/install/FusionInsight-Spark2x-\*/spark/conf/log4j-executor.properties**) specified by the **-Dlog4j.configuration** parameter in **spark.executor.extraJavaOptions** to **DEBUG**:

```
log4j.rootCategory=DEBUG, sparklog
```

It takes several seconds for the DEBUG level to take effect.

### 26.2.7.13 Configuring Whether Spark Obtains HBase Tokens

#### Scenario

When Spark is used to submit tasks, the driver obtains tokens from HBase by default. To access HBase, you need to configure the **jaas.conf** file for security authentication. If the **jaas.conf** file is not configured, the application will fail to run.

Therefore, perform the following operations based on whether the application involves HBase:

- If the application does not involve HBase, you do not need to obtain the HBase tokens. In this case, set **spark.yarn.security.credentials.hbase.enabled** to **false**.
- If the application involves HBase, set **spark.yarn.security.credentials.hbase.enabled** to **true** and configure the **jaas.conf** file on the driver as follows:

```
{client}/spark/bin/spark-sql --master yarn-client --principal {principal} --keytab {keytab} --driver-java-options "-Djava.security.auth.login.config={LocalPath}/jaas.conf"
```

Specify Keytab and Principal in the **jaas.conf** file. The following is an example:

```
Client {
 com.sun.security.auth.module.Krb5LoginModule required
 useKeyTab=true
 keyTab = "{LocalPath}/user.keytab"
```



```
principal="super@<System domain name>"
useTicketCache=false
debug=false;
};
```

## Configuration

Configure the following parameter in the **spark-defaults.conf** file of the Spark client.

**Table 26-56** Parameter description

| Parameter                                     | Description                                                                                                                                                                                     | Default Value |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.yarn.security.credentials.hbase.enabled | <p>Indicates whether HBase obtains a token.</p> <ul style="list-style-type: none"> <li>• <b>true</b>: HBase obtains a token.</li> <li>• <b>false</b>: HBase does not obtain a token.</li> </ul> | false         |


### 26.2.7.14 Configuring LIFO for Kafka

#### Scenario

If the Spark Streaming application is connected to Kafka, after the Spark Streaming application is terminated abnormally and restarted from the checkpoint, the system preferentially processes the tasks that are not completed before the application is terminated (Period A) and the tasks generated based on data that enters Kafka during the period (Period B) from the application termination to the restart. Then the application processes the tasks generated based on data that enters Kafka after the application is restarted (Period C). For data that enters Kafka in period B, Spark generates a corresponding number of tasks based on the end time (**batch** time). The first task reads all data, but other tasks may not read data. As a result, the task processing pressure is uneven.

If the tasks in Period A and Period B are processed slowly, the processing of tasks in period C is affected. To cope with the preceding scenario, Spark provides the last-in first-out (LIFO) function for Kafka.

**Figure 26-5** Time axis for restarting the Spark Streaming application



After this function is enabled, Spark preferentially schedules tasks in Period C. If there are multiple tasks in Period C, Spark schedules and executes the tasks in the sequence of task generation. Then Spark executes the tasks in Periods A and B. For data that enters Kafka in Period B, Spark generates tasks based on the end time and evenly distributes all data that enters Kafka in this period to each task to avoid uneven task processing pressure.

Constraints:

- This function applies only to the direct mode of Spark Streaming, and the execution result does not depend on the processing result of the previous batch (that is, stateless operation, for example, **updatestatebykey**). Multiple data input streams must be comparatively independent from each other. Otherwise, the result may change after the data is divided.
- The Kafka LIFO function can be enabled only when the application is connected to the Kafka input source.
- If both Kafka LIFO and flow control functions are enabled when the application is submitted, the flow control function is not enabled for the data that enters Kafka in Period B to ensure that the task scheduling priority for reading the data is the lowest. Flow control is enabled for the tasks in Period C after the application is restarted.

## Configuration

Configure the following parameters in the **spark-defaults.conf** file on the Spark driver.

**Table 26-57** Parameter description

| Parameter                         | Description                                             | Default Value |
|-----------------------------------|---------------------------------------------------------|---------------|
| spark.streaming.kafka.direct.lifo | Specifies whether to enable the LIFO function of Kafka. | false         |

| Parameter                                  | Description                                   | Default Value                                               |
|--------------------------------------------|-----------------------------------------------|-------------------------------------------------------------|
| spark.streaming.kafka010.inputstream.class | Obtains the decoupled class on FusionInsight. | org.apache.spark.streaming.kafka010.DirectKafkaInputDStream |

## 26.2.7.15 Configuring Reliability for Connected Kafka

### Scenario

When the Spark Streaming application is connected to Kafka and the application is restarted, the application reads data from Kafka based on the last read topic offset and the latest offset of the current topic.

If the leader of a Kafka topic fails and the offset of the Kafka leader is greatly different from that of the Kafka follower, the Kafka follower and leader are switched over after the Kafka service is restarted. As a result, the offset of the topic decreases after the Kafka service is restarted.

- If the Spark Streaming application keeps running, the start position for reading Kafka data is greater than the end position because the offset of the topic in Kafka decreases. As a result, the application cannot read data from Kafka and reports an error.
- Before restarting the Kafka service, stop the Spark Streaming application. After the Kafka service is restarted, restart the Spark Streaming application to restore the application from the checkpoint. In this case, the Spark Streaming application records the offset position read before the termination and uses the position as the reference to read subsequent data. The Kafka offset decreases (for example, from 100,000 to 10,000). Spark Streaming consumes data only after the offset of the Kafka leader increases to 100,000. As a result, the newly sent data whose offset is between 10,000 and 100,000 is lost.

To resolve the preceding problem, you can configure reliability for Kafka connected to Spark Streaming. After the reliability function of connected Kafka is enabled:

- If the offset of a topic in Kafka decreases when the Spark Streaming application is running, the latest offset of the topic in Kafka is used as the start position for reading Kafka data and subsequent data is read.

For a task that has been generated but has not been scheduled, if the read Kafka offset is greater than the latest offset of the topic in Kafka, the task fails to be executed.

#### NOTE

If a large number of tasks fail, the Executor is added to the blacklist. As a result, subsequent tasks cannot be deployed and run. If this happens, you can set **spark.blacklist.enabled** to disable the blacklist function. The blacklist function is enabled by default.

- If the offset of a topic in Kafka decreases, the Spark Streaming application restarts to restore the unfinished tasks. If the read Kafka offset range is greater than the latest offset of the topic in Kafka, the task is directly discarded.

 **NOTE**

If the state function is used in the Spark Streaming application, do not enable the reliability function of connected Kafka.

## Configuration

Configure the following parameter in the **spark-defaults.conf** file of the Spark client.

**Table 26-58** Parameter description

| Parameter                         | Description                                                                                                                                                                                                                                                    | Default Value |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.streaming.Kafka.reliability | Indicates whether to enable the reliability function for Kafka connected to Spark Streaming. <ul style="list-style-type: none"><li>• <b>true</b>: The reliability function is enabled.</li><li>• <b>false</b>: The reliability function is disabled.</li></ul> | false         |

### 26.2.7.16 Configuring Streaming Reading of Driver Execution Results

#### Scenario

When a query statement is executed, the returned result may be large (containing more than 100,000 records). In this case, JDBCServer out of memory (OOM) may occur. Therefore, the data aggregation function is provided to avoid OOM without sacrificing the performance.

#### Configuration

Two data aggregation function configuration parameters are provided. The two parameters are set in the **tunning** option on the Spark JDBCServer server. After the setting is complete, restart JDBCServer.

**Table 26-59** Parameter description

| Parameter                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Default Value |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.bigdata.thriftServer.useHdfsCollect | <p>Indicates whether to save result data to HDFS instead of the memory.</p> <p>Advantages: The query result is stored in HDFS. Therefore, JDBCServer OOM does not occur.</p> <p>Disadvantages: The query is slow.</p> <ul style="list-style-type: none"> <li>● <b>true</b>: Result data is saved to HDFS.</li> <li>● <b>false</b>: This function is disabled.</li> </ul> <p><b>NOTICE</b><br/>When <b>spark.sql.bigdata.thriftServer.useHdfsCollect</b> is set to <b>true</b>, result data is saved to HDFS. However, the job description on the native JobHistory page cannot be associated with the corresponding SQL statement. In addition, the execution ID in the spark-beeline command output is null. To solve the JDBCServer OOM problem and ensure correct information display, you are advised to set <b>spark.sql.userlocalFileCollect</b>.</p> | false         |
| spark.sql.userlocalFileCollect                | <p>Indicates whether to save result data to the local disk instead of memory.</p> <p>Advantages: In the case of small data volume, the performance loss can be ignored compared with the data storage mode using the native memory. In the case of large data volume (hundreds of millions of data records), the performance is much better than that when data is stored in the HDFS and native memory.</p> <p>Disadvantages: Optimization is required. In the case of large data volume, it is recommended that the JDBCServer driver memory be 10 GB and each core of the executor be allocated with 3 GB memory.</p> <ul style="list-style-type: none"> <li>● <b>true</b>: This function is enabled.</li> <li>● <b>false</b>: This function is disabled.</li> </ul>                                                                                     | false         |

| Parameter                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Default Value |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.collect.Hive      | <p>This parameter is valid only when <b>spark.sql.useLocalFileCollect</b> is set to <b>true</b>. It indicates whether to save the result data to a disk in direct serialization mode or in indirect serialization mode.</p> <p>Advantage: For queries of tables with a large number of partitions, the aggregation performance of the query results is better than that of the storage mode that query results are directly stored on the disk.</p> <p>Disadvantages: The disadvantages are the same as those when <b>spark.sql.useLocalFileCollect</b> is enabled.</p> <ul style="list-style-type: none"> <li>• <b>true</b>: This function is enabled.</li> <li>• <b>false</b>: This function is disabled.</li> </ul> | false         |
| spark.sql.collect.serialize | <p>This parameter takes effect only when both <b>spark.sql.useLocalFileCollect</b> and <b>spark.sql.collect.Hive</b> are set to <b>true</b>.</p> <p>The function is to further improve performance.</p> <ul style="list-style-type: none"> <li>• <b>java</b>: Data is collected in Java serialization mode.</li> <li>• <b>kryo</b>: Data is collected in kryo serialization mode. The performance is better than that when the Java serialization mode is used.</li> </ul>                                                                                                                                                                                                                                             | java          |

 NOTE

**spark.sql.bigdata.thriftServer.useHdfsCollect** and **spark.sql.useLocalFileCollect** cannot be set to **true** at the same time.

### 26.2.7.17 Filtering Partitions without Paths in Partitioned Tables

#### Scenario

When you perform the *select* query in Hive partitioned tables, the **FileNotFoundException** exception is displayed if a specified partition path does not exist in HDFS. To avoid the preceding exception, configure **spark.sql.hive.verifyPartitionPath** parameter to filter partitions without paths.

#### Procedure

Perform either of the following methods to filter partitions without paths:

- Configure the following parameter in the **spark-defaults.conf** file on Spark client.

**Table 26-60** Parameter description

| Parameter                          | Description                                                                                                                                                      | Default Value |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.hive.verifyPartitionPath | Whether to filter partitions without paths when reading Hive partitioned tables.<br><b>true</b> : enables the filtering<br><b>false</b> : disables the filtering | false         |

- When running the spark-submit command to submit an application, configure the **--conf** parameter to filter partitions without paths.

For example:

```
spark-submit --class org.apache.spark.examples.SparkPi --conf spark.sql.hive.verifyPartitionPath=true $SPARK_HOME/lib/spark-examples_*.jar
```

## 26.2.7.18 Configuring Spark2x Web UI ACLs

### Scenario

Users need to implement security protection for Spark2x web UI when some data on the UI cannot be viewed by other users. Once a user attempts to log in to the UI, Spark2x can check the view ACL of the user to determine whether to allow the user to access the UI.

Spark2x has two types of web UI. One is for running tasks. You can access the web UI using the application link on the native Yarn page or the REST APIs. The other one is for ended tasks. You can access the web UI using the Spark2x JobHistory service or the REST APIs.

#### NOTE

This section applies only to clusters in security mode (with Kerberos authentication enabled).

- Configuring the ACL of the web UI for running tasks  
For a running task, you can set the following parameters on the server:
  - spark.admin.acls**: specifies the web UI administrator list.
  - spark.admin.acls.groups**: specifies the administrator group list.
  - spark.ui.view.acls**: specifies the Yarn page visitor list.
  - spark.modify.acls.groups**: specifies the Yarn page visitor group list.
  - spark.modify.acls**: specifies the web UI modifier list.
  - spark.ui.view.acls.groups**: specifies the web UI modifier group list.
- Configuring the ACL of the web UI for ended tasks  
For ended tasks, use client parameter **spark.history.ui.acls.enable** to enable or disable the ACL access permission.  
If ACL control is enabled, configure client parameters **spark.admin.acls** and **spark.admin.acls.groups** to specify the web UI administrator list and administrator group list. Use client parameters **spark.ui.view.acls** and

**spark.modify.acls.groups** to specify the visitor list and visitor group list that view web UI task details. Use client parameters **spark.modify.acls** and **spark.ui.view.acls.groups** to specify the visitor list and group list that modify web UI task details.

## Configuration

Log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Spark2x > Configurations**, click **All Configurations**, search for **acl**, and modify the following parameters on the JobHistory, JDBCServer, SparkResource, and Spark pages.

**Table 26-61** Parameter description

| Parameter                    | Description                                                                                                                                                                                                                                      | Default Value |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.history.ui.acls.enable | Indicates whether JobHistory supports the permission verification of a single task.                                                                                                                                                              | true          |
| spark.acls.enable            | Indicates whether to enable Spark permission management.<br>If this function is enabled, the system checks whether the user has the permission to access and modify task information.                                                            | true          |
| spark.admin.acls             | Indicates the list of Spark administrators. All members in the list have the rights to manage all Spark tasks. You can configure multiple administrators and separate them from each other using commas (,).                                     | admin         |
| spark.admin.acls.groups      | Indicates the list of Spark administrator groups. All groups in the list have the permission to manage all Spark tasks. You can configure multiple administrator groups and separate them from each other using commas (,).                      | -             |
| spark.modify.acls            | Indicates the list of members that have the permission to modify Spark tasks. By default, the user who starts a task has the permission to modify the task. You can configure multiple users and separate them from each other using commas (,). | -             |



| Parameter                 | Description                                                                                                                                                                                                                                      | Default Value |
|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.modify.acls.groups  | Indicates the list of groups that have the permission to modify Spark tasks. You can configure multiple groups and separate them from each other using commas (,).                                                                               | -             |
| spark.ui.view.acls        | Indicates the list of members that have the permission to access Spark tasks. By default, the user who starts a task has the permission to modify the task. You can configure multiple users and separate them from each other using commas (,). | -             |
| spark.ui.view.acls.groups | Indicates the list of groups that have the permission to access Spark tasks. You can configure multiple groups and separate them from each other using commas (,).                                                                               | -             |

 NOTE

If you use a client to submit tasks, you must download the client again after modifying the `spark.admin.acls`, `spark.admin.acls.groups`, `spark.modify.acls`, `spark.modify.acls.groups`, `spark.ui.view.acls`, and `spark.ui.view.acls.groups` parameters.

## 26.2.7.19 Configuring Vector-based ORC Data Reading

### Scenario

ORC is a column-based storage format in the Hadoop ecosystem. It originates from Apache Hive and is used to reduce the Hadoop data storage space and accelerate the Hive query speed. Similar to Parquet, ORC is not a pure column-based storage format. In the ORC format, the entire table is split based on the row group, data in each row group is stored by column, and data is compressed as much as possible to reduce storage space consumption. Vector-based ORC data reading significantly improves the ORC data reading performance. In Spark2.3, SparkSQL supports vector-based ORC data reading (this function is supported in earlier Hive versions). Vector-based ORC data reading improves the data reading performance by multiple times.

This feature can be enabled by using the following parameter.

- **spark.sql.orc.enableVectorizedReader**: specifies whether vector-based ORC data reading is supported. The default value is **true**.
- **spark.sql.codegen.wholeStage**: specifies whether to compile all stages of multiple operations into a Java method. The default value is **true**.
- **spark.sql.codegen.maxFields**: specifies the maximum number of fields (including nested fields) supported by all stages of codegen. The default value is **100**.

- **spark.sql.orc.impl**: specifies whether Hive or Spark SQL native is used as the SQL execution engine to read ORC data. The default value is **hive**.

## Parameters

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x**, click the **Configurations** tab and then **All Configurations**, and search for the following parameters.

| Parameter                            | Description                                                                                                                            | Default Value | Value Range    |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|---------------|----------------|
| spark.sql.orc.enableVectorizedReader | Specifies whether vector-based ORC data reading is supported. The default value is <b>true</b> .                                       | true          | [true,false]   |
| spark.sql.codegen.wholeStage         | Specifies whether to compile all stages of multiple operations into a Java method. The default value is <b>true</b> .                  | true          | [true,false]   |
| spark.sql.codegen.maxFields          | Specifies the maximum number of fields (including nested fields) supported by all stages of codegen. The default value is <b>100</b> . | 100           | Greater than 0 |
| spark.sql.orc.impl                   | Specifies whether Hive or Spark SQL native is used as the SQL execution engine to read ORC data. The default value is <b>hive</b> .    | hive          | [hive,native]  |

### NOTE

1. To use vector-based ORC data reading of SparkSQL, the following conditions must be met:
  - **spark.sql.orc.enableVectorizedReader** must be set to **true** (default value). Generally, the value is not changed.
  - **spark.sql.codegen.wholeStage** must be set to **true** (default value). Generally, the value is not changed.
  - The value of **spark.sql.codegen.maxFields** must be greater than or equal to the number of columns in scheme.
  - All data is of the AtomicType. Specifically, data is not null or of the UDT, array, or map type. If there is data of the preceding types, expected performance cannot be obtained.
  - **spark.sql.orc.impl** must be set to **native**. The default value is **hive**.
2. If a task is submitted using the client, modification of the following parameters takes effect only after you download the client again: **spark.sql.orc.enableVectorizedReader**, **spark.sql.codegen.wholeStage**, **spark.sql.codegen.maxFields**, and **spark.sql.orc.impl**.

## 26.2.7.20 Broaden Support for Hive Partition Pruning Predicate Pushdown

### Scenario

In earlier versions, the predicate for pruning Hive table partitions is pushed down. Only comparison expressions between column names and integers or character strings can be pushed down. In version 2.3, pushdown of the null, in, and, or expressions are supported.

### Parameters

Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x**. On the page that is displayed, click the **Configurations** tab then the **All Configurations** sub-tab, and search for the following parameters:

| Parameter                                                 | Description                                                                             | Default Value | Value Range  |
|-----------------------------------------------------------|-----------------------------------------------------------------------------------------|---------------|--------------|
| spark.sql.hive.advancedPartitionPredicatePushdown.enabled | Specifies whether to broaden the support for Hive partition pruning predicate pushdown. | true          | [true,false] |

## 26.2.7.21 Hive Dynamic Partition Overwriting Syntax

### Scenario

In earlier versions, when the **insert overwrite** syntax is used to overwrite partition tables, only partitions with specified expressions are matched, and partitions without specified expressions are deleted. In Spark2.3, partitions without specified expressions are automatically matched. The syntax is the same as that of the dynamic partition matching syntax of Hive.

### Parameters

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, click **All Configurations**, and search for the following parameter.

| Parameter                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                         | Default Value | Value Range      |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|------------------|
| spark.sql.sources.partitionOverwrite-Mode | Specifies the mode for inserting data in partition tables by running the <b>insert overwrite</b> command, which can be <b>STATIC</b> or <b>DYNAMIC</b> . When it is set to <b>STATIC</b> , Spark deletes all partitions based on the matching conditions. When it is set to <b>DYNAMIC</b> , Spark matches partitions based on matching conditions and dynamically matches partitions without specified conditions. | STATIC        | [STATIC,DYNAMIC] |

## 26.2.7.22 Configuring the Column Statistics Histogram to Enhance the CBO Accuracy

### Scenarios

The execution plan for SQL statements is optimized in Spark. Common optimization rules are heuristic optimization rules. Heuristic optimization rules are provided based on the characteristics of logical plans, and the data characteristics and the execution costs of operators are not considered. Spark 2.20 introduces the Cost-Based Optimization (CBO). CBO collects statistics on tables and columns and estimates the number of output records and size of each operator in bytes based on the input data sets of operators, which is the cost of executing an operator.

CBO will adjust the execution plan to minimize the end-to-end query time. The main points are as follows:

- Filter out irrelevant data as soon as possible.
- Minimize the cost of each operator.

The CBO optimization process is divided into two steps:

1. Collect statistics.
2. Estimate the output data sets of a specific operator based on the input data sets.

Table-level statistics includes: number of records and the total size of a table data file.

Column-level statistics includes: number of unique values, maximum value, minimum value, number of null values, average length, maximum length, and the histogram.

After the statistics is obtained, the execution cost of operators can be estimated. Common operators include filter and join operators.

Histogram is a type of column statistics. It can clearly describe the distribution of column data. The column data is distributed to a specified number of bins that are displayed in ascending order by size. The upper and lower limits of each bin are calculated. The amount of data in all bins is the same (a contour histogram). After the data is distributed, the cost estimation of each operator is more accurate and the optimization effect is better.

This feature can be enabled by using the following parameter.

**spark.sql.statistics.histogram.enabled:** specifies whether to enable the histogram function. The default value is **false**.

## Parameter Configuration

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, click **All Configurations**, and search for the following parameters.

| Parameter                                    | Description                                                                                                                                                                                                                                | Default Value | Value Range  |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|--------------|
| spark.sql.cbo.enabled                        | Enables CBO to estimate the statistics for the execution plan.                                                                                                                                                                             | false         | [true,false] |
| spark.sql.cbo.joinReorder.enabled            | Enables CBO join for reordering.                                                                                                                                                                                                           | false         | [true,false] |
| spark.sql.cbo.joinReorder.dp.threshold       | Specifies the maximum number of nodes that can be joined in the dynamic planning algorithm.                                                                                                                                                | 12            | >=1          |
| spark.sql.cbo.joinReorder.card.weight        | Specifies the proportion of dimension (number of rows) in the comparison of planned cost during reconnection: Number of rows x Proportion of dimension + File size x (1 - Proportion of dimension)                                         | 0.7           | 0-1          |
| spark.sql.statistics.size.autoUpdate.enabled | Enables the function of automatically updating the table size when the table data volume changes.<br>Note: If there are a large number of data files in a table, this operation is time consume, and the data processing speed is reduced. | false         | [true,false] |

| Parameter                                | Description                                                                                                                                                                                                                                                          | Default Value | Value Range  |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|--------------|
| spark.sql.statistics.histogram.enabled   | After this function is enabled, a histogram is generated when column statistics is collected. Histogram can improve the estimation accuracy, but collecting histogram information requires additional workload.                                                      | false         | [true,false] |
| spark.sql.statistics.histogram.numBins   | Specifies the number of bins for the generated histogram.                                                                                                                                                                                                            | 254           | >=2          |
| spark.sql.statistics.ndv.maxError        | Specifies the maximum estimation deviation allowed by the HyperLogLog++ algorithm when the column level statistics is generated.                                                                                                                                     | 0.05          | 0-1          |
| spark.sql.statistics.percentile.accuracy | Specifies the accuracy rate of the percentile estimation when the contour histogram is generated. The larger the value is, the more accurate the estimation is. The estimation error value can be obtained through (1.0/Accuracy rate of the percentile estimation). | 10000         | >=1          |

 NOTE

- If you want the histogram to take effect in CBO, the following conditions must be met:
  - Set **spark.sql.statistics.histogram.enabled** to **true**. The default value is **false**. Change the value to **true** to enable the histogram function.
  - Set **spark.sql.cbo.enabled** to **true**. The default value is **false**. Change the value to **true** to enable CBO.
  - Set **spark.sql.cbo.joinReorder.enabled** to **true**. The default value is **false**. Change the value to **true** to enable connection reordering.
- If a client is used to submit a task, you need to download the client again after configuring the following parameters: **spark.sql.cbo.enabled**, **spark.sql.cbo.joinReorder.enabled**, **spark.sql.cbo.joinReorder.dp.threshold**, **spark.sql.cbo.joinReorder.card.weight**, **spark.sql.statistics.size.autoUpdate.enabled**, **spark.sql.statistics.histogram.enabled**, **spark.sql.statistics.histogram.numBins**, **spark.sql.statistics.ndv.maxError**, and **spark.sql.statistics.percentile.accuracy**.

### 26.2.7.23 Configuring Local Disk Cache for JobHistory

#### Scenarios

JobHistory can use local disks to cache the historical data of Spark applications to prevent the JobHistory memory from loading a large amount of application data, reducing the memory pressure. In addition, the cached data can be reused to improve the speed for subsequent application access.

#### Parameter Configuration

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, click the **All Configurations** tab, and search for the following parameters:

| Parameter                        | Description                                                                                                                                                                                             | Default Value                           |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| spark.history.store.path         | Specifies the local directory for storing historical information for JobHistory. If this parameter is specified, JobHistory caches historical application data in the local disk instead of the memory. | \${BIGDATA_HOME}/tmp/spark2x_JobHistory |
| spark.history.store.maxDiskUsage | Specifies the maximum available space of the local disk cache.                                                                                                                                          | 10 GB                                   |

### 26.2.7.24 Configuring Spark SQL to Enable the Adaptive Execution Feature

#### Scenario

The Spark SQL adaptive execution feature enables Spark SQL to optimize subsequent execution processes based on intermediate results to improve overall execution efficiency. The following features have been implemented:

1. Automatic configuration of the number of shuffle partitions  
Before the adaptive execution feature is enabled, Spark SQL specifies the number of partitions for a shuffle process by specifying the **spark.sql.shuffle.partitions** parameter. This method lacks flexibility when multiple SQL queries are performed on an application and cannot ensure optimal performance in all scenarios. After adaptive execution is enabled, Spark SQL automatically configures the number of partitions for each shuffle process, instead of using the general configuration. In this way, the proper number of partitions is automatically used during each shuffle process.
2. Dynamic adjusting of the join execution plan  
Before the adaptive execution feature is enabled, Spark SQL creates an execution plan based on the optimization results of rule-based optimization (RBO) and Cost-Based Optimization (CBO). This method ignores changes of result sets during data execution. For example, when a view created based on a large table is joined with other large tables, the execution plan cannot be

adjusted to BroadcastJoin even if the result set of the view is small. After the adaptive execution feature is enabled, Spark SQL can dynamically adjust the execution plan based on the execution result of the previous stage to obtain better performance.

3. Automatic processing of data skew

If data skew occurs during SQL statement execution, the memory overflow of an executor or slow task execution may occur. After the adaptive execution feature is enabled, Spark SQL can automatically process data skew scenarios. Multiple tasks are started for partitions where data skew occurs. Each task reads several output files obtained from the shuffle process and performs union operations on the join results of these tasks to eliminate data skew.

## Parameters

Log in to FusionInsight Manager, choose **Cluster > Services > Spark2x > Configurations**, click **All Configurations**, and search for the following parameter.

| Parameter                                                 | Description                                                                                                                                                                                                                                                                                                                                                                          | Default Value |
|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.adaptive.enabled                                | Specifies whether to enable the adaptive execution function.<br><br>Note: If AQE and Static Partition Pruning (DPP) are enabled at the same time, DPP takes precedence over AQE during SparkSQL task execution. As a result, AQE does not take effect.                                                                                                                               | false         |
| spark.sql.optimizer.dynamicPartitionPruning.enabled       | The switch to enable DPP.                                                                                                                                                                                                                                                                                                                                                            | true          |
| spark.sql.adaptive.coalescePartitions.enabled             | If this parameter is set to <b>true</b> and <b>spark.sql.adaptive.enabled</b> is set to <b>true</b> , Spark combines partitions that are consecutively random played based on the target size (specified by <b>spark.sql.adaptive.advisoryPartitionSizeInBytes</b> ) to prevent too many small tasks from being executed.                                                            | true          |
| spark.sql.adaptive.coalescePartitions.initialPartitionNum | Initial number of shuffle partitions before merge. The default value is the same as the value of <b>spark.sql.shuffle.partitions</b> . This parameter is valid only when <b>spark.sql.adaptive.enabled</b> and <b>spark.sql.adaptive.coalescePartitions.enabled</b> are set to <b>true</b> . This parameter is optional. The initial number of partitions must be a positive number. | 200           |



| Parameter                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Default Value |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.adaptive.coalescePartitions.minPartitionNum | Minimum number of shuffle partitions after merge. If this parameter is not set, the default degree of parallelism (DOP) of the Spark cluster is used. This parameter is valid only when <b>spark.sql.adaptive.enabled</b> and <b>spark.sql.adaptive.coalescePartitions.enabled</b> are set to <b>true</b> . This parameter is optional. The initial number of partitions must be a positive number.                                                                                                                                                                     | 1             |
| spark.sql.adaptive.shuffle.targetPostShuffleInputSize | Target size of a partition after shuffling. Spark 3.0 and later versions do not support this parameter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 64MB          |
| spark.sql.adaptive.advisoryPartitionSizeInBytes       | Size of a shuffle partition (unit: byte) during adaptive optimization ( <b>spark.sql.adaptive.enabled</b> is set to <b>true</b> ). This parameter takes effect when Spark aggregates small shuffle partitions or splits shuffle partitions where skew occurs.                                                                                                                                                                                                                                                                                                           | 64MB          |
| spark.sql.adaptive.fetchShuffleBlocksInBatch          | Whether to obtain consecutive shuffle blocks in batches. For the same map job, reading consecutive shuffle blocks in batches can reduce I/Os and improve performance, instead of reading blocks one by one. Note that multiple consecutive blocks exist in a single read request only when <b>spark.sql.adaptive.enabled</b> and <b>spark.sql.adaptive.coalescePartitions.enabled</b> are set to <b>true</b> . This feature also relies on a relocatable serializer that uses cascading to support the codec and the latest version of the shuffle extraction protocol. | true          |
| spark.sql.adaptive.localShuffleReader.enabled         | If the value of this parameter is <b>true</b> and the value of <b>spark.sql.adaptive.enabled</b> is <b>true</b> , Spark attempts to use the local shuffle reader to read shuffle data when shuffling of partitions is not required, for example, after sort-merge join is converted to broadcast-hash join.                                                                                                                                                                                                                                                             | true          |

| Parameter                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                       | Default Value |
|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.sql.adaptive.skewJoin.enabled                         | Specifies whether to enable the function of automatic processing of the data skew in join operations. The function is enabled when this parameter is set to <b>true</b> and <b>spark.sql.adaptive.enabled</b> is set to <b>true</b> .                                                                                                                                                             | true          |
| spark.sql.adaptive.skewJoin.skewedPartitionFactor           | This parameter is a multiplier used to determine whether a partition is a data skew partition. If the data size of a partition exceeds the value of this parameter multiplied by the median of the all partition sizes except this partition and exceeds the value of <b>spark.sql.adaptive.skewJoin.skewedPartitionThresholdInBytes</b> , this partition is considered as a data skew partition. | 5             |
| spark.sql.adaptive.skewJoin.skewedPartitionThresholdInBytes | If the partition size (unit: byte) is greater than the threshold as well as the product of the <b>spark.sql.adaptive.skewJoin.skewedPartitionFactor</b> value and the median partition size, skew occurs in the partition. Ideally, the value of this parameter should be greater than that of <b>spark.sql.adaptive.advisoryPartitionSizeInBytes</b> .                                           | 256MB         |
| spark.sql.adaptive.nonEmptyPartitionRatioForBroadcastJoin   | If the ratio of non-null partitions is less than the value of this parameter when two tables are joined, broadcast hash join cannot be properly performed regardless of the partition size. This parameter is valid only when <b>spark.sql.adaptive.enabled</b> is set to <b>true</b> .                                                                                                           | 0.2           |

### 26.2.7.25 Configuring Event Log Rollover

#### Scenario

When the event log mode is enabled for Spark, that is, **spark.eventLog.enabled** is set to **true**, events are written to a configured log file to record the program running process. If a program, for example JDBCServer or Spark Streaming, runs

for a long period of time and has run many jobs and tasks during this period, many events are recorded in the log file, significantly increasing the file size.

When log rollover is enabled, metadata events are written into the log file and job events are written into a new log file (whether a job event is written to the new log file depends on the file size). Metadata events include EnvironmentUpdate, BlockManagerAdded, BlockManagerRemoved, UnpersistRDD, ExecutorAdded, ExecutorRemoved, MetricsUpdate, ApplicationStart, ApplicationEnd, and LogStart. Job events include StageSubmitted, StageCompleted, TaskResubmit, TaskStart, TaskEnd, TaskGettingResult, JobStart, and JobEnd. For Spark SQL applications, job events also include ExecutionStart and ExecutionEnd.

The UI for the HistoryServer service of Spark is obtained by reading and parsing these log files. The memory size is preset before the HistoryServer process starts. Therefore, when the size of log files is large, loading and parsing these files may cause problems such as insufficient memory and driver GC.

To load large log files in small memory mode, you need to enable log rollover for large applications. Generally, it is recommended that this function be enabled for long-running applications.

## Parameters

Log in to FusionInsight Manager, choose **Cluster > Services > Spark2x > Configurations**, click **All Configurations**, and search for the following parameters.

| Parameter                              | Description                                                                                                                                                                                    | Default Value |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| spark.eventLog.rolling.enabled         | Whether to enable rollover for event log files. If this parameter is set to <b>true</b> , the size of each event log file is reduced to the configured size.                                   | true          |
| spark.eventLog.rolling.maxFileSize     | Maximum size of the event log file to be rolled over when <b>spark.eventlog.rolling.enabled</b> is set to <b>true</b> .                                                                        | 128M          |
| spark.eventLog.compression.codec       | Codec used to compress event logs. By default, Spark provides four types of codecs: LZ4, LZF, Snappy, and ZSTD. If this parameter is not specified, <b>spark.io.compression.codec</b> is used. | None          |
| spark.eventLog.logStageExecutorMetrics | Whether to write each stage peak value (for each executor) of executor metrics to the event log.                                                                                               | false         |

## 26.2.8 Adapting to the Third-party JDK When Ranger Is Used

### Scenarios

When Ranger is used as the permission management service of Spark SQL, the certificate in the cluster is required for accessing RangerAdmin. If you use a third-party JDK instead of the JDK or JRE in the cluster, RangerAdmin fails to be accessed. As a result, the Spark application fails to be started.

In this scenario, you need to perform the following operations to import the certificate in the cluster to the third-party JDK or JRE.

### Configuration Method

**Step 1** Run the following command to export the certificate from the cluster:

1. Install the cluster client in a client, for example, `/opt/client`.
2. Run the following command to switch to the client installation directory:  
**cd /opt/client**
3. Run the following command to configure environment variables:  
**source bigdata\_env**
4. Generate the certificate file.  
**keytool -export -alias fusioninsightsubroot -storepass changeit -keystore /opt/client/JRE/jre/lib/security/cacerts -file fusioninsightsubroot.crt**

**Step 2** Import the certificate in the cluster to the third-party JDK or JRE.

Copy the `fusioninsightsubroot.crt` file generated in [Step 1](#) to the third-party JRE node, set the `JAVA_HOME` environment variable of the node, and run the following command to import the certificate:

```
keytool -import -trustcacerts -alias fusioninsightsubroot -storepass changeit -file fusioninsightsubroot.crt -keystore MY_JRE/lib/security/cacerts
```

#### NOTE

`MY_JRE` indicates the installation path of the third-party JRE. Change it based on the site requirements.

----End

## 26.3 Spark2x Logs

### Log Description

Log paths:

- Executor run log: `${BIGDATA_DATA_HOME}/hadoop/data${i}/nm/containerlogs/application_${appid}/container_${contid}`

 **NOTE**

The logs of running tasks are stored in the preceding path. After the running is complete, the system determines whether to aggregate the logs to an HDFS directory based on the Yarn configuration. For details, see [Common YARN Parameters](#).

- Other logs: `/var/log/Bigdata/spark2x`

**Log archiving rule:**

- When tasks are submitted in **yarn-client** or **yarn-cluster** mode, executor log files are stored each time when the size of the log files reaches 50 MB. A maximum of 10 log files can be reserved without being compressed.
- The JobHistory2x log file is backed up each time when the size of the log file reaches 100 MB. A maximum of 100 log files can be reserved without being compressed.
- The JDBCServer2x log file is backed up each time when the size of the log file reaches 100 MB. A maximum of 100 log files can be reserved without being compressed.
- The IndexServer2x log file is backed up each time when the size of the log file reaches 100 MB. A maximum of 100 log files can be reserved without being compressed.
- The JDBCServer2x audit log file is backed up each time when the size of the log file reaches 20 MB by default. A maximum of 20 log files can be reserved without being compressed.
- The log file size and the number of compressed files to be reserved can be configured on FusionInsight Manager.

**Table 26-62** Spark2x log list

| Log Type             | Name                                    | Description                                                   |
|----------------------|-----------------------------------------|---------------------------------------------------------------|
| SparkResource2x logs | spark.log                               | Spark2x service initialization log                            |
|                      | prestart.log                            | Prestart script log                                           |
|                      | cleanup.log                             | Cleanup log file for instance installation and uninstallation |
|                      | spark-availability-check.log            | Spark2x service health check log                              |
|                      | spark-service-check.log                 | Spark2x service check log                                     |
| JDBCServer2x logs    | JDBCServer-start.log                    | JDBCServer2x startup log                                      |
|                      | JDBCServer-stop.log                     | JDBCServer2x stop log                                         |
|                      | JDBCServer.log                          | JDBCServer2x run log on the server                            |
|                      | jdbc-state-check.log                    | JDBCServer2x health check log                                 |
|                      | jdbcservice-omm-pid***-gc.log.*.current | JDBCServer2x process GC log                                   |

| Log Type              | Name                                                                                         | Description                                                                                              |
|-----------------------|----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
|                       | spark-omm-<br>org.apache.spark.sql.hive.t<br>hriftserver.HiveThriftProxy<br>Server2-***.out* | JDBCServer2x process startup<br>log. If the process stops, the<br><b>jstack</b> information is printed.  |
| JobHistory2x logs     | jobHistory-start.log                                                                         | JobHistory2x startup log                                                                                 |
|                       | jobHistory-stop.log                                                                          | JobHistory2x stop log                                                                                    |
|                       | JobHistory.log                                                                               | JobHistory2x running process<br>log                                                                      |
|                       | jobhistory-omm-pid***-<br>gc.log.*.current                                                   | JobHistory2x process GC log                                                                              |
|                       | spark-omm-<br>org.apache.spark.deploy.hi<br>story.HistoryServer-<br>***.out*                 | JobHistory2x process startup<br>log. If the process stops, the<br><b>jstack</b> information is printed.  |
| IndexServer2x<br>logs | IndexServer-start.log                                                                        | IndexServer2x startup log                                                                                |
|                       | IndexServer-stop.log                                                                         | IndexServer2x stop log                                                                                   |
|                       | IndexServer.log                                                                              | IndexServer2x run log on the<br>server                                                                   |
|                       | indexserver-state-<br>check.log                                                              | IndexServer2x health check log                                                                           |
|                       | indexserver-omm-pid***-<br>gc.log.*.current                                                  | IndexServer2x process GC log                                                                             |
|                       | spark-omm-<br>org.apache.spark.sql.hive.t<br>hriftserver.IndexServerPro<br>xy-***.out*       | IndexServer2x process startup<br>log. If the process stops, the<br><b>jstack</b> information is printed. |
| Audit Log             | jdbcserver-audit.log                                                                         | JDBCServer2x audit log                                                                                   |
|                       | ranger-audit.log                                                                             |                                                                                                          |

## Log levels

**Table 26-63** describes the log levels supported by Spark2x. The priorities of log levels are ERROR, WARN, INFO, and DEBUG in descending order. Logs whose levels are higher than or equal to the specified level are printed. The number of printed logs decreases as the specified log level increases.

**Table 26-63** Log levels

| Level | Description                                                                              |
|-------|------------------------------------------------------------------------------------------|
| ERROR | Error information about the current event processing                                     |
| WARN  | Exception information about the current event processing                                 |
| INFO  | Logs of this level record normal running status information about the system and events. |
| DEBUG | Logs of this level record the system information and system debugging information.       |

To modify log levels, perform the following operations:

 **NOTE**

By default, the service does not need to be restarted after the Spark2x log levels are configured.

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration**.
- Step 3** Select **All Configurations**.
- Step 4** On the menu bar on the left, select the log menu of the target role.
- Step 5** Select a desired log level.
- Step 6** Click **Save**. Then, click **OK**.

----End

## Log Format

**Table 26-64** Log Format

| Type    | Format                                                                                                                                                            | Example                                                                                              |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| Run log | <yyyy-MM-dd<br>HH:mm:ss,SSS> <Log level> <br><Name of the thread that<br>generates the log> <br><Message in the log> <br><Location where the log<br>event occurs> | 2014-09-22 11:16:23,980<br>INFO DAGScheduler: Final<br>stage: Stage 0(reduce at<br>SparkPi.scala:35) |

## 26.4 Obtaining Container Logs of a Running Spark Application

Container logs of running Spark applications are distributed on multiple nodes. This section describes how to quickly obtain container logs.

### Scenario Description

You can run the **yarn logs** command to obtain the logs of applications running on Yarn. In different scenarios, you can run the following commands to obtain required logs:

1. Obtain complete logs of the application: **yarn logs --applicationId <appld> -out <outputDir>**.

Example: **yarn logs --applicationId application\_1574856994802\_0016 -out /opt/test**

The following figure shows the command output.

- a. If the application is running, container logs in the **dead** state cannot be obtained.
- b. If the application is stopped, all archived container logs can be obtained.

2. Obtain logs of a specified container: **yarn logs -applicationId <appld> -containerId <containerId>**.

Example: **yarn logs -applicationId application\_1574856994802\_0018 -containerId container\_e01\_1574856994802\_0018\_01\_000003**

The following figure shows the command output.

- a. If the application is running, container logs in the **dead** state cannot be obtained.
- b. If the application is stopped, you can obtain logs of any container.

3. Obtain container logs in any state: **yarn logs -applicationId <appld> -containerId <containerId> -nodeAddress <nodeAddress>**

Example: **yarn logs -applicationId application\_1574856994802\_0019 -containerId container\_e01\_1574856994802\_0019\_01\_000003 -nodeAddress 192-168-1-1:8041**

Execution result: Logs of any container can be obtained.

#### NOTE

You need to set *nodeAddress* in the command. You can run the following command to obtain the value:

```
yarn node -list -all
```

## 26.5 Small File Combination Tools

### Tool Overview

In a large-scale Hadoop production cluster, HDFS metadata is stored in the NameNode memory, and the cluster scale is restricted by the memory limitation



of each NameNode. If there are a large number of small files in the HDFS, a large amount of NameNode memory is consumed, which greatly reduces the read and write performance and prolongs the job running time. Based on the preceding information, the small file problem is a key factor that restricts the expansion of the Hadoop cluster.

This tool provides the following functions:

1. Checks the number of small files whose size is less than the threshold configured by the user in tables and returns the average size of all data files in the table directory.
2. Provides the function of combination table files. Users can set the average file size after combination.

## Supported Table Types

Spark: Parquet, ORC, CSV, Text, and Json.

Hive: Parquet, ORC, CSV, Text, RCFile, Sequence and Bucket.

### NOTE

1. After tables with compressed data are merged, Spark uses the default compression format Snappy for data compression. You can configure **spark.sql.parquet.compression.codec** (available values: **uncompressed**, **gzip**, and **snappy**) and **spark.sql.orc.compression.codec** (available values: **uncompressed**, **zlib**, **lzo**, and **snappy**) on the client to select the compression format for the Parquet and ORC tables. Compression formats available for Hive and Spark tables are different. Except the preceding compression formats, other compression formats are not supported.
2. To merge bucket table data, you need to add the following configurations to the **hive-site.xml** file on the Spark2x client:

```
<property>
<name>hive.enforce.bucketing</name>
<value>>false</value>
</property>
<property>
<name>hive.enforce.sorting</name>
<value>>false</value>
</property>
```
3. Spark does not support the feature of encrypting data columns in Hive.

## Tool Usage

Download and install the client. For example, the installation directory is **/opt/client**. Go to **/opt/client/Spark2x/spark/bin** and run the **mergetool.sh** script.

### Environment variables loading

```
source /opt/client/bigdata_env
```

```
source /opt/client/Spark2x/component_env
```

### Scanning function

Command: **sh mergetool.sh scan <db.table> <filesize>**

The format of *db.table* is *Database name, Table name*. *filesize* is the user-defined threshold of the small file size (unit: MB). The returned result is the number of

files that is smaller than the threshold and the average size of data files in the table directory.

Example: **sh mergetool.sh scan default.table1 128**

### Combination function

Command: **sh mergetool.sh merge <db.table> <filesize> <shuffle>**

The format of *db.table* is *Database name, Table name*. **filesize** is the user-defined average file size after file combination (unit: MB). **shuffle** is a Boolean value, and the value is **true** or **false**, which is used to configure whether to allow data to be shuffled during the merge.

Example: **sh mergetool.sh merge default.table1 128 false**

If the following information is displayed, the operation is successful:

SUCCESS: Merge succeeded

#### NOTE

1. Ensure that the current user is the owner of the merged table.
2. Before combination, ensure that HDFS has sufficient storage space, greater than the size of the combined table.
3. Table data must be combined separately. If a table is read during table data combination, the file may not be found temporarily. After the combination is complete, this problem is resolved. During the combination, do not write data to the corresponding tables. Otherwise, data inconsistency may occur.
4. If an error occurs indicating that the file does not exist when the query of data in a partitioned table is performed on the session spark-beeline/spark-sql that is always in the connected status. You can run the **refresh table** *Table name* command as prompted to query the data again.
5. Configure **filesize** based on the site requirements. For example, you can set **filesize** to a value greater than the average during file merging after obtaining the average file size by file scan. Otherwise, the number of files may increase after the file merging.
6. During the file merging, data in the original tables is removed to the recycle bin. In the case of any exception occurs on the data after file merging, the data in the original tables is used to replace the damaged data. If an exception occurs during the process, restore the data in the trash directory by using the **mv** command in HDFS.
7. In the HDFS router federation scenario, if the target NameService of the table root path is different from that of the root path **/user**, you need to manually clear the original table files stored in the recycle bin during the second combination. Otherwise, the combination fails.
8. This tool uses the configuration of the client. Performance optimization can be performed modifying required configuration in the client configuration file.

### shuffle configuration

For the combination function, you can roughly estimate the change on the number of partitions before and after the combination.

Generally, if the number of old partitions is greater than the number of new partitions, set **shuffle** to **false**. However, if the number of old partitions is much greater than that of new partitions (for example, more than 100 times), you can set **shuffle** to **true** to increase the degree of parallelism and improve the combination speed.

**NOTICE**

- If **shuffle** is set to **true** (repartition), the performance is improved. However, due to the particularity of the Parquet and ORC storage modes, repartition will reduce the compression ratio and the total size of the table in HDFS increases by 1.3 times.
- If **shuffle** is set to **false** (coalesce), the merged files may have some difference in size, which is close to the value of the configured **filesize**.

**Log storage location**

The default log storage path is `/tmp/SmallFilesLog.log4j`. To customize the log storage path, you can configure `log4j.appender.logfile.File` in `/opt/client/Spark2x/spark/tool/log4j.properties`.

## 26.6 Using CarbonData for First Query

### Tool Overview

The first query of CarbonData is slow, which may cause a delay for nodes that have high requirements on real-time performance.

The tool provides the following functions:

- Preheat the tables that have high requirements on query delay for the first time.

### Tool Usage

Download and install the client. For example, the installation directory is `/opt/client`. Go to the `/opt/client/Spark2x/spark/bin` directory and run `start-prequery.sh`.

Configure `prequeryParams.properties` by referring to [Table 26-65](#).

**Table 26-65** Parameters

Parameter	Description	Example
<code>spark.prequery.period.max.minute</code>	Maximum preheating duration, in minutes.	60
<code>spark.prequery.tables</code>	Table name configuration, <i>database.table:int</i> . The table name supports the wildcard (*). <b>int</b> indicates the duration (unit: day) within which the table is updated before it is preheated.	default.test*:10

Parameter	Description	Example
spark.prequery.maxThreads	Maximum number of concurrent threads during preheating	50
spark.prequery.sslEnable	The value is <b>true</b> in security mode and <b>false</b> in non-security mode.	true
spark.prequery.driver	IP address and port number of JDBCServer. The format is <i>IP address:Port number</i> . If multiple servers need to be preheated, enter multiple <i>IP address:Port number</i> of the servers and separate them with commas (,).	192.168.0.2:22550
spark.prequery.sql	SQL statement for preheating. Different statements are separated by colons (:).	SELECT COUNT(*) FROM %s;SELECT * FROM %s LIMIT 1
spark.security.url	URL required by JDBC in security mode	;sasLQop=auth-conf;auth=KERBEROS;principal=spark2x/hadoop.hadoop.com@HADOOP.COM;

 NOTE

The statement configured in **spark.prequery.sql** is executed in each preheated table. The table name is replaced with **%s**.

**Script Usage**

Command format: **sh start-prequery.sh**

To run this command, place **user.keytab** or **jaas.conf** (either of them) and **krb5.conf** (mandatory) in the **conf** directory.

 NOTE

- Currently, this tool supports only Carbon tables.
- This tool initializes the Carbon environment and pre-reads table metadata to JDBCServer. Therefore, this tool is more suitable for multi-active instances and static allocation mode.

## 26.7 Spark2x Performance Tuning

## 26.7.1 Spark Core Tuning

### 26.7.1.1 Data Serialization

#### Scenario

Spark supports the following types of serialization:

- JsonSerializer
- KryoSerializer

Data serialization affects the Spark application performance. In specific data format, KryoSerializer offers 10X higher performance than JsonSerializer. For Int data, performance optimization can be ignored.

KryoSerializer depends on Chill of Twitter. Not all Java Serializable objects support KryoSerializer. Therefore, class must be manually registered.

Serialization involves task serialization and data serialization. Only JsonSerializer can be used for Spark task serialization. JsonSerializer and KryoSerializer can be used for data serialization.

#### Procedure

When the Spark program is running, a large volume of data needs to be serialized during the shuffle and RDD cache procedures. By default, JsonSerializer is used. You can also configure KryoSerializer as the data serializer to improve serialization performance.

Add the following code to enable KryoSerializer to be used:

- Implement the class registrar and manually register the class.

```
package com.etl.common;

import com.esotericsoftware.kryo.Kryo;
import org.apache.spark.serializer.KryoRegistrator;

public class DemoRegistrator implements KryoRegistrator
{
 @Override
 public void registerClasses(Kryo kryo)
 {
 //Class examples are given below. Register the custom classes.
 kryo.register(AggrateKey.class);
 kryo.register(AggrateValue.class);
 }
}
```

You can configure **spark.kryo.registrationRequired** on Spark client. Whether to require registration with Kryo. If set to 'true', Kryo will throw an exception if an unregistered class is serialized. If set to false (the default), Kryo will write unregistered class names along with each object. Writing class names can cause significant performance overhead. This operation will affect the system performance. If the value of **spark.kryo.registrationRequired** is configured to **true**, you need to manually register the class. For a class that is not serialized, the system will not automatically write the class name, but display an exception. Compare the configuration of **true** with that of **false**, the configuration of **true** has the better performance.

- Configure KryoSerializer as the data serializer and class registrar.

```
val conf = new SparkConf()
conf.set("spark.serializer", "org.apache.spark.serializer.KryoSerializer")
.set("spark.kryo.registrator", "com.etl.common.DemoRegistrator")
```

## 26.7.1.2 Optimizing Memory Configuration

### Scenario

Spark is a memory-based computing frame. If the memory is insufficient during computing, the Spark execution efficiency will be adversely affected. You can determine whether memory becomes the performance bottleneck by monitoring garbage collection (GC) and evaluating the resilient distributed dataset (RDD) size in the memory, and take performance optimization measures.

To monitor GC of node processes, add the `-verbose:gc -XX:+PrintGCDetails -XX:+PrintGCTimeStamps` parameter to the **spark.driver.extraJavaOptions** and **spark.executor.extraJavaOptions** in the client configuration file **conf/spark-default.conf**. If "Full GC" is frequently reported, GC needs to be optimized. Cache the RDD and query the RDD size in the log. If a large value is found, change the RDD storage level.

### Procedure

- To optimize GC, adjust the ratio of the young generation and tenured generation. Add **-XX:NewRatio** parameter to the **spark.driver.extraJavaOptions** and **spark.executor.extraJavaOptions** in the client configuration file **conf/spark-default.conf**. For example, export `SPARK_JAVA_OPTS="-XX:NewRatio=2"`. The new generation accounts for 1/3 of the heap, and the tenured generation accounts for 2/3.
- Optimize the RDD data structure when compiling Spark programs.
  - Use primitive arrays to replace fastutil arrays, for example, use fastutil library.
  - Avoid nested structure.
  - Avoid using String in keys.
- Suggest serializing the RDDs when developing Spark programs.

By default, data is not serialized when RDDs are cached. You can set the storage level to serialize the RDDs and minimize memory usage. For example:

```
testRDD.persist(StorageLevel.MEMORY_ONLY_SER)
```

## 26.7.1.3 Setting the DOP

### Scenario

The degree of parallelism (DOP) specifies the number of tasks to be executed concurrently. It determines the number of data blocks after the shuffle operation. Configure the DOP to improve the processing capability of the system.

Query the CPU and memory usage. If the tasks and data are not evenly distributed among nodes, increase the DOP. Generally, set the DOP to two or three times that of the total CPUs in the cluster.

## Procedure

Configure the DOP parameter using one of the following methods based on the actual memory, CPU, data, and application logic conditions:

- Configure the DOP parameter in the operation function that generates the shuffle. This method has the highest priority.  
`testRDD.groupByKey(24)`
- Configure the DOP using **spark.default.parallelism**. This method has the lower priority than the preceding one.  
`val conf = new SparkConf();  
conf.set("spark.default.parallelism", 24);`
- Configure the value of **spark.default.parallelism** in the **\$SPARK\_HOME/conf/spark-defaults.conf** file. This method has the lowest priority.  
`spark.default.parallelism 24`

### 26.7.1.4 Using Broadcast Variables

#### Scenario

Broadcast distributes data sets to each node. It allows data to be obtained locally when a data set is needed during a Spark task. If broadcast is not used, data serialization will be scheduled to tasks each time when a task requires data sets. It is time-consuming and makes the task get bigger.

1. If a data set will be used by each slice of a task, broadcast the data set to each node.
2. When small and big tables need to be joined, broadcast small tables to each node. This eliminates the shuffle operation, changing the join operation into a common operation.

#### Procedure

Add the following code to broadcast the testArr data to each node:

```
def main(args: Array[String]) {
 ...
 val testArr: Array[Long] = new Array[Long](200)
 val testBroadcast: Broadcast[Array[Long]] = sc.broadcast(testArr)
 val resultRdd: RDD[Long] = inpputRdd.map(input => handleData(testBroadcast, input))
 ...
}

def handleData(broadcast: Broadcast[Array[Long]], input: String) {
 val value = broadcast.value
 ...
}
```

### 26.7.1.5 Using the external shuffle service to improve performance

#### Scenario

When the Spark system runs applications that contain a shuffle process, an executor process also writes shuffle data and provides shuffle data for other executors in addition to running tasks. If the executor is heavily loaded and GC is triggered, the executor cannot provide shuffle data for other executors, affecting task running.

The external shuffle service is an auxiliary service in NodeManager. It captures shuffle data to reduce the load on executors. If GC occurs on an executor, tasks on other executors are not affected.

## Procedure

- Step 1** Log in to FusionInsight Manager.
- Step 2** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**. Select **All Configurations**.
- Step 3** Choose **SparkResource2x** > **Default** and modify the following parameters.

**Table 26-66** Parameter list

Parameter	Default Value	Changed To
spark.shuffle.service.enabled	false	true

- Step 4** Restart the Spark2x service for the configuration to take effect.

### NOTE

To use the External Shuffle Service function on the Spark2x client, you need to download and install the Spark2x client again.

----End

## 26.7.1.6 Configuring Dynamic Resource Scheduling in Yarn Mode

### Scenario

Resources are a key factor that affects Spark execution efficiency. When a long-running service (such as the JDBCServer) is allocated with multiple executors without tasks but resources of other applications are insufficient, resources are wasted and scheduled improperly.

Dynamic resource scheduling can add or remove executors of applications in real time based on the task load. In this way, resources are dynamically scheduled to applications.

## Procedure

- Step 1** Configure the external shuffle service.
- Step 2** Log in to FusionInsight Manager, and choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration** > **All Configurations**. Enter **spark.dynamicAllocation.enabled** in the search box and set its value to **true** to enable the dynamic resource scheduling function. This function is disabled by default.

----End

[Table 26-67](#) lists some optional configuration items.



**Table 26-67** Parameters for dynamic resource scheduling

Configuration Item	Description	Default Value
spark.dynamicAllocation.minExecutors	Indicates the minimum number of executors.	0
spark.dynamicAllocation.initialExecutors	Indicates the number of initial executors.	0
spark.dynamicAllocation.maxExecutors	Indicates the maximum number of executors.	2048
spark.dynamicAllocation.schedulerBacklogTimeout	Indicates the first timeout period for scheduling.	1s
spark.dynamicAllocation.sustainedSchedulerBacklogTimeout	Indicates the second and later timeout interval for scheduling.	1s
spark.dynamicAllocation.executorIdleTimeout	Indicates the idle timeout interval for common executors.	60s
spark.dynamicAllocation.cachedExecutorIdleTimeout	Indicates the idle timeout interval for executors with cached blocks.	<ul style="list-style-type: none"> <li>• JDBCServer2x: 2147483647s</li> <li>• IndexServer2x: 2147483647s</li> <li>• SparkResource2x: 120</li> </ul>

 **NOTE**

The external shuffle service must be configured before using the dynamic resource scheduling function.

### 26.7.1.7 Configuring Process Parameters

#### Scenario

There are three processes in Spark on Yarn mode: driver, ApplicationMaster, and executor. The Driver and Executor handle the scheduling and running of the task. The ApplicationMaster handles the start and stop of the container.

Therefore, the configuration of the driver and executor is very important to run the Spark application. You can optimize the performance of the Spark cluster according to the following procedure.

#### Procedure

**Step 1** Configure the driver memory.

The driver schedules tasks and communicates with the executor and the ApplicationMaster. Add driver memory when the number and parallelism level of the tasks increases.

You can configure the driver memory based on the number of the tasks.

- Set **spark.driver.memory** in **spark-defaults.conf** to a proper value.
- Add the **--driver-memory MEM** parameter to configure the memory when using the **spark-submit** command.

### Step 2 Configure the number of the executors.

One core in an executor can run one task at the same time. Therefore, more tasks can be processed at the same time if you increase the number of the executors. You can add the number of the executors to increase the efficiency if resources are sufficient.

- Set **spark.executor.instance** in **spark-defaults.conf** or **SPARK\_EXECUTOR\_INSTANCES** in **spark-env.sh** to a proper value.
- Add the **--num-executors NUM** parameter to configure the number of the executors when using the **spark-submit** command.

### Step 3 Configure the number of the executor cores.

Multiple cores in an executor can run multiple tasks at the same time, which increases the task concurrency. However, because all cores share the memory of an executor, you need to balance the memory and the number of cores.

- Set **spark.executor.cores** in **spark-defaults.conf** or **SPARK\_EXECUTOR\_CORES** in **spark-env.sh** to a proper value.
- When you run the **spark-submit** command, add the **--executor-cores NUM** parameter to set the number of executor cores.

### Step 4 Configure the executor memory.

The executor memory is used for task execution and communication. You can increase the memory for a big task that needs more resources, and reduce the memory to increase the concurrency level for a small task that runs fast.

- Set **spark.executor.memory** in **spark-defaults.conf** or **SPARK\_EXECUTOR\_MEMORY** in **spark-env.sh** to a proper value.
- When you run the **spark-submit** command, add the **--executor-memory MEM** parameter to set the memory.

----End

## Example

- During the **spark wordcount** calculation, the amount of data is 1.6 TB and the number of the executors is 250.

The execution fails under the default configuration, and the **Futures timed out** and **OOM** errors occur.

However each task of wordcount is small and runs fast, the amount of the data is big and the tasks are too many. Therefore the objects on the driver end become huge when there are many tasks. Besides the fact that the executor communicates with the driver once each task is finished, the

problem of disconnection between processes caused by insufficient memory occurs.

The application runs successfully when the memory of the Driver is set to 4 GB.

- Many errors still occurred in the default configuration when running TPC-DS test on JDBCServer, such as "Executor Lost". When there is 30 GB of driver memory, 2 executor cores, 125 executors, and 6 GB of executor memory, all tasks can be successfully executed.

### 26.7.1.8 Designing the Direction Acyclic Graph (DAG)

#### Scenario

Optimal program structure helps increase execution efficiency. During application programming, avoid shuffle operations and combine narrow-dependency operations.

#### Procedure

This topic describes how to design the DAG using the following example:

- **Data format:** Time when a vehicle passes a toll station, license plate number, toll station number, and more
- **Logic:** Two vehicles are determined to be traveling together if the following conditions are met:
  - Both vehicles pass the same toll stations in the same sequence.
  - The difference between the time that the vehicles pass the same toll station is smaller than a specified value.

There are two implementation ways for this example. [Figure 26-6](#) shows the logic of implementation 1 and [Figure 26-7](#) shows logic of implementation 2.

**Figure 26-6** Implementation logic 1



Logic description:

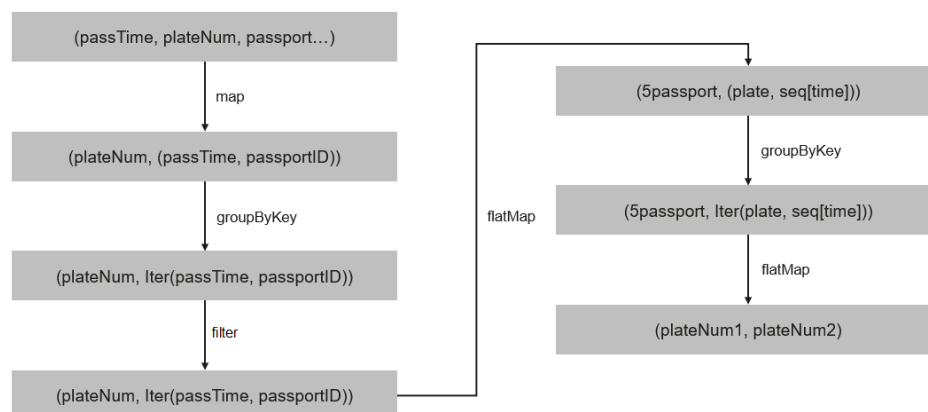
1. Collect information about the toll stations passed by each vehicle based on the vehicle license plate number and sort the toll stations.

- The following data is obtained: vehicle license plate number 1, [(time, toll station 3), (time, toll station 2), (time, toll station 4), (time, toll station 5)]
2. Determine the sequence in which the vehicle passed through.
    - (toll station 3, (vehicle license plate number 1, time, 1st toll station))
    - (toll station 2, (vehicle license plate number 1, time, 2nd toll station))
    - (toll station 4, (vehicle license plate number 1, time, 3rd toll station))
    - (toll station 5, (vehicle license plate number 1, time, 4th toll station))
  3. Aggregate data by toll station.
    - toll station 1, [(vehicle license plate number 1, time, 1st toll station), (vehicle license plate number 2, time, 5th toll station), (vehicle license plate number 3, time, 2nd toll station)]
  4. Determine whether the time difference that two vehicles passed through the same toll station is below the specified value. If yes, fetch information about the two vehicles.
    - (vehicle license plate number 1, vehicle license plate number 2),(1st toll station, 5th toll station)
    - (vehicle license plate number 1, vehicle license plate number 3),(1st toll station, 2nd toll station)
  5. Aggregate data based on the vehicle license plate numbers that passed through the same toll stations.
    - (vehicle license plate number 1, vehicle license plate number 2), [(1st toll station, 5th toll station), (2nd toll station, 6th toll station), (1st toll station, 7th toll station), (3rd toll station, 8th toll station)]
  6. If the two vehicles pass through the same toll stations in sequence, for example, toll stations 3, 4, 5 are the first, second, and third toll station passed by vehicle 1 and the 6th, 7th, and 8th toll station passed by vehicle 2, and the number of toll stations meets the specified requirements, the two vehicles are determined to be traveling together.

The logic of implementation 1 has the following disadvantages:

- The logic is complex.
- Too many shuffle operations affect performance.

**Figure 26-7** Implementation logic 2



Logic description:

1. Collect information about the toll stations passed by each vehicle based on the vehicle license plate number and sort the toll stations.  
The following data is obtained: vehicle license plate number 1, [(time, toll station 3), (time, toll station 2), (time, toll station 4), (time, toll station 5)]
2. Based on the number of toll stations (the number is 3 in this example) that must be passed by these vehicles, divide the toll station sequence as follows:  
toll station 3 > toll station 2 > toll station 4, (vehicle license plate number 1, [time passing through toll station 3, time passing through toll station 2, time passing through toll station 4])  
toll station 2 > toll station 4 > toll station 5, (vehicle license plate number 1, [time passing through toll station 2, time passing through toll station 4, time passing through toll station 5])
3. Aggregate information about vehicles that pass the same toll stations in the same sequence.  
toll station 3 > toll station 2 > toll station 4, [(vehicle license plate number 1, [time passing through toll station 3, time passing through toll station 2, time passing through toll station 4]), (vehicle license plate number 2, [time passing through toll station 3, time passing through toll station 2, time passing through toll station 4]), (vehicle license plate number 3, [time passing through toll station 3, time passing through toll station 2, time passing through toll station 4])]
4. Determine whether the time difference that these vehicles passed through the same toll station is below the specified value. If yes, the vehicles are determined to be traveling together.

The logic of implementation 2 has the following advantages:

- The logic is simplified.
- One **groupByKey** is reduced, that is, one less shuffle operation is performed. It helps improve performance.

### 26.7.1.9 Experience

#### Use mapPartitions to calculate data by partition.

If the overhead of each record is high, for example:

```
rdd.map{x=>conn=getDBConn;conn.write(x.toString);conn.close}
```

Use mapPartitions to calculate data by partition.

```
rdd.mapPartitions(records => conn.getDBConn;for(item <- records)
write(item.toString); conn.close)
```

Use mapPartitions to flexibly operate data. For example, to calculate the TopN of a large data, mapPartitions can be used to calculate the TopN of each partition and then sort the TopN of all partitions when N is small. Compared with sorting full data for the TopN, this method has the higher efficiency.

#### Use coalesce to adjust the number of slices.

Use coalesce to adjust the number of slices. There are two coalesce functions:

```
coalesce(numPartitions: Int, shuffle: Boolean = false)
```

When **shuffle** is set to **true**, the function is the same as `repartition(numPartitions: Int)`. Partitions are recreated using the shuffle. When **shuffle** is set to **false**, partitions of the parent resilient distributed datasets (RDD) are calculated in the same task. In this case, if the value of **numPartitions** is larger than the number of sections of the parent RDD, partitions will not be recreated.

The following scenario is encountered, you can choose the coalesce operator:

- If the previous operation involves a large number of filters, use coalesce to minimize the number of zero-loaded tasks. In `coalesce(numPartitions, false)`, the value of **numPartitions** is smaller than the number of sections of the parent RDD.
- Use coalesce when the number of slices entered is too big to execute.
- Use coalesce when the programs are suspended in the shuffle operation because of a large number of tasks or the Linux resources are limited. In this case, use `coalesce(numPartitions, true)` to recreate partitions.

## Configure a localDir for each disk.

During the shuffle procedure of Spark, data needs to be written into local disks. The performance bottleneck of Spark is shuffle, and the bottleneck of shuffle is the I/O. To improve the I/O performance, you can configure multiple disks to implement concurrent data writing. If a node is mounted with multiple disks, configure a Spark local Dir for each disk. This can effectively distribute shuffle files in multiple locations, improving disk I/O efficiency. The performance cannot be improved effectively if a disk is configured with multiple directories.

## Collect small data sets.

The collect operation does not apply to a large data volume.

When the collect operation is performed, the Executor data will be sent to the Driver. Before performing this operation, ensure that the memory of Driver is sufficient. Otherwise, the Driver process may encounter an OutOfMemory error. If the data volume is unknown, perform the `saveAsTextFile` operation to write data into the HDFS. If the data volume is known and the Driver has sufficient memory, perform the collect operation.

## Use reduceByKey

`reduceByKey` causes local aggregation on the Map side, which offers a smooth shuffle procedure. The shuffle operations, like `groupByKey`, will not perform aggregation on the Map side. Therefore, use `reduceByKey` as possible as you can, and avoid `groupByKey().map(x=>(x._1,x._2.size))`.

## Broadcast map instead of array.

If table query is required for each record of the data transmitted from the Driver side, broadcast the data in the set/map instead of Iterator. The query speed of Set/Map is approximately  $O(1)$ , while the query speed of Iterator is  $O(n)$ .

## Avoid data skew.

If data skew occurs (certain data volume is extremely large), the execution time of tasks is inconsistent even if there is no Garbage Collection (GC).

- Redefine the keys. Use keys of smaller granularity to optimize the task size.
- Modify the degree of parallelism (DOP).

## Optimize the data structure.

- Store data by column. As a result, only the required columns are scanned when data is read.
- When using the Hash Shuffle, set **spark.shuffle consolidateFiles** to **true** to combine the intermediate files of shuffle, minimize the number of shuffle files and file I/O operations, and improve performance. The number of final files is the number of reduce tasks.

## 26.7.2 Spark SQL and DataFrame Tuning

### 26.7.2.1 Optimizing the Spark SQL Join Operation

#### Scenario

When two tables are joined in Spark SQL, the broadcast function (see section "Using Broadcast Variables") can be used to broadcast tables to each node. This minimizes shuffle operations and improves task execution efficiency.

#### NOTE

The join operation refers to the inner join operation only.

#### Procedure

The following describes how to optimize the join operation in Spark SQL. Assume that both tables A and B have the **name** column. Join tables A and B as follows:

1. Estimate the table sizes.

Estimate the table size based on the size of data loaded each time.

You can also check the table size in the directory of the Hive database. In the **hive-site.xml** configuration file of Spark, view the Hive database directory, which is **/user/hive/warehouse** by default. The default Hive database directory for multi-instance Spark is **/user/hive/warehouse**, for example, **/user/hive1/warehouse**.

```
<property>
 <name>hive.metastore.warehouse.dir</name>
 <value>${test.warehouse.dir}</value>
 <description></description>
</property>
```

Run the **hadoop** command to check the size of the table. For example, run the following command to view the size of table **A**:

```
hadoop fs -du -s -h ${test.warehouse.dir}/a
```

 **NOTE**

To perform the broadcast operation, ensure that at least one table is not empty.

2. Configure a threshold for automatic broadcast.

The threshold for triggering broadcast for a table is 10485760 (that is, 10 MB) in Spark. If either of the table sizes is smaller than 10 MB, skip this step.

**Table 26-68** lists configuration parameters of the threshold for automatic broadcasting.

**Table 26-68** Parameter description

Parameter	Default Value	Description
spark.sql.autoBroadcastJoinThreshold	10485760	<p>Indicates the maximum value for the broadcast configuration when two tables are joined.</p> <ul style="list-style-type: none"> <li>• When the size of a field in a table involved in an SQL statement is less than the value of this parameter, the system broadcasts the SQL statement.</li> <li>• If the value is set to <b>-1</b>, broadcast is not performed.</li> </ul>

Methods for configuring the threshold for automatic broadcasting:

- Set **spark.sql.autoBroadcastJoinThreshold** in the **spark-defaults.conf** configuration file of Spark.

```
spark.sql.autoBroadcastJoinThreshold = <size>
```

- Run the Hive command to set the threshold. Before joining the tables, run the following command:

```
SET spark.sql.autoBroadcastJoinThreshold=<size>;
```

3. Join the tables.

- The size of each table is smaller than the threshold.
  - If the size of table A is smaller than that of table B, run the following command:  
SELECT A.name FROM B JOIN A ON A.name = B.name;
  - If the size of table B is smaller than that of table A, run the following command:  
SELECT A.name FROM A JOIN B ON A.name = B.name;
- One table size is smaller than the threshold, while the other table size is greater than the threshold.  
Broadcast the smaller table.
- The size of each table is greater than the threshold.  
Compare the size of the field involved in the query with the threshold.
  - If the values of the fields in a table are smaller than the threshold, the corresponding data in the table is broadcast.



- If the values of the fields in the two tables are greater than the threshold, do not broadcast either of the table.
4. (Optional) In the following scenarios, you need to run the Analyze command (***ANALYZE TABLE tableName COMPUTE STATISTICS noscan;***) to update metadata before performing the broadcast operation:
- The table to be broadcasted is a newly created partitioned table and the file type is non-Parquet.
  - The table to be broadcasted is a newly updated partitioned table.

## Reference

A task is ended if a timeout occurs during the execution of the to-be-broadcasted table.

By default, BroadCastJoin allows only 5 minutes for the to-be-broadcasted table calculation. If the time is exceeded, a timeout will occur. However, the broadcast task of the to-be-broadcasted table calculation is still being executed, resulting in resource waste.

The following methods can be used to address this issue:

- Modify the value of **spark.sql.broadcastTimeout** to increase the timeout duration.
- Reduce the value of **spark.sql.autoBroadcastJoinThreshold** to disable the optimization of BroadCastJoin.

### 26.7.2.2 Improving Spark SQL Calculation Performance Under Data Skew

#### Scenario

When multiple tables are joined in Spark SQL, skew occurs in join keys and the data volume in some Hash buckets is much higher than that in other buckets. As a result, some tasks with a large amount of data run slowly, resulting low computing performance. Other tasks with a small amount of data are quickly completed, which frees many CPUs and results in a waste of CPU resources.

If the automatic data skew function is enabled, data that exceeds the bucketing threshold is bucketed. Multiple tasks proceed data in one bucket. Therefore, CUP usage is enhanced and the system performance is improved.

#### NOTE

Data that has no skew is bucketed and run in the original way.

Restrictions:

- Only the join between two tables is supported.
- FULL OUTER JOIN data does not support data skew.  
For example, the following SQL statement indicates that the skew of table **a** or table **b** cannot trigger the optimization.  
***select aid FROM a FULL OUTER JOIN b ON aid=bid;***
- LEFT OUTER JOIN data does not support the data skew of the right table.

For example, the following SQL statement indicates that the skew of table **b** cannot trigger the optimization.

***select aid FROM a LEFT OUTER JOIN b ON aid=bid;***

- RIGHT OUTER JOIN does not support the data skew of the left table.

For example, the following SQL statement indicates that the skew of table **a** cannot trigger the optimization.

***select aid FROM a RIGHT OUTER JOIN b ON aid=bid;***

## Configuration Description

Add the following parameters in the following table to the **spark-defaults.conf** configuration file on the Spark driver.

**Table 26-69** Parameter description

Parameter	Description	Default Value
spark.sql.adaptive.enabled	The switch to enable the adaptive execution feature.  Note: If AQE and Static Partition Pruning (DPP) are enabled at the same time, DPP takes precedence over AQE during SparkSQL task execution. As a result, AQE does not take effect. The DPP in the cluster is enabled by default. Therefore, you need to disable it when enabling the AQE.	false
spark.sql.optimizer.dynamicPartitionPruning.enabled	The switch to enable DPP.	true
spark.sql.adaptive.skewJoin.enabled	Specifies whether to enable the function of automatic processing of the data skew in join operations. The function is enabled when this parameter is set to <b>true</b> and <b>spark.sql.adaptive.enabled</b> is set to <b>true</b> .	true
spark.sql.adaptive.skewJoin.skewedPartitionFactor	This parameter is a multiplier used to determine whether a partition is a data skew partition. If the data size of a partition exceeds the value of this parameter multiplied by the median of the all partition sizes except this partition and exceeds the value of <b>spark.sql.adaptive.skewJoin.skewedPartitionThresholdInBytes</b> , this partition is considered as a data skew partition.	5

Parameter	Description	Default Value
spark.sql.adaptive.skewjoin.skewedPartitionThresholdInBytes	If the partition size (unit: byte) is greater than the threshold as well as the product of the <b>spark.sql.adaptive.skewJoin.skewedPartitionFactor</b> value and the median partition size, skew occurs in the partition. Ideally, the value of this parameter should be greater than that of <b>spark.sql.adaptive.advisoryPartitionSizeInBytes..</b>	256MB
spark.sql.adaptive.shuffle.targetPostShuffleInputSize	Minimum amount of shuffle data processed by each task. The unit is byte.	67108864

### 26.7.2.3 Optimizing Spark SQL Performance in the Small File Scenario

#### Scenario

A Spark SQL table may have many small files (far smaller than an HDFS block), each of which maps to a partition on the Spark by default. In other words, each small file is a task. If the small files are great in number, Spark must initiate a large number of tasks. If shuffle operations exist in Spark SQL, the number of hash buckets increases, affecting performance.

In this scenario, you can manually specify the split size of each task to avoid an excessive number of tasks and improve performance.

#### NOTE

If the SQL logic does not involve shuffle operations, this optimization does not improve performance.

#### Configuration

If you want to enable small file optimization, configure the **spark-defaults.conf** file on the Spark client.

**Table 26-70** Parameter description

Parameter	Description	Default Value
spark.sql.files.maxPartitionBytes	The maximum number of bytes that can be packed into a single partition when a file is read. Unit: byte	134217728 (128 MB)

Parameter	Description	Default Value
spark.files.openCostInBytes	The estimated cost to open a file, measured by the number of bytes that can be scanned in the same time. This is used when putting multiple files into a partition. It is better to over estimate, then the partitions with small files will be faster than partitions with larger files.	4 MB

## 26.7.2.4 Optimizing the INSERT...SELECT Operation

### Scenario

The INSERT...SELECT operation needs to be optimized if any of the following conditions is true:

- Many small files need to be queried.
- A few large files need to be queried.
- The INSERT...SELECT operation is performed by a non-spark user in Beeline/JDBCServer mode.

### Procedure

Optimize the INSERT...SELECT operation as follows:

- If the table to be created is the Hive table, set the storage type to Parquet. This enables INSERT...SELECT statements to be run faster.
- Perform the INSERT...SELECT operation as a spark-sql user or spark user (if in Beeline/JDBCServer mode). In this way, it is no longer necessary to change the file owner repeatedly, accelerating the execution of INSERT...SELECT statements.

#### NOTE

In Beeline/JDBCServer mode, the executor user is the same as the driver user. The driver user is a spark user because the driver is a part of JDBCServer service and started by a spark user. If the Beeline user is not a spark user, the file owner must be changed to the Beeline user (actual user) because the executor is unaware of the Beeline user.

- If many small files need to be queried, set spark.sql.files.maxPartitionBytes and spark.files.openCostInBytes to set the maximum size in bytes of partition and combine multiple small files in a partition to reduce file amount. This accelerates file renaming, ultimately enabling INSERT...SELECT statements to be run faster.

#### NOTE

The preceding optimizations are not a one-size-fits-all solution. In the following scenario, it still takes long to perform the INSERT...SELECT operation:  
The dynamic partitioned table contains many partitions.

## 26.7.2.5 Multiple JDBC Clients Concurrently Connecting to JDBCServer

### Scenario

Multiple clients can be connected to JDBCServer at the same time. However, if the number of concurrent tasks is too large, the default configuration of JDBCServer must be optimized to adapt to the scenario.

### Procedure

1. Set the fair scheduling policy of JDBCServer.  
The default scheduling policy of Spark is **FIFO**, which may cause a failure of short tasks in multi-task scenarios. Therefore, the fair scheduling policy must be used in multi-task scenarios to prevent task failure.
  - a. For details about how to configure Fair Scheduler in Spark, visit <http://spark.apache.org/docs/3.1.1/job-scheduling.html#scheduling-within-an-application>.
  - b. Configure Fair Scheduler on the JDBC client.
    - i. In the Beeline command line client or the code defined by JDBC, run the following statement:  
**PoolName** is a scheduling pool for Fair Scheduler.  
`SET spark.sql.thriftserver.scheduler.pool=PoolName;`
    - ii. Run the SQL command. The Spark task will be executed in the preceding scheduling pool.
2. Set the **BroadCastHashJoin** timeout interval.  
There is a timeout parameter of **BroadCastHashJoin**. The task query fails if the query period exceeds the preset timeout interval. In multi-task scenarios, the Spark task of BroadCastHashJoin may fail due to resource preemption. Therefore, it is necessary to modify the timeout interval in the **spark-defaults.conf** file of JDBCServer.

**Table 26-71** Parameter description

Parameter	Description	Default Value
spark.sql.broadcastTimeout	The timeout interval in the broadcast table of <b>BroadCastHashJoin</b> . If there are many concurrent tasks, set the parameter to a larger value or a negative number.	-1 (Numeral type. The actual value is 5 minutes.)

## 26.7.2.6 Optimizing Memory when Data Is Inserted into Dynamic Partitioned Tables

### Scenario

When SparkSQL inserts data to dynamic partitioned tables, the more partitions there are, the more HDFS files a single task generates and the more memory

metadata occupies. In this case, Garbage Collection (GC) is severe and Out of Memory (OOM) may occur.

Assume there are 10240 tasks and 2000 partitioned. Before the rename operation of HDFS files from a temporary directory to the target directory, there is about 29 GB FileStatus metadata.

## Procedure

Insert **distribute by** followed by partition fields into dynamic partition statements.

For example:

```
insert into table store_returns partition (sr_returned_date_sk) select
sr_return_time_sk,sr_item_sk,sr_customer_sk,sr_demo_sk,sr_hdemo_sk,sr_addr_sk,
sr_store_sk,sr_reason_sk,sr_ticket_number,sr_return_quantity,sr_return_amt,sr_return
n_tax,sr_return_amt_inc_tax,sr_fee,sr_return_ship_cost,sr_refunded_cash,sr_reversed
_charge,sr_store_credit,sr_net_loss,sr_returned_date_sk from $
{SOURCE}.store_returns distribute by sr_returned_date_sk;
```

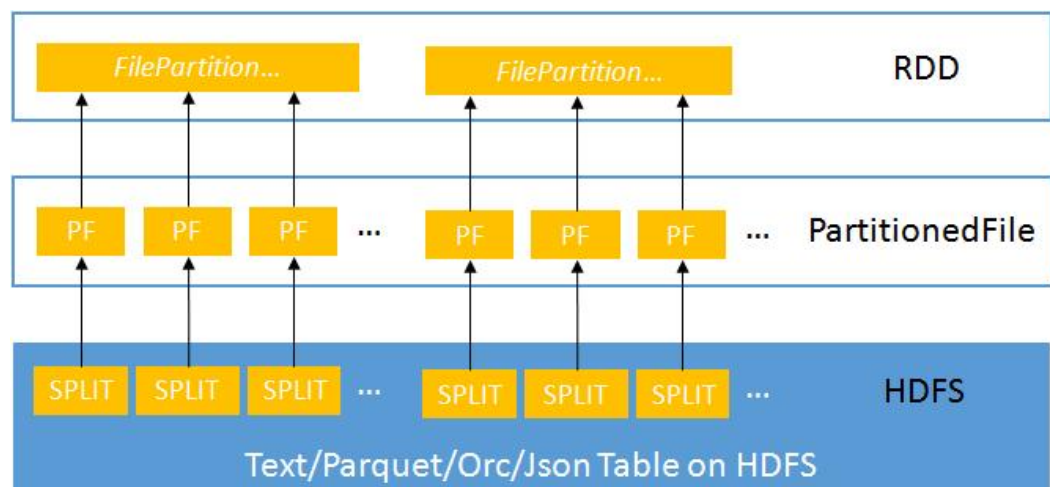
### 26.7.2.7 Optimizing Small Files

#### Scenario

A Spark SQL table may have many small files (far smaller than an HDFS block), each of which maps to a partition on the Spark by default. In other words, each small file is a task. In this way, Spark has to start many such tasks. If a shuffle operation is involved in the SQL logic, the number of hash buckets soars, severely hindering system performance.

In case of massive number of small files, when DataSource creates an RDD, it splits small files in the Spark SQL table to PartitionedFiles and then merges the PartitionedFiles to a partition to avoid generating too many hash buckets during the shuffle operation. See [Figure 26-8](#).

Figure 26-8 Merging small files



## Procedure

If you want to enable small file optimization, configure the **spark-defaults.conf** file on the Spark client.

**Table 26-72** Parameter description

Parameter	Description	Default Value
spark.sql.files.maxPartitionBytes	The maximum number of bytes that can be packed into a single partition when a file is read. Unit: byte	134217728 (128 MB)
spark.files.openCostInBytes	The estimated cost to open a file, measured by the number of bytes that can be scanned in the same time. This is used when putting multiple files into a partition. It is better to over estimate, then the partitions with small files will be faster than partitions with larger files.	4 MB

### 26.7.2.8 Optimizing the Aggregate Algorithms

#### Scenario

Spark SQL supports hash aggregate algorithm. Namely, use fast aggregate hashmap as cache to improve aggregate performance. The hashmap replaces the previous ColumnarBatch to avoid performance problems caused by the wide mode (multiple key or value fields) of an aggregate table.

#### Procedure

If you want to enable optimization of aggregate algorithm, configure following parameters in the **spark-defaults.conf** file on the Spark client.

**Table 26-73** Parameter description

Parameter	Description	Default Value
spark.sql.codegen.aggregate.map.twolevel.enabled	Specifies whether to enable aggregation algorithm optimization. <ul style="list-style-type: none"> <li><b>true:</b> Enable</li> <li><b>false:</b> Disable</li> </ul>	true

## 26.7.2.9 Optimizing Datasource Tables

### Scenario

Save the partition information about the datasource table to the Metastore and process partition information in the Metastore.

- Optimize the datasource tables, support syntax such as adding, deletion, and modification in the table based on partitions, improving compatibility with Hive.
- Support statements of partition tailoring and push down to the Metastore to filter unmatched partitions.

Example:

```
select count(*) from table where partCol=1; //partCol (partition column)
```

You need only to process data corresponding to partCol=1 when performing the TableScan operation in the physical plan.

### Procedure

If you want to enable Datasource table optimization, configure the **spark-defaults.conf** file on the Spark client.

**Table 26-74** Parameter description

Parameter	Description	Default Value
spark.sql.hive.manageFilesourcePartitions	Specifies whether to enable Metastore partition management (including datasource tables and converted Hive). <ul style="list-style-type: none"> <li>• <b>true</b> indicates enabling Metastore partition management. In this case, datasource tables are stored in Hive and Metastore is used to tailor partitions in query statements.</li> <li>• <b>false</b> indicates disabling Metastore partition management.</li> </ul>	true
spark.sql.hive.metastorePartitionPruning	Specifies whether to support pushing down predicate to Hive Metastore. <ul style="list-style-type: none"> <li>• <b>true</b> indicates supporting pushing down predicate to Hive Metastore. Only the predicate of Hive tables is supported.</li> <li>• <b>false</b> indicates not supporting pushing down predicate to Hive Metastore.</li> </ul>	true



Parameter	Description	Default Value
spark.sql.hive.filesourcePartitionFileCacheSize	<p>The cache size of the partition file metadata in the memory.</p> <p>All tables share a cache that can use up to specified num bytes for file metadata.</p> <p>This parameter is valid only when <b>spark.sql.hive.manageFilesourcePartitions</b> is set to <b>true</b>.</p>	250 * 1024 * 1024
spark.sql.hive.convertMetastoreOrc	<p>The processing approach of ORC tables.</p> <ul style="list-style-type: none"> <li>• <b>false</b>: Spark SQL uses Hive SerDe to process ORC tables.</li> <li>• <b>true</b>: Spark SQL uses the Spark built-in mechanism to process ORC tables.</li> </ul>	true

## 26.7.2.10 Merging CBO

### Scenario

Spark SQL supports rule-based optimization by default. However, the rule-based optimization cannot ensure that Spark selects the optimal query plan. Cost-Based Optimizer (CBO) is a technology that intelligently selects query plans for SQL statements. After CBO is enabled, the CBO optimizer performs a series of estimations based on the table and column statistics to select the optimal query plan.

### Procedure

Perform the following steps to enable CBO:

1. You need to run corresponding SQL commands to collect required table and column statistics.

SQL commands are as follows (to be chosen as required):

- Generate table-level statistics (table scanning):

***ANALYZE TABLE src COMPUTE STATISTICS***

This command generates **sizeInBytes** and **rowCount**.

When you use the ANALYZE statement to collect statistics, sizes of tables not from HDFS cannot be calculated.

- Generate table-level statistics (no table scanning):

***ANALYZE TABLE src COMPUTE STATISTICS NOSCAN***

This command generates only **sizeInBytes**. Compared with the originally generated **sizeInBytes** and **rowCount** if the **sizeInBytes** remains unchanged, **rowCount** (if any) reserves. Otherwise, **rowCount** is cleared.

- Generate column-level statistics:  
***ANALYZE TABLE src COMPUTE STATISTICS FOR COLUMNS a, b, c***  
This command generates column statistics and updates table statistics for consistency. Statistics of complicated data types (such as Seq and Map) and HiveStringType cannot be generated.
- Display statistics:  
***DESC FORMATTED src***  
This command displays *xxx* bytes and *xxx* rows in **Statistics** to indicate table-level statistics. You can also run the following command to display column statistics:  
***DESC FORMATTED src a***

**Limitation:** The current statistics collection does not support statistics for partition levels for partitioned tables.

2. Configure parameters in [Table 26-75](#) in the **spark-defaults.conf** file on the Spark client.

**Table 26-75** Parameter description

Parameter	Description	Default Value
spark.sql.cbo.enabled	The switch to enable or disable CBO. <ul style="list-style-type: none"> <li>• <b>true:</b> Enable</li> <li>• <b>false:</b> Disable</li> </ul> To enable this function, ensure that statistics of related tables and columns are generated.	false
spark.sql.cbo.joinReorder.enabled	Specifies whether to automatically adjust the sequence of consecutive inner joins by using CBO. <ul style="list-style-type: none"> <li>• <b>true:</b> Enable</li> <li>• <b>false:</b> Disable</li> </ul> To enable this function, ensure that statistics of related tables and columns are generated and CBO is enabled.	false
spark.sql.cbo.joinReorder.default.threshold	Specifies the threshold of the number of tables that the sequence of consecutive inner joins is automatically adjusted by CBO. If the threshold is exceeded, the sequence of joins is not adjusted.	12

## 26.7.2.11 Optimizing SQL Query of Data of Multiple Sources

### Scenario

This section describes how to enable or disable the query optimization for inter-source complex SQL.

### Procedure

- (Optional) Prepare for connecting to the MPPDB data source.  
If the data source to be connected is MPPDB, a class name conflict occurs because the MPPDB Driver file **gsjdbc4.jar** and the Spark JAR package **gsjdbc4-VXXXRXXXCXXSPCXXX.jar** contain the same class name. Therefore, before connecting to the MPPDB data source, perform the following steps:
  - a. Move **gsjdbc4-VXXXRXXXCXXSPCXXX.jar** from Spark. Spark running does not depend on this JAR file. Therefore, moving this JAR file to another directory (for example, the **/tmp** directory) will not affect Spark running.
    - i. Log in to the Spark server and remove **gsjdbc4-VXXXRXXXCXXSPCXXX.jar** from the **/\${BIGDATA\_HOME}/FusionInsight\_Spark2x\_8.1.0.1/install/FusionInsight-Spark2x-\*/spark/jars** directory.
    - ii. Log in to the Spark client host and move **gsjdbc4-VXXXRXXXCXXSPCXXX.jar** from the **/opt/client/Spark2x/spark/jars** directory.
  - b. Obtain the MPPDB Driver file **gsjdbc4.jar** from the MPPDB installation package and upload the file to the following directories:

#### NOTE

Obtain **gsjdbc4.jar** from **FusionInsight\_MPPDB\software\components\package\FusionInsight-MPPDB-xxx\package\Gauss-MPPDB-ALL-PACKAGES\GaussDB-xxx-REDHAT-xxx-Jdbc\jdbc**, the directory where the MPPDB installation package is stored.

- **/\${BIGDATA\_HOME}/FusionInsight\_Spark2x\_8.1.0.1/install/FusionInsight-Spark2x-\*/spark/jars** on the Spark server
  - **/opt/client/Spark2x/spark/jars** on the Spark client
- c. Update the **/user/spark2x/jars/8.1.0.1/spark-archive-2x.zip** package stored in the HDFS.

#### NOTE

The version 8.1.0.1 is used as an example. Replace it with the actual version number.

- i. Log in to the node where the client is installed as a client installation user. Run the following command to switch to the client installation directory, for example, **/opt/client**:  
**cd /opt/client**
- ii. Run the following command to configure environment variables:  
**source bigdata\_env**

- iii. If the cluster is in security mode, run the following command to get authenticated:  
**kinit Component service user**
  - iv. Run the following commands to create the temporary file **./tmp**, obtain **spark-archive-2x.zip** from HDFS, and decompress it to the **tmp** directory:  
**mkdir tmp**  
**hdfs dfs -get /user/spark2x/jars/8.1.0.1/spark-archive-2x.zip ./**  
**unzip spark-archive-2x.zip -d ./tmp**
  - v. Switch to the **tmp** directory, delete the **gsjdbc4-VXXXXXXCXXSPCXXX.jar** file, upload the MPPDB Driver file **gsjdbc4.jar** to the **tmp** directory, and run the following command to compress the file again:  
**zip -r spark-archive-2x.zip \*.jar**
  - vi. Delete **spark-archive-2x.zip** from the HDFS and update the **spark-archive-2x.zip** package generated in **c.v** to the **/user/spark2x/jars/8.1.0.1/** directory in the HDFS.  
**hdfs dfs -rm /user/spark2x/jars/8.1.0.1/spark-archive-2x.zip**  
**hdfs dfs -put ./spark-archive-2x.zip /user/spark2x/jars/8.1.0.1**
- d. Restart the Spark service. After the Spark service is restarted, restart the Spark client.
- Enable the optimization function.  
For all modules that support query pushdown, you can run the **SET** command on the **spark-beeline** client to enable the cross-source query optimization function. By default, the function is disabled.  
Pushdown configurations can be performed in dimensions of global, data sources, and tables. Commands are as follows:
    - Global (valid for all data sources):  
**SET spark.sql.datasource.jdbc = project,aggregate,orderby-limit**
    - Data sources:  
**SET spark.sql.datasource.\${url} = project,aggregate,orderby-limit**
    - Tables:  
**SET spark.sql.datasource.\${url}.\${table} = project,aggregate,orderby-limit**

When you run the **SET** command to configure preceding parameters, you are allowed to specify multiple pushdown modules and separate them by commas. The following table lists parameters of corresponding pushdown modules.

**Table 26-76** Parameters of modules

Module	Parameter Value in the SET Command
project	project

Module	Parameter Value in the SET Command
aggregate	aggregate
order by, limit over project or aggregate	orderby-limit

The following is a statement for creating an external table of MySQL:

```
create table if not exists pdmysql using org.apache.spark.sql.jdbc
options(driver "com.mysql.jdbc.Driver", url "jdbc:mysql://ip2:3306/test",
user "hive", password "xxx", dbtable "mysqldata");
```

In the preceding statement:

- `${url}` = `jdbc:mysql://ip2:3306/test`
- `${table}` = `mysqldata`

 **NOTE**

- On the right of the equal sign (=) is the operators (separated by commas) to be enabled by pushdown.
- Priority: table > data source > global. If the table switch is set, the global switch of the data source is invalid for the table. If a data source switch is set, the global switch is invalid for the data source.
- The equal sign (=) is not allowed in URL. Equal signs (=) are automatically deleted in the SET clause.
- After multiple SET operations, results with different keys will not overwrite each other.

- Add functions that support query pushdown.

In addition to query pushdown of mathematical, time, and string functions such as `abs()`, `month()`, and `length()`, you can run the **SET** command to add a data source that supports query pushdown. Run the following command on the Spark-beeline client:

```
SET spark.sql.datasource.${datasource}.functions = fun1,fun2
```

- Reset the configuration set by the **SET** command.

Currently, you can only run the **RESET** command on the **spark-beeline** client to cancel all **SET** content. After running the **RESET** command, all values in the **SET** command will be cleared. Exercise caution when performing this operation.

The **SET** command is valid in the current session on the client. After the client is shut down, the content in the **SET** command turns invalid.

Alternatively, change the value of `spark.sql.locale.support` in the `spark-defaults.conf` file to **true**.

## Precautions

Only MySQL, MPPDB, Hive, oracle, and PostgreSQL data sources are supported.

## 26.7.2.12 SQL Optimization for Multi-level Nesting and Hybrid Join

### Scenario

This section describes the optimization suggestions for SQL statements in multi-level nesting and hybrid join scenarios.

### Prerequisites

The following provides an example of complex query statements:

```
select
s_name,
count(1) as numwait
from (
select s_name from (
select
s_name,
t2.l_orderkey,
l_suppkey,
count_suppkey,
max_suppkey
from
test2 t2 right outer join (
select
s_name,
l_orderkey,
l_suppkey from (
select
s_name,
t1.l_orderkey,
l_suppkey,
count_suppkey,
max_suppkey
from
test1 t1 join (
select
s_name,
l_orderkey,
l_suppkey
from
orders o join (
select
s_name,
l_orderkey,
l_suppkey
from
nation n join supplier s
on
s.s_nationkey = n.n_nationkey
and n.n_name = 'SAUDI ARABIA'
join lineitem l
on
s.s_suppkey = l.l_suppkey
where
l.l_receiptdate > l.l_commitdate
and l.l_orderkey is not null
) l1 on o.o_orderkey = l1.l_orderkey and o.o_orderstatus = 'F'
) l2 on l2.l_orderkey = t1.l_orderkey
) a
where
(count_suppkey > 1)
or ((count_suppkey=1)
and (l_suppkey <> max_suppkey))
) l3 on l3.l_orderkey = t2.l_orderkey
) b
where
```

```
(count_suppkey is null)
or ((count_suppkey=1)
and (L_suppkey = max_suppkey))
) c
group by
s_name
order by
numwait desc,
s_name
limit 100;
```

## Procedure

### Step 1 Analyze services.

Analyze business to determine whether SQL statements can be simplified through measures, for example, by combining tables to reduce the number of nesting levels layers and join times.

### Step 2 If the SQL statements cannot be simplified, configure the driver memory.

- Use **spark-submit** or **spark-sql** to run SQL statements and go to [Step 3](#).
- Use **spark-beeline** to run SQL statements and go to [Step 4](#).

### Step 3 During execution of SQL statements, specify the **driver-memory** parameter. An example of SQL statements is as follows:

```
/spark-sql --master=local[4] --driver-memory=512M -f /tpch.sql
```

### Step 4 Before you run SQL statements, change the memory size as the MRS cluster administrator.

1. Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x**. On the page that is displayed, click the **Configuration** tab.
2. Click the **All Configurations** sub-tab and search for **SPARK\_DRIVER\_MEMORY**.
3. Set the parameter to a larger value to increase the memory size. The value must be an integer, and the unit must be MB or GB. For example, enter **512 MB**.

----End

## Related Information

In the event of insufficient DRIVER memory, the following error may be displayed during the query:

```
2018-02-11 09:13:14,683 | WARN | Executor task launch worker for task 5 | Calling spill() on
RowBasedKeyValueBatch. Will not spill but return 0. |
org.apache.spark.sql.catalyst.expressions.RowBasedKeyValueBatch.spill(RowBasedKeyValueBatch.java:173)
2018-02-11 09:13:14,682 | WARN | Executor task launch worker for task 3 | Calling spill() on
RowBasedKeyValueBatch. Will not spill but return 0. |
org.apache.spark.sql.catalyst.expressions.RowBasedKeyValueBatch.spill(RowBasedKeyValueBatch.java:173)
2018-02-11 09:13:14,704 | ERROR | Executor task launch worker for task 2 | Exception in task 2.0 in stage
1.0 (TID 2) | org.apache.spark.internal.Logging$class.logError(Logging.scala:91)
java.lang.OutOfMemoryError: Unable to acquire 262144 bytes of memory, got 0
at org.apache.spark.memory.MemoryConsumer.allocateArray(MemoryConsumer.java:100)
at org.apache.spark.unsafe.map.BytesToBytesMap.allocate(BytesToBytesMap.java:791)
at org.apache.spark.unsafe.map.BytesToBytesMap.<init>(BytesToBytesMap.java:208)
at org.apache.spark.unsafe.map.BytesToBytesMap.<init>(BytesToBytesMap.java:223)
```

```

 at
org.apache.spark.sql.execution.UnsafeFixedWidthAggregationMap.<init>(UnsafeFixedWidthAggregationMap.j
ava:104)
 at
org.apache.spark.sql.execution.aggregate.HashAggregateExec.createHashMap(HashAggregateExec.scala:307)
 at org.apache.spark.sql.catalyst.expressions.GeneratedClass
$GeneratedIterator.agg_doAggregateWithKeys$(Unknown Source)
 at org.apache.spark.sql.catalyst.expressions.GeneratedClass$GeneratedIterator.processNext(Unknown
Source)
 at org.apache.spark.sql.execution.BufferedRowIterator.hasNext(BufferedRowIterator.java:43)
 at org.apache.spark.sql.execution.WholeStageCodegenExec$$anonfun$8$$anon
$1.hasNext(WholeStageCodegenExec.scala:381)
 at scala.collection.Iterator$$anon$11.hasNext(Iterator.scala:408)
 at
org.apache.spark.shuffle.sort.BypassMergeSortShuffleWriter.write(BypassMergeSortShuffleWriter.java:126)
 at org.apache.spark.scheduler.ShuffleMapTask.runTask(ShuffleMapTask.scala:96)
 at org.apache.spark.scheduler.ShuffleMapTask.runTask(ShuffleMapTask.scala:53)
 at org.apache.spark.scheduler.Task.run(Task.scala:99)
 at org.apache.spark.executor.Executor$TaskRunner.run(Executor.scala:325)
 at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)
 at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624)
 at java.lang.Thread.run(Thread.java:748)

```

## 26.7.3 Spark Streaming Tuning

### Scenario

Streaming is a mini-batch streaming processing framework that features second-level delay and high throughput. To optimize Streaming is to improve its throughput while maintaining second-level delay so that more data can be processed per unit time.

#### NOTE

This section applies to the scenario where the input data source is Kafka.

### Procedure

A simple streaming processing system consists of a data source, a receiver, and a processor. The data source is Kafka, the receiver is the Kafka data source receiver of Streaming, and the processor is Streaming.

Streaming optimization is to optimize the performance of the three components.

- **Data source optimization**

In actual application scenarios, the data source stores the data in the local disks to ensure the error tolerance of the data. However, the calculation results of the Streaming are stored in the memory, and the data source may become the largest bottleneck of the streaming system.

Kafka can be optimized from the following aspects:

- Use Kafka-0.8.2 or later version that allows you to use new Producer APIs in asynchronous mode.
- Configure multiple Broker directories, multiple I/O threads, and a proper number of partitions for a topic.

- **Receiver optimization**

Streaming has multiple data source receivers, such as Kafka, Flume, MQTT, and ZeroMQ. Kafka has the most receiver types and is the most mature receiver.



Kafka provides three types of receiver APIs:

- KafkaReceiver directly receives Kafka data. If the process is abnormal, data may be lost.
- ReliableKafkaReceiver receives data displacement through ZooKeeper records.
- DirectKafka reads data from each partition of Kafka through the RDD, ensuring high reliability.

According to the implementation mechanism and test results, DirectKafka provides better performance than the other two APIs. Therefore, the DirectKafka API is recommended to implement the receiver.

Kafka receivers function as Kafka consumers.

- **Processor optimization**

The bottom layer of Spark Streaming is executed by Spark. Therefore, most optimization measures for Spark can also be applied to Spark Streaming. The following is an example:

- Data serialization
- Memory configuration
- Configuring DOP
- Using the external shuffle service to improve performance

 **NOTE**

Higher performance of Spark Streaming indicates lower overall reliability. Examples:

If `spark.streaming.receiver.writeAheadLog.enable` is set to **false**, disk I/Os are reduced and performance is improved. However, because WAL is disabled, data is lost during fault recovery.

Therefore, do not disable configuration items that ensure data reliability in production environments during Spark Streaming tuning.

- **Log archive optimization**

The `spark.eventLog.group.size` parameter is used to group **JobHistory** logs of an application based on the specified number of jobs. Each group creates a file recording log to prevent **JobHistory** reading failures caused by an oversized log generated during the long-term running of the application. If this parameter is set to **0**, logs are not grouped.

Most Spark Streaming jobs are small jobs and are generated at a high speed. As a result, frequent grouping is performed and a large number of small log files are generated, consuming disk I/O resources. You are advised to increase the parameter value to, for example, **1000** or greater.

## 26.8 Common Issues About Spark2x

### 26.8.1 Spark Core

### 26.8.1.1 How Do I View Aggregated Spark Application Logs?

#### Question

How do I view the aggregated container logs on the page when the log aggregation function is enabled on YARN?

#### Answer

For details, see [Viewing Aggregated Container Logs on the Web UI](#).

### 26.8.1.2 Why Cannot Exit the Driver Process?

#### Question

Why cannot exit the Driver process after running the **yarn application -kill applicationID** command to stop the Spark Streaming application?

#### Answer

Running the **yarn application -kill applicationID** command can only stop the SparkContext corresponding to Spark Streaming application, but cannot exit the current Driver process. If there are other permanent threads in the Driver process (for example, the spark shell is continually checking command input or Spark Streaming is continually reading data from data source), the Driver process will not be killed when the SparkContext is stopped. To exit the Driver process, you are advised to run the **kill -9 pid** command to kill the current Driver process by hand.

### 26.8.1.3 Why Does FetchFailedException Occur When the Network Connection Is Timed out

#### Question

On a large cluster of 380 nodes, run the ScalaSort test case in the HiBench test that runs the 29T data, and configure Executor as **--executor-cores 4**. The following abnormality is displayed:

```
org.apache.spark.shuffle.FetchFailedException: Failed to connect to /192.168.114.12:23242
 at
 org.apache.spark.storage.ShuffleBlockFetcherIterator.throwFetchFailedException(ShuffleBlockFetcherIterator.scala:321)
 at org.apache.spark.storage.ShuffleBlockFetcherIterator.next(ShuffleBlockFetcherIterator.scala:306)
 at org.apache.spark.storage.ShuffleBlockFetcherIterator.next(ShuffleBlockFetcherIterator.scala:51)
 at scala.collection.Iterator$$anon$11.next(Iterator.scala:328)
 at scala.collection.Iterator$$anon$13.hasNext(Iterator.scala:371)
 at scala.collection.Iterator$$anon$11.hasNext(Iterator.scala:327)
 at org.apache.spark.util.CompletionIterator.hasNext(CompletionIterator.scala:32)
 at org.apache.spark.InterruptibleIterator.hasNext(InterruptibleIterator.scala:39)
 at org.apache.spark.util.collection.ExternalSorter.insertAll(ExternalSorter.scala:217)
 at org.apache.spark.shuffle.hash.HashShuffleReader.read(HashShuffleReader.scala:102)
 at org.apache.spark.rdd.ShuffledRDD.compute(ShuffledRDD.scala:90)
 at org.apache.spark.rdd.RDD.computeOrReadCheckpoint(RDD.scala:301)
 at org.apache.spark.rdd.RDD.iterator(RDD.scala:265)
 at org.apache.spark.rdd.MapPartitionsRDD.compute(MapPartitionsRDD.scala:38)
 at org.apache.spark.rdd.RDD.computeOrReadCheckpoint(RDD.scala:301)
 at org.apache.spark.rdd.RDD.iterator(RDD.scala:265)
 at org.apache.spark.rdd.MapPartitionsRDD.compute(MapPartitionsRDD.scala:38)
```

```

at org.apache.spark.rdd.RDD.computeOrReadCheckpoint(RDD.scala:301)
at org.apache.spark.rdd.RDD.iterator(RDD.scala:265)
at org.apache.spark.rdd.UnionRDD.compute(UnionRDD.scala:87)
at org.apache.spark.rdd.RDD.computeOrReadCheckpoint(RDD.scala:301)
at org.apache.spark.rdd.RDD.iterator(RDD.scala:265)
at org.apache.spark.scheduler.ShuffleMapTask.runTask(ShuffleMapTask.scala:73)
at org.apache.spark.scheduler.ShuffleMapTask.runTask(ShuffleMapTask.scala:41)
at org.apache.spark.scheduler.Task.run(Task.scala:87)
at org.apache.spark.executor.Executor$TaskRunner.run(Executor.scala:213)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
at java.lang.Thread.run(Thread.java:745)
Caused by: java.io.IOException: Failed to connect to /192.168.114.12:23242
at org.apache.spark.network.client.TransportClientFactory.createClient(TransportClientFactory.java:214)
at org.apache.spark.network.client.TransportClientFactory.createClient(TransportClientFactory.java:167)
at org.apache.spark.network.netty.NettyBlockTransferService$$anon
$1.createAndStart(NettyBlockTransferService.scala:91)
at
org.apache.spark.network.shuffle.RetryingBlockFetcher.fetchAllOutstanding(RetryingBlockFetcher.java:140)
at org.apache.spark.network.shuffle.RetryingBlockFetcher.access$200(RetryingBlockFetcher.java:43)
at org.apache.spark.network.shuffle.RetryingBlockFetcher$1.run(RetryingBlockFetcher.java:170)
at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511)
at java.util.concurrent.FutureTask.run(FutureTask.java:266)
... 3 more
Caused by: java.net.ConnectException: Connection timed out: /192.168.114.12:23242
at sun.nio.ch.SocketChannelImpl.checkConnect(Native Method)
at sun.nio.ch.SocketChannelImpl.finishConnect(SocketChannelImpl.java:717)
at io.netty.channel.socket.nio.NioSocketChannel.doFinishConnect(NioSocketChannel.java:224)
at io.netty.channel.nio.AbstractNioChannel
$AbstractNioUnsafe.finishConnect(AbstractNioChannel.java:289)
at io.netty.channel.nio.NioEventLoop.processSelectedKey(NioEventLoop.java:528)
at io.netty.channel.nio.NioEventLoop.processSelectedKeysOptimized(NioEventLoop.java:468)
at io.netty.channel.nio.NioEventLoop.processSelectedKeys(NioEventLoop.java:382)
at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:354)
at io.netty.util.concurrent.SingleThreadEventExecutor$2.run(SingleThreadEventExecutor.java:111)
... 1 more

```

## Answer

When an application is run, configure the Executor parameter as **--executor-cores 4**. The degree of parallelism (DOP) is high in a single process, resulting in that the IO is highly occupied and the task works slowly.

```
16/02/26 10:04:53 INFO TaskSetManager: Finished task 2139.0 in stage 1.0 (TID 151149) in 376455 ms on 10-196-115-2 (694/153378)
```

Because running a single task takes more than 6 minutes. The network connection is timed out and the running task fails.

Set the number of cores as 1, which is **--executor-cores 1**. A task is executed smoothly in proper time (within 15s).

```
16/02/29 02:24:46 INFO TaskSetManager: Finished task 59564.0 in stage 1.0 (TID 208574) in 15088 ms on 10-196-115-6 (59515/153378)
```

Therefore, to process the task of network connection timed out and avoid such error, you can reduce the core number of a single Executor.

### 26.8.1.4 How to Configure Event Queue Size If Event Queue Overflows?

#### Question

How to configure the event queue size if the following Driver log information is displayed indicating that the event queue overflows?

- **Common applications**  
Dropping SparkListenerEvent because no remaining room in event queue.  
This likely means one of the SparkListeners is too slow and cannot keep up with the rate at which tasks are being started by the scheduler.
- **Spark Streaming applications**  
Dropping StreamingListenerEvent because no remaining room in event queue.  
This likely means one of the StreamingListeners is too slow and cannot keep up with the rate at which events are being started by the scheduler.

## Answer

1. Stop the application. Set the configuration option **spark.event.listener.logEnable** in the Spark configuration file **spark-defaults.conf** to **true**. And set the configuration option **spark.eventQueue.size** to **1000W**. If you need to control the logging rate (in milliseconds), also change the value of the configuration option **spark.event.listener.logRate**.  
By default, the logging rate is 1000 ms, which means that one log is printed out every 1000 ms.
2. Start the application.  
The following log information is displayed, including the event consumption rate, event production rate, and **MaxSize** (maximum size of messages in the queue).  
INFO LiveListenerBus: [SparkListenerBus]:16044 events are consumed in 5000 ms.  
INFO LiveListenerBus: [SparkListenerBus]:51381 events are produced in 5000 ms, eventQueue still has 86417 events, MaxSize: 171764.
3. Change the value of the configuration option **spark.eventQueue.size** in the Spark configuration file **spark-defaults.conf** based on the **MaxSize** in the log information.  
For example, if **MaxSize** is 250000, the appropriate message queue size is 300000.

### 26.8.1.5 What Can I Do If the `getApplicationReport` Exception Is Recorded in Logs During Spark Application Execution and the Application Does Not Exit for a Long Time?

#### Question

During Spark application execution, if the driver fails to connect to ResourceManager, the following error is reported and it does not exit for a long time. What can I do?

```
16/04/23 15:31:44 INFO RetryInvocationHandler: Exception while invoking getApplicationReport of class ApplicationClientProtocolPBClientImpl over 37 after 1 fail over attempts. Trying to fail over after sleeping for 44160ms.
java.net.ConnectException: Call From vm1/192.168.39.30 to vm1:8032 failed on connection exception: java.net.ConnectException: Connection refused; For more details see: http://wiki.apache.org/hadoop/ConnectionRefused
```

#### Answer

In Spark, there is a scheduled thread that listens to the status of ApplicationMaster by connecting to ResourceManager. The connection to the ResourceManager times out. As a result, the preceding error is reported and the

system keeps trying to connect to the ResourceManager. In the ResourceManager, the number of retry times is limited. By default, the number of retry times is 30 and the retry interval is about 30 seconds. The preceding error is reported during each retry. The driver exits only after the number of times is exceeded.

**Table 26-77** describes the retry-related configuration items in the ResourceManager.

**Table 26-77** Parameter description

Parameter	Description	Default Value
yarn.resourcemanager.connect.max-wait.ms	Maximum waiting time for connecting to the ResourceManager.	900000
yarn.resourcemanager.connect.retry-interval.ms	Interval for reconnecting to the ResourceManager.	30000

Number of retries (**yarn.resourcemanager.connect.max-wait.ms/ yarn.resourcemanager.connect.retry-interval.ms**) = Maximum waiting time for connecting to the ResourceManager/Interval for reconnecting to the ResourceManager

On the Spark client, modify the **conf/yarn-site.xml** file to add and configure **yarn.resourcemanager.connect.max-wait.ms** and **yarn.resourcemanager.connect.retry-interval.ms**. In this way, the number of retry times can be changed, and the Spark application can exit in advance.

### 26.8.1.6 What Can I Do If "Connection to ip:port has been quiet for xxx ms while there are outstanding requests" Is Reported When Spark Executes an Application and the Application Ends?

#### Question

When Spark executes an application, an error similar to the following is reported and the application ends. What can I do?

```
2016-04-20 10:42:00,557 | ERROR | [shuffle-server-2] | Connection to 10-91-8-208/10.18.0.115:57959 has
been quiet for 180000 ms while there are outstanding requests. Assuming connection is dead; please adju
st spark.network.timeout if this is wrong. |
org.apache.spark.network.server.TransportChannelHandler.userEventTriggered(TransportChannelHandler.java:
128)
2016-04-20 10:42:00,558 | ERROR | [shuffle-server-2] | Still have 1 requests outstanding when connection
from 10-91-8-208/10.18.0.115:57959 is closed | org.apache.spark.network.client.TransportResponseHandl
er.channelUnregistered(TransportResponseHandler.java:102)
2016-04-20 10:42:00,562 | WARN | [yarn-scheduler-ask-am-thread-pool-160] | Error sending message
[message = DoShuffleClean(application_1459995017785_0108,319)] in 1 attempts |
org.apache.spark.Logging$class
s.logWarning(Logging.scala:92)
java.io.IOException: Connection from 10-91-8-208/10.18.0.115:57959 closed
 at
 org.apache.spark.network.client.TransportResponseHandler.channelUnregistered(TransportResponseHandler.j
ava:104)
 at
```

```

org.apache.spark.network.server.TransportChannelHandler.channelUnregistered(TransportChannelHandler.java:94)
 at
io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
 at
io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
 at
io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
 at
io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
 at
io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
 at
io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
 at
io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
 at
io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
 at
io.netty.channel.ChannelInboundHandlerAdapter.channelUnregistered(ChannelInboundHandlerAdapter.java:53)
 at
io.netty.channel.AbstractChannelHandlerContext.invokeChannelUnregistered(AbstractChannelHandlerContext.java:158)
 at
io.netty.channel.AbstractChannelHandlerContext.fireChannelUnregistered(AbstractChannelHandlerContext.java:144)
 at io.netty.channel.DefaultChannelPipeline.fireChannelUnregistered(DefaultChannelPipeline.java:739)
 at io.netty.channel.AbstractChannel$AbstractUnsafe$8.run(AbstractChannel.java:659)
 at io.netty.util.concurrent.SingleThreadEventExecutor.runAllTasks(SingleThreadEventExecutor.java:357)
 at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:357)
 at io.netty.util.concurrent.SingleThreadEventExecutor$2.run(SingleThreadEventExecutor.java:111)
 at java.lang.Thread.run(Thread.java:745)
2016-04-20 10:42:00,573 | INFO | [dispatcher-event-loop-14] | Starting task 177.0 in stage 1492.0 (TID 1996351, linux-254, PROCESS_LOCAL, 2106 bytes) | org.apache.spark.Logging$class.logInfo(Logging.scala:59)
2016-04-20 10:42:00,574 | INFO | [task-result-getter-0] | Finished task 85.0 in stage 1492.0 (TID 1996259) in 191336 ms on linux-254 (106/3000) | org.apache.spark.Logging$class.logInfo(Logging.scala:59)
2016-04-20 10:42:00,811 | ERROR | [Yarn application state monitor] | Yarn application has already exited with state FINISHED! | org.apache.spark.Logging$class.logError(Logging.scala:75)

```

## Answer

Symptom: The value of **spark.rpc.io.connectionTimeout** is less than the value of **spark.rpc.askTimeout**. In full GC or network delay scenarios, when the channel reaches the expiration time and still receives no response, the channel is terminated. When detecting that the channel is terminated, the AM considers the driver as disconnected, and the entire application is stopped.

Solution: Set the parameter in the **spark-defaults.conf** file on the Spark client by running the **set** command. During parameter configuration, ensure that the channel expiration time (**spark.rpc.io.connectionTimeout**) is greater than or equal to the RPC response timeout (**spark.rpc.askTimeout**).

**Table 26-78** Parameter description

Parameter	Description	Default Value
spark.rpc.askTimeout	RPC response timeout. If this parameter is not set, the value of <b>spark.network.timeout</b> is used by default.	120s

### 26.8.1.7 Why Do Executors Fail to be Removed After the NodeManager Is Shut Down?

#### Question

If the NodeManager is shut down with the Executor dynamic allocation enabled, the Executors on the node where the NodeManager is shut down fail to be removed from the driver page after the idle time expires.

#### Answer

When the ResourceManager detects that the NodeManager is shut down, the driver has requested to kill Executors due to idle time expiry. However, the Executors cannot actually be killed because the NodeManager is shut down. The driver cannot detect the LOST events of these Executors and does not remove Executors from its Executor list. Therefore, the Executors are not removed from the driver page. This phenomenon is normal after the YARN NodeManager is shut down. The Executors will be removed after the NodeManager restarts.

### 26.8.1.8 What Can I Do If the Message "Password cannot be null if SASL is enabled" Is Displayed?

#### Question

ExternalShuffle is enabled for the application that runs Spark. Task loss occurs in the application because the message "java.lang.NullPointerException: Password cannot be null if SASL is enabled" is displayed. The following shows some key logs:

```
2016-05-13 12:05:27.093 | WARN | [task-result-getter-2] | Lost task 98.0 in stage 22.1 (TID 193603, linux-173, 2): FetchFailed(BlockManagerId(13, 172.168.100.13, 27337)), org.apache.spark.shuffle.FetchFailedException: java.lang.NullPointerException: Password cannot be null if SASL is enabled
 at org.spark-project.guava.base.Preconditions.checkNotNull(Preconditions.java:208)
 at org.apache.spark.network.sasl.SparkSaslServer.encodePassword(SparkSaslServer.java:196)
 at org.apache.spark.network.sasl.SparkSaslServer$DigestCallbackHandler.handle(SparkSaslServer.java:166)
 at com.sun.security.sasl.digest.DigestMD5Server.validateClientResponse(DigestMD5Server.java:589)
 at com.sun.security.sasl.digest.DigestMD5Server.evaluateResponse(DigestMD5Server.java:244)
 at org.apache.spark.network.sasl.SparkSaslServer.response(SparkSaslServer.java:119)
 at org.apache.spark.network.sasl.SaslRpcHandler.receive(SaslRpcHandler.java:100)
 at org.apache.spark.network.server.TransportRequestHandler.processRpcRequest(TransportRequestHandler.java:128)
 at org.apache.spark.network.server.TransportRequestHandler.handle(TransportRequestHandler.java:99)
 at org.apache.spark.network.server.TransportChannelHandler.channelRead0(TransportChannelHandler.java:104)
```

#### Answer

The cause is that NodeManager restarts. When ExternalShuffle is used, Spark uses NodeManager to transmit shuffle data. Therefore, the memory of NodeManager may be seriously insufficient.

In the FusionInsight of the current version, the default memory of NodeManager is only 1 GB. When the data volume of Spark tasks is large (greater than 1 TB),

the memory is severely insufficient and the message response is slow. As a result, the FusionInsight health check determines that the NodeManager process exits and forcibly restarts the NodeManager, causing the preceding problem.

Solution

Adjust the memory of the NodeManager. If the data volume is large (greater than 1 TB), the memory of NodeManager must be greater than 4 GB.

### 26.8.1.9 What Should I Do If the Message "Failed to CREATE\_FILE" Is Displayed in the Restarted Tasks When Data Is Inserted Into the Dynamic Partition Table?

#### Question

When inserting data into the dynamic partition table, a large number of shuffle files are damaged due to the disk disconnection, node error, and the like. In this case, why the message **Failed to CREATE\_FILE** is displayed in the restarted tasks?

```
2016-06-25 15:11:31,323 | ERROR | [Executor task launch worker-0] | Exception in task 15.0 in stage 10.1 (TID 1258) | org.apache.spark.Logging$class.logError(Logging.scala:96)
org.apache.hadoop.hive.ql.metadata.HiveException:
org.apache.hadoop.ipc.RemoteException(org.apache.hadoop.hdfs.protocol.AlreadyBeingCreatedException):
Failed to CREATE_FILE /user/hive/warehouse/testdb.db/web_sales/.hive-staging_hive_2016-06-25_15-09-16_999_8137121701603617850-1/-ext-10000/_temporary/0/_temporary/attempt_201606251509_0010_m_000015_0/ws_sold_date=1999-12-17/part-00015 for
DFSClient_attempt_2016
06251509_0010_m_000015_0_353134803_151 on 10.1.1.5 because this file lease is currently owned by
DFSClient_attempt_201606251509_0010_m_000015_0_-848353830_156 on 10.1.1.6
```

#### Answer

The last step of inserting data into the dynamic partition table is to read shuffle files and then write the data to the mapped partition files.

After a large number of shuffle files are damaged, a large number of tasks fail, causing the restart of jobs. Before the restart of jobs, Spark closes the handles that write table partition files. However, the HDFS cannot process the scenario of batch tasks closing handles. After tasks restart next time, the handles are not released in a timely manner on the NameNode. As a result, the message **Failed to CREATE\_FILE** is displayed.

This error only occurs when a large number of shuffle files are damaged. The tasks will restart after the error occurs and the restart can be completed within milliseconds.

### 26.8.1.10 Why Tasks Fail When Hash Shuffle Is Used?

#### Question

When Hash shuffle is used to run a job that consists of 1000000 map tasks x 100000 reduce tasks, run logs report many message failures and Executor heartbeat timeout, leading to task failures. Why does this happen?



## Answer

During the shuffle process, Hash shuffle just writes the data of different reduce partitions to their respective disk files according to hash results without sorting the data.

If there are many reduce partitions, a large number of disk files will be generated. In your case,  $10^{11}$  shuffle files, that is,  $1000000 * 100000$  shuffle files, will be generated. The sheer number of disk files will have a great impact on the file read and write performance. In addition, the operations such as sorting and compressing will consume a large amount of temporary memory space because a large number of file handles are open, presenting great challenges to memory management and garbage collection and incurring the possibility that the Executor fails to respond to Driver.

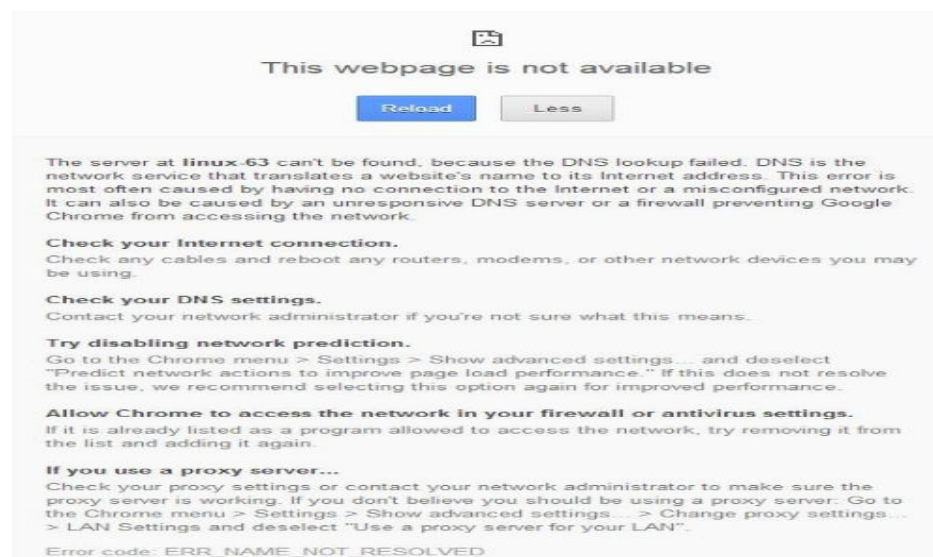
Sort shuffle, instead of Hash shuffle, is recommended to run a job.

### 26.8.1.11 What Can I Do If the Error Message "DNS query failed" Is Displayed When I Access the Aggregated Logs Page of Spark Applications?

#### Question

When the `http(s)://<spark ip>:<spark port>` mode is used to access the Spark JobHistory page, if the displayed Spark JobHistory page is not the page of FusionInsight Manager (the URL of FusionInsight Manager is similar to `https://<oms ip>:20026/Spark2x/JobHistory2x/xx/`), click an application and click **AggregatedLogs**, click the logs of an executor to be viewed. An error message in [Figure 26-9](#) is displayed.

Figure 26-9 DNS query failure



#### Answer

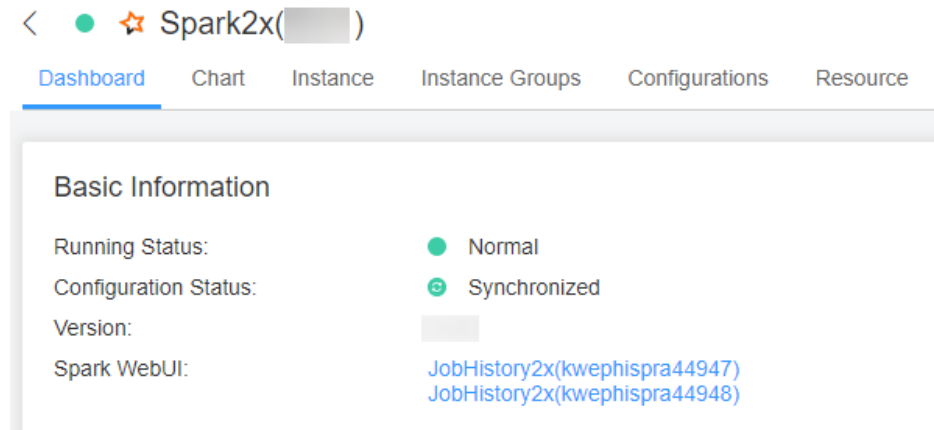
**Cause:** The domain name is not added to the `hosts` file of the Windows OS in the pop-up URL (for example, `https://<hostname>:20026/Spark2x/JobHistory2x/xx/`)

`history/application_xxx/jobs/`). As a result, the DNS query fails and the web page cannot be displayed.

**Solution:**

- You are advised to visit the **Spark JobHistory** page using FusionInsight Manager. Click the links in the blue box in **Figure 26-10**.

**Figure 26-10** Spark2x page of FusionInsight Manager



- If you do not want to access the **Spark JobHistory** page using the FusionInsight Manager, change `<hostname>` in the URL to the IP address or add the domain name to the `hosts` file of the Windows OS.

### 26.8.1.12 What Can I Do If Shuffle Fetch Fails Due to the "Timeout Waiting for Task" Exception?

#### Question

When I execute a 100 TB TPC-DS test suite in the JDBCServer mode, the "Timeout waiting for task" is displayed. As a result, shuffle fetch fails, the stage keeps retrying, and the task cannot be completed properly. What can I do?

#### Answer

The ShuffleService function is used in JDBCServer mode. In the reduce phase, all executors obtain data from NodeManager. When the data volume reaches a level (more than 10 TB), the NodeManager may reach the bottleneck (ShuffleService is in the NodeManager process). As a result, some tasks for obtaining data time out. Therefore, the problem occurs.

You are advised to disable ShuffleService for Spark tasks whose data volume is greater than 10 TB. That is, set `spark.shuffle.service.enabled` in the `Spark-defaults.conf` configuration file to `false`.

### 26.8.1.13 Why Does the Stage Retry due to the Crash of the Executor?

#### Question

When I run Spark tasks with a large data volume, for example, 100 TB TPCDS test suite, why does the Stage retry due to Executor loss sometimes? The message

"Executor 532 is lost rpc with driver, but is still alive, going to kill it" is displayed, indicating that the loss of the Executor is caused by a JVM crash.

The log of the key JVM crash is as follows:

```

A fatal error has been detected by the Java Runtime Environment:

Internal Error (sharedRuntime.cpp:834), pid=241075, tid=140476258551552
fatal error: exception happened outside interpreter, nmethods and vtable stubs at pc
0x00007fcda9eb8eb1
```

## Answer

This error does not affect services. This error is caused by defects of the Oracle JVM, but not the platform code. There is the fault tolerance mechanism for Executors in Spark: the Stage retries in case of an Executor crash to ensure the success execution of tasks.

### 26.8.1.14 Why Do the Executors Fail to Register Shuffle Services During the Shuffle of a Large Amount of Data?

#### Question

When more than 50 terabytes of data is shuffled, some executors fail to register shuffle services due to timeout. The shuffle tasks then fail. Why? The error log is as follows:

```
2016-10-19 01:33:34,030 | WARN | ContainersLauncher #14 | Exception from container-launch with
container ID: container_e1452_1476801295027_2003_01_004512 and exit code: 1 |
LinuxContainerExecutor.java:397
ExitCodeException exitCode=1:
at org.apache.hadoop.util.Shell.runCommand(Shell.java:561)
at org.apache.hadoop.util.Shell.run(Shell.java:472)
at org.apache.hadoop.util.Shell$ShellCommandExecutor.execute(Shell.java:738)
at
org.apache.hadoop.yarn.server.nodemanager.LinuxContainerExecutor.launchContainer(LinuxContainerExecuto
r.java:381)
at
org.apache.hadoop.yarn.server.nodemanager.containermanager.launcher.ContainerLaunch.call(ContainerLaun
ch.java:312)
at
org.apache.hadoop.yarn.server.nodemanager.containermanager.launcher.ContainerLaunch.call(ContainerLaun
ch.java:88)
at java.util.concurrent.FutureTask.run(FutureTask.java:266)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
at java.lang.Thread.run(Thread.java:745)
2016-10-19 01:33:34,031 | INFO | ContainersLauncher #14 | Exception from container-launch. |
ContainerExecutor.java:300
2016-10-19 01:33:34,031 | INFO | ContainersLauncher #14 | Container id:
container_e1452_1476801295027_2003_01_004512 | ContainerExecutor.java:300
2016-10-19 01:33:34,031 | INFO | ContainersLauncher #14 | Exit code: 1 | ContainerExecutor.java:300
2016-10-19 01:33:34,031 | INFO | ContainersLauncher #14 | Stack trace: ExitCodeException exitCode=1: |
ContainerExecutor.java:300
```

## Answer

The imported data exceeds 50 TB, which exceeds the shuffle processing capability. The shuffle may fail to respond to the registration request of an executor in a timely manner due to the heavy load.

The timeout interval for an executor to register the shuffle service is 5 seconds. The maximum number of retries is 3. This parameter is not configurable.

You are advised to increase the number of task retry times and the number of allowed executor failure times.

Configure the following parameters in the **spark-defaults.conf** file on the client: If **spark.yarn.max.executor.failures** does not exist, manually add it.

**Table 26-79** Parameter Description

Parameter	Description	Default Value
spark.task.maxFailures	Specifies task retry times.	4
spark.yarn.max.executor.failures	Specifies executor failure attempt times. Set <b>spark.dynamicAllocation.enabled</b> to <b>false</b> , to disable the dynamic allocation of executors.	numExecutors * 2, with minimum of 3
	Specifies executor failure attempt times. Set <b>spark.dynamicAllocation.enabled</b> to <b>true</b> , to enable the dynamic allocation of executors.	3

### 26.8.1.15 Why Does the Out of Memory Error Occur in NodeManager During the Execution of Spark Applications

#### Question

During the execution of Spark applications, if the YARN External Shuffle service is enabled and there are too many shuffle tasks, the **java.lang.OutOfMemoryError: Direct buffer Memory** error occurs, indicating insufficient memory. The error log is as follows:

```
2016-12-06 02:01:00,768 | WARN | shuffle-server-38 | Exception in connection from /192.168.101.95:53680 |
TransportChannelHandler.java:79
io.netty.handler.codec.DecoderException: java.lang.OutOfMemoryError: Direct buffer memory
 at io.netty.handler.codec.ByteToMessageDecoder.channelRead(ByteToMessageDecoder.java:153)
 at
io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:333)
 at
io.netty.channel.AbstractChannelHandlerContext.fireChannelRead(AbstractChannelHandlerContext.java:319)
 at io.netty.channel.DefaultChannelPipeline.fireChannelRead(DefaultChannelPipeline.java:787)
 at io.netty.channel.nio.AbstractNioByteChannel$NioByteUnsafe.read(AbstractNioByteChannel.java:130)
 at io.netty.channel.nio.NioEventLoop.processSelectedKey(NioEventLoop.java:511)
 at io.netty.channel.nio.NioEventLoop.processSelectedKeysOptimized(NioEventLoop.java:468)
```

```

at io.netty.channel.nio.NioEventLoop.processSelectedKeys(NioEventLoop.java:382)
at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:354)
at io.netty.util.concurrent.SingleThreadEventExecutor$2.run(SingleThreadEventExecutor.java:116)
at java.lang.Thread.run(Thread.java:745)
Caused by: java.lang.OutOfMemoryError: Direct buffer memory
at java.nio.Bits.reserveMemory(Bits.java:693)
at java.nio.DirectByteBuffer.<init>(DirectByteBuffer.java:123)
at java.nio.ByteBuffer.allocateDirect(ByteBuffer.java:311)
at io.netty.buffer.PoolArena$DirectArena.newChunk(PoolArena.java:434)
at io.netty.buffer.PoolArena.allocateNormal(PoolArena.java:179)
at io.netty.buffer.PoolArena.allocate(PoolArena.java:168)
at io.netty.buffer.PoolArena.reallocate(PoolArena.java:277)
at io.netty.buffer.PooledByteBuf.capacity(PooledByteBuf.java:108)
at io.netty.buffer.AbstractByteBuf.ensureWritable(AbstractByteBuf.java:251)
at io.netty.buffer.AbstractByteBuf.writeBytes(AbstractByteBuf.java:849)
at io.netty.buffer.AbstractByteBuf.writeBytes(AbstractByteBuf.java:841)
at io.netty.buffer.AbstractByteBuf.writeBytes(AbstractByteBuf.java:831)
at io.netty.handler.codec.ByteToMessageDecoder.channelRead(ByteToMessageDecoder.java:146)
... 10 more

```

## Answer

In the Shuffle Service of YARN, the number of started threads are twice of the number of available CPU cores. The default size of direct buffer memory is 128 MB. If there are too many shuffle tasks connected at the same time, the direct buffer memory allocated to each thread service is insufficient. For example, if there are 40 CPU cores and there are 80 threads started by the Shuffle Service of YARN, the direct buffer memory allocated to each thread is less than 2 MB.

To solve this problem, increase the directory buffer memory based on the number of CPU cores in NodeManager. For example, if there are 40 of CPU cores, increase the direct buffer memory to 512 MB, that is, configure the **GC\_OPTS** parameter of NodeManager as follows:

```
-XX:MaxDirectMemorySize=512M
```

### NOTE

By default, **-XX:MaxDirectMemorySize** is not configured in the **GC\_OPTS** parameter. To configure it, you need to add it to the **GC\_OPTS** parameter as an custom option.

To configure the **GC\_OPTS** parameter, log in to FusionInsight Manager, choose **Cluster > Name of the desired cluster > Services > Yarn > Configurations**, click **All Configurations**, and choose **NodeManager > System**, and then modify the **GC\_OPTS** parameter.

**Table 26-80** Parameter description

Parameter	Description	Default Value
GC_OPTS	The GC parameter of YARN NodeManger.	128M

## 26.8.1.16 Why Does the Realm Information Fail to Be Obtained When SparkBench is Run on HiBench for the Cluster in Security Mode?

### Question

Execution of the sparkbench task (for example, Wordcount) of HiBench6 fails. The bench.log indicates that the Yarn task fails to be executed. The failure information displayed on the Yarn UI is as follows:

```
Exception in thread "main" org.apache.spark.SparkException: Unable to load YARN support
 at org.apache.spark.deploy.SparkHadoopUtil$.liftedTree$1$1 (SparkHadoopUtil.scala:390)
 at org.apache.spark.deploy.SparkHadoopUtil$.yarn$lzycompute (SparkHadoopUtil.scala:385)
 at org.apache.spark.deploy.SparkHadoopUtil$.yarn (SparkHadoopUtil.scala:385)
 at org.apache.spark.deploy.SparkHadoopUtil$.get (SparkHadoopUtil.scala:410)
 at org.apache.spark.deploy.yarn.ApplicationMaster$.main (ApplicationMaster.scala:796)
 at org.apache.spark.deploy.yarn.ExecutorLauncher$.main (ApplicationMaster.scala:821)
 at org.apache.spark.deploy.yarn.ExecutorLauncher.main (ApplicationMaster.scala)
Caused by: java.lang.IllegalArgumentException: Can't get Kerberos realm
 at org.apache.hadoop.security.HadoopKerberosName.setConfiguration (HadoopKerberosName.java:65)
 at org.apache.hadoop.security.UserGroupInformation.initialize (UserGroupInformation.java:288)
 at org.apache.hadoop.security.UserGroupInformation.setConfiguration (UserGroupInformation.java:336)
 at org.apache.spark.deploy.SparkHadoopUtil.<init> (SparkHadoopUtil.scala:51)
 at org.apache.spark.deploy.yarn.YarnSparkHadoopUtil.<init> (YarnSparkHadoopUtil.scala:49)
 at sun.reflect.NativeConstructorAccessorImpl.newInstance0 (Native Method)
 at sun.reflect.NativeConstructorAccessorImpl.newInstance (NativeConstructorAccessorImpl.java:62)
 at sun.reflect.DelegatingConstructorAccessorImpl.newInstance (DelegatingConstructorAccessorImpl.java:45)
 at java.lang.reflect.Constructor.newInstance (Constructor.java:423)
 at java.lang.Class.newInstance (Class.java:442)
 at org.apache.spark.deploy.SparkHadoopUtil$.liftedTree$1$1 (SparkHadoopUtil.scala:387)
 ... 6 more
Caused by: java.lang.reflect.InvocationTargetException
 at sun.reflect.NativeMethodAccessorImpl.invoke0 (Native Method)
 at sun.reflect.NativeMethodAccessorImpl.invoke (NativeMethodAccessorImpl.java:62)
 at sun.reflect.DelegatingMethodAccessorImpl.invoke (DelegatingMethodAccessorImpl.java:43)
 at java.lang.reflect.Method.invoke (Method.java:498)
 at org.apache.hadoop.security.authentication.util.KerberosUtil.getDefaultRealm (KerberosUtil.java:88)
 at org.apache.hadoop.security.HadoopKerberosName.setConfiguration (HadoopKerberosName.java:63)
 ... 16 more
Caused by: KrbException: Cannot locate default realm
 at sun.security.krb5.Config.getDefaultRealm (Config.java:1029)
 ... 22 more
```

### Answer

In C80SPC200 and later, the file stored in the **/etc/krb5.conf** directory is no longer replaced during cluster installation. Instead, the file is stored in the corresponding path on the client through parameter configurations, and HiBench does not reference the client configuration file. Solution: Use the file stored in the **/opt/client/KrbClient/kerberos/var/krb5kdc/krb5.conf** directory on the client to overwrite that in the **/etc/krb5.conf** directories of all nodes. Make a backup before the overwriting.

## 26.8.2 Spark SQL and DataFrame

### 26.8.2.1 What Do I have to Note When Using Spark SQL ROLLUP and CUBE?

#### Question

Suppose that there is a table src(d1, d2, m) with the following data:

```
1 a 1
1 b 1
2 b 2
```

The results for statement "select d1, sum(d1) from src group by d1, d2 with rollup" are shown as below:

```
NULL 0
1 2
2 2
1 1
1 1
2 2
```

Why the first line of the above results is (NULL,0), rather than (NULL,4)?

## Answer

When conducting the rollup and cube operation, we usually perform the dimension-based analysis and what we need is the measurement result, so we would not conduct aggregation operation on the dimension.

Suppose that there is a table src(d1, d2, m), so the statement 1 "select d1, sum(m) from src group by d1, d2 with rollup" conducts the rollup operation on the dimension d1 and d2 to compute the result of m. It has actual business meaning, and its results are in line with the expectation. However, the statement 2 "select d1, sum(d1) from src group by d1, d2 with rollup" cannot be explained from the business perspective. For the statement 2, the result for all aggregations (sum/avg/max/min) is 0.

### NOTE

Only when there is an aggregation operation for fields in "group by" in the rollup and cube operation, the result is 0. For non-rollup and non-cube operations, the result will be in line with the expectation.

## 26.8.2.2 Why Spark SQL Is Displayed as a Temporary Table in Different Databases?

### Question

Why temporary tables of the previous database are displayed after the database is switched?

1. Create a temporary DataSource table, for example:

```
create temporary table ds_parquet
using org.apache.spark.sql.parquet
options(path '/tmp/users.parquet');
```

2. Switch to another database, and run **show tables**. The temporary table created in the previous table is displayed.

```
0: jdbc:hive2://192.168.169.84:22550/default> show tables;
+-----+
| tableName | isTemporary |
+-----+
| ds_parquet | true |
| cmb_tbl_carbon | false |
+-----+
2 rows selected (0.109 seconds)
0: jdbc:hive2://192.168.169.84:22550/default>
```

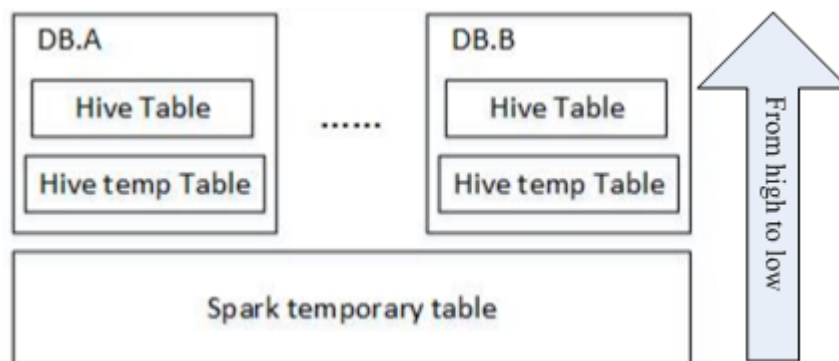
## Answer

The table management hierarchy of Spark is shown in [Figure 26-11](#). The lowest layer stores all temporary DataSource tables. There is no such concept as database at this layer. DataSource tables are visible in various databases.

The MetaStore of Hive is located at the upper layer. This layer distinguishes among databases. In each database, there are two types of Hive table, permanent and temporary. Therefore, Spark supports data tables of the same name at three layers.

During query, SparkSQL first checks for temporary Spark tables, then temporary Hive tables in the current database, and at last the permanent tables in the current database.

**Figure 26-11** Spark table management hierarchy



When a session quits, temporary tables related to the user operation are automatically deleted. Manual deletion of temporary files is not recommended.

When deleting temporary files, use the same priority as that for query. The priorities are temporary Spark table, temporary Hive table, and permanent Hive table ranging from high to low. If you want to directly delete Hive tables but not temporary Spark tables, you can directly use the ***drop table dbName.TableName*** command.

### 26.8.2.3 How to Assign a Parameter Value in a Spark Command?

#### Question

Is it possible to assign parameter values through Spark commands, in addition to through a user interface or a configuration file?

#### Answer

Spark configuration options can be defined either in a configuration file or in Spark commands.

To assign a parameter value, run the `--conf` command on a Spark client. The parameter value takes effect immediately after the command is run.

The command format is `--conf + parameter name + parameter value`. Example command:



```
--conf spark.eventQueue.size=50000
```

## 26.8.2.4 What Directory Permissions Do I Need to Create a Table Using SparkSQL?

### Question

The following error information is displayed when a new user creates a table using SparkSQL:

```
0: jdbc:hive2://192.168.169.84:22550/default> create table testACL(c string);
Error: org.apache.spark.sql.execution.QueryExecutionException: FAILED: Execution Error, return code 1 from
org.apache.hadoop.hive.ql.exec.DDLTask. MetaException(message:Got exception:
org.apache.hadoop.security.AccessControlException
Permission denied: user=testACL, access=EXECUTE, inode="/user/hive/warehouse/
testacl":spark:hadoop:drwxrwx---
 at
org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkAccessAcl(FSPermissionChecker.java:403
)
 at org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.check(FSPermissionChecker.java:306)
 at
org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkTraverse(FSPermissionChecker.java:259)
 at
org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkPermission(FSPermissionChecker.java:20
5)
 at
org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkPermission(FSPermissionChecker.java:19
0)
 at org.apache.hadoop.hdfs.server.namenode.FSDirectory.checkPermission(FSDirectory.java:1710)
 at
org.apache.hadoop.hdfs.server.namenode.FSDirStatAndListingOp.getFileInfo(FSDirStatAndListingOp.java:109)
 at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getFileInfo(FSNamesystem.java:3762)
 at
org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.getFileInfo(NameNodeRpcServer.java:1014)
 at
org.apache.hadoop.hdfs.protocolPB.ClientNamenodeProtocolServerSideTranslatorPB.getFileInfo(ClientNamen
odeProtocolServerSideTranslatorPB.java:853)
 at org.apache.hadoop.hdfs.protocol.proto.ClientNamenodeProtocolProtos$ClientNamenodeProtocol
$2.callBlockingMethod(ClientNamenodeProtocolProtos.java)
 at org.apache.hadoop.ipc.ProtobufRpcEngine$Server$ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:616)
 at org.apache.hadoop.ipc.RPC$Server.call(RPC.java:973)
 at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2089)
 at org.apache.hadoop.ipc.Server$Handler$1.run(Server.java:2085)
 at java.security.AccessController.doPrivileged(Native Method)
 at javax.security.auth.Subject.doAs(Subject.java:422)
 at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1675)
 at org.apache.hadoop.ipc.Server$Handler.run(Server.java:2083)
) (state=,code=0)
```

### Answer

When you create a table using Spark SQL, the interface of Hive is called by the underlying system and a directory named after the table will be created in the **/user/hive/warehouse** directory. Therefore, you must have the permissions to read, write, and execute the **/user/hive/warehouse** directory or the group permission of Hive.

The **/user/hive/warehouse** is specified by the `hive.metastore.warehouse.dir` parameter.

### 26.8.2.5 Why Do I Fail to Delete the UDF Using Another Service?

#### Question

Why do I fail to delete the UDF using another service, for example, delete the UDF created by Hive using Spark SQL.

#### Answer

The UDF can be created using any of the following services:

1. Hive client.
2. JDBCServer API. You can connect JDBCServer to Spark Beeline or JDBC client code, and run SQL statements to create the UDF.
3. spark-sql.

The scenarios in which the UDF failed to be deleted may be as follows:

- If you use Spark Beeline to delete the UDF created by other services, you must restart the JDBCServer before the deletion. Otherwise, the deletion fails. If you use spark-sql to delete the UDF created by other services, you must restart the spark-sql before the deletion. Otherwise, the deletion fails.

Cause: After the UDF is created, if the JDBCServer or the spark-sql has not been restarted, the newly created UDF will not be saved by the FunctionRegistry object in the thread where Spark locates. As a result, the UDF failed to be deleted.

Solution: Restart the JDBCServer and spark-sql of the Spark client and delete the UDF.

- When creating UDF on the Hive client, the **add jar** command (e.g. **add jar /opt/test/two\_udfs.jar**) is used to add the **.jar** package instead of specifying the path of **.jar** package in creating UDF statement. As a result, the **ClassNotFound** error occurs when you use other services to delete the UDF.

Cause: When you use a service to delete the UDF, the service will load the class that corresponds to the UDF to obtain the UDF. However, the **.jar** package is added by the **add jar** command and jar package does not exist in the classpath of other services. As a result, the **ClassNotFound** error occurs and the UDF failed to be deleted.

Solution: The UDF created using the preceding approach must be deleted using the same approach. No other approaches are allowed.

### 26.8.2.6 Why Cannot I Query Newly Inserted Data in a Parquet Hive Table Using SparkSQL?

#### Question

Why cannot I query newly inserted data in a parquet Hive table using SparkSQL? This problem occurs in the following scenarios:

1. For partitioned tables and non-partitioned tables, after data is inserted on the Hive client, the latest inserted data cannot be queried using SparkSQL.

2. After data is inserted into a partitioned table using SparkSQL, if the partition information remains unchanged, the newly inserted data cannot be queried using SparkSQL.

## Answer

To improve Spark performance, parquet metadata is cached. When the parquet table is updated by Hive or another means, the cached metadata remains unchanged, resulting in SparkSQL failing to query the newly inserted data.

For a parquet Hive partition table, if the partition information remains unchanged after data is inserted, the cached metadata is not updated. As a result, the newly inserted data cannot be queried by SparkSQL.

To solve the query problem, update metadata before starting a Spark SQL query.

***REFRESH TABLE table\_name;***

*table\_name* indicates the name of the table to be updated. The table must exist. Otherwise, an error is reported.

When the query statement is executed, the latest inserted data can be obtained.

### 26.8.2.7 How to Use Cache Table?

#### Question

What is cache table used for? Which point should I pay attention to while using cache table?

#### Answer

Spark SQL caches tables into memory so that data can be directly read from memory instead of disks, reducing memory overhead due to disk reads.

Note that cached tables consume Executor's memory. This means that caching large or many tables compromises Executor's stability even if compressed storage has been used to reduce memory overhead as much as possible.

If it is no longer necessary to accelerate data query by means of cache table, run the following command to uncache tables to free up memory:

***uncache table table\_name***

#### NOTE

The Storage tab page of the Spark Driver user interface displays the cached tables.

### 26.8.2.8 Why Are Some Partitions Empty During Repartition?

#### Question

During the repartition operation, the number of blocks (**`spark.sql.shuffle.partitions`**) is set to 4,500, and the number of keys used by repartition exceeds 4,000. It is expected that data corresponding to different keys

can be allocated to different partitions. However, only 2,000 partitions have data, and data corresponding to different keys is allocated to the same partition.

## Answer

This is normal.

The partition to which data is distributed is obtained by performing a modulo operation on hashcode of a key. Different hashcodes may have the same modulo result. In this case, data is distributed to the same partition, as a result, some partitions do not have data, and some partitions have data corresponding to multiple keys.

You can adjust the value of **spark.sql.shuffle.partitions** to adjust the cardinality during modulo operation and improve the unevenness of data blocks. After multiple verifications, it is found that the effect is good when the parameter is set to a prime number or an odd number.

Configure the following parameters in the **spark-defaults.conf** file on the Driver client.

**Table 26-81** Parameter Description

Parameter	Description	Default Value
spark.sql.shuffle.partitions	Number of shuffle data blocks during the shuffle operation.	200

### 26.8.2.9 Why Does 16 Terabytes of Text Data Fails to Be Converted into 4 Terabytes of Parquet Data?

#### Question

When the default configuration is used, 16 terabytes of text data fails to be converted into 4 terabytes of parquet data, and the error information below is displayed. Why?

```
Job aborted due to stage failure: Task 2866 in stage 11.0 failed 4 times, most recent failure: Lost task 2866.6 in stage 11.0 (TID 54863, linux-161, 2): java.io.IOException: Failed to connect to /10.16.1.11:23124 at org.apache.spark.network.client.TransportClientFactory.createClient(TransportClientFactory.java:214) at org.apache.spark.network.client.TransportClientFactory.createClient(TransportClientFactory.java:167) at org.apache.spark.network.netty.NettyBlockTransferService$anon$1.createAndStart(NettyBlockTransferService.scala:92)
```

**Table 26-82** lists the default configuration.

**Table 26-82** Parameter Description

Parameter	Description	Default Value
spark.sql.shuffle.partitions	Number of shuffle data blocks during the shuffle operation.	200

Parameter	Description	Default Value
spark.shuffle.sasl.timeout	Timeout interval of SASL authentication for the shuffle operation. Unit: second	120s
spark.shuffle.io.connectionTimeout	Timeout interval for connecting to a remote node during the shuffle operation. Unit: second	120s
spark.network.timeout	Timeout interval for all network connection operations. Unit: second	360s

## Answer

The current data volume is 16 TB, but the number of partitions is only 200. As a result, each task is overloaded and the preceding problem occurs.

To solve the preceding problem, you need to adjust the parameters.

- Increase the number of partitions to divide the task into smaller ones.
- Increase the timeout interval during task execution.

Configure the following parameters in the **spark-defaults.conf** file on the client:

**Table 26-83** Parameter Description

Parameter	Description	Recommended Value
spark.sql.shuffle.partitions	Number of shuffle data blocks during the shuffle operation.	4501
spark.shuffle.sasl.timeout	Timeout interval of SASL authentication for the shuffle operation. Unit: second	2000s
spark.shuffle.io.connectionTimeout	Timeout interval for connecting to a remote node during the shuffle operation. Unit: second	3000s
spark.network.timeout	Timeout interval for all network connection operations. Unit: second	360s

### 26.8.2.10 Why the Operation Fails When the Table Name Is TABLE?

#### Question

When the table name is set to **table**, why the error information similar to the following is displayed after the **drop table table** command or other command is run?

```
16/07/12 18:56:29 ERROR SparkSQLDriver: Failed in [drop table table]
java.lang.RuntimeException: [1.1] failure: identifier expected
table
^
at scala.sys.package$.error(package.scala:27)
at org.apache.spark.sql.catalyst.SqlParserTrait$class.parseTableIdentifier(SqlParser.scala:56)
at org.apache.spark.sql.catalyst.SqlParser$.parseTableIdentifier(SqlParser.scala:485)
```

#### Answer

The word **table** is a keyword of Spark SQL statements and must not be used as a table name.

### 26.8.2.11 Why Is a Task Suspended When the ANALYZE TABLE Statement Is Executed and Resources Are Insufficient?

#### Question

When the **analyze table** statement is executed using **spark-sql**, the task is suspended and the information below is displayed. Why?

```
spark-sql> analyze table hivetable2 compute statistics;
Query ID = root_20160716174218_90f55869-000a-40b4-a908-533f63866fed
Total jobs = 1
Launching Job 1 out of 1
Number of reduce tasks is set to 0 since there's no reduce operator
16/07/20 17:40:56 WARN JobResourceUploader: Hadoop command-line option parsing not performed.
Implement the Tool interface and execute your application with ToolRunner to remedy this.
Starting Job = job_1468982600676_0002, Tracking URL = http://10-120-175-107:8088/proxy/
application_1468982600676_0002/
Kill Command = /opt/client/HDFS/hadoop/bin/hadoop job -kill job_1468982600676_0002
```

#### Answer

When the statement is executed, the SQL statement starts the **analyze table hivetable2 compute statistics** MapReduce tasks. On the ResourceManager Web UI of Yarn, the task is not executed due to insufficient resources. As a result, the task is suspended.

Figure 26-12 ResourceManager web UI

application_	Type	default	Start Time	End Time	Status	Modes
analyze table hivetable2 compute statistics(Stage=0)	MAPREDUCE	default	Wed Jul 20 17:40:56 +0800 2016	Wed Jul 20 17:40:56	ACCEPTED UNDEFINED	ApplicationMaster 0
SparkSQL::192.168.169.84	SPARK	default	Wed Jul 20 17:40:56	Wed Jul 20 17:40:56	RUNNING UNDEFINED	ApplicationMaster 0

You are advised to add **noscan** when running the **analyze table** statement. The function of this statement is the same as that of the **analyze table hivetable2 compute statistics** statement. The command is as follows:

```
spark-sql> analyze table hivetable2 compute statistics noscan
```

This command does not start MapReduce tasks and does not occupy Yarn resources. Therefore, the tasks can be executed.

### 26.8.2.12 If I Access a parquet Table on Which I Do not Have Permission, Why a Job Is Run Before "Missing Privileges" Is Displayed?

#### Question

If I access a parquet table on which I do not have permission, why a job is run before "Missing Privileges" is displayed?

#### Answer

The execution sequence of Spark SQL statement parse the table in the statement first, then obtain the metadata in the table, and finally check the permission.

The metadata of a parquet table contains the Split information (which is read by HDFS API) about files. If the table contains many files, the HDFS API reads data in serial mode, in which degrades the performance. If the number of files in the table exceeds the threshold *spark.sql.sources.parallelSplitDiscovery.threshold*, a job will be generated to use Executor to read the data in parallel mode.

The permission authentication is executed after the metadata is obtained. Therefore, when the number of files in the table exceeds the threshold, a job is run before the permission authentication error message **Missing Privileges**.

### 26.8.2.13 Why Do I Fail to Modify MetaData by Running the Hive Command?

#### Question

When do I fail to modify the metadata in the datasource and Spark on HBase table by running the Hive command?

#### Answer

The current Spark version does not support modifying the metadata in the datasource and Spark on HBase tables by running the Hive command.

### 26.8.2.14 Why Is "RejectedExecutionException" Displayed When I Exit Spark SQL?

#### Question

After successfully running Spark tasks with large data volume, for example, 2-TB TPCDS test suite, why is the abnormal stack information **"RejectedExecutionException"** displayed sometimes? The log is as follows:

```
16/07/16 10:19:56 ERROR TransportResponseHandler: Still have 2 requests outstanding when connection from linux-192/10.1.1.5:59250 is closed
java.util.concurrent.RejectedExecutionException: Task scala.concurrent.impl.CallbackRunnable@5fc1ab rejected from java.util.concurrent.ThreadPoolExecutor@52fa7e19[Terminated, pool size = 0, active threads = 0, queued tasks = 0, completed tasks = 3025]
```

## Answer

When Spark SQL is closed, the application and the message channel are closed. If there are unprocessed messages, the connection should be closed to rectify the exception. If the thread pool inside Scala is closed, the abnormal stack information "RejectedExecutionException" is displayed. This abnormal stack information will not be displayed if the thread pool inside Scala is not closed.

The error occurs when the application is successfully run and closed. Therefore, the error will not affect the services.

### 26.8.2.15 How Do I Do If I Incidentally Kill the JDBCServer Process During Health Check?

#### Question

In the health check solution, when the number of concurrently executed statements reaches the upper limit of the thread pool, the health check command fails to run. As a result, the health check program times out and the Spark JDBCServer process is killed.

#### Answer

Currently, JDBCServer has two thread pools **HiveServer2-Handler-Pool** and **HiveServer2-Background-Pool**. HiveServer2-Handler-Pool is used to process session connections, and HiveServer2-Background-Pool is used to execute SQL statements.

In the current health check mechanism, a session connection is created and the health check command **HEALTHCHECK** is executed in the thread where the session is located to determine the health status of Spark JDBCServer. Therefore, HiveServer2-Handler-Pool must reserve a thread to process the health check session connection and execute the health check command, otherwise, the health check session cannot be established or the health check command cannot be executed. As a result, Spark JDBCServer is regarded as unhealthy and then killed. That is, if the number of thread pools of HiveServer2-Handler-Pool is 100, a maximum of 99 sessions can be connected.

### 26.8.2.16 Why No Result Is found When 2016-6-30 Is Set in the Date Field as the Filter Condition?

#### Question

Why no result is found when 2016-6-30 is set in the date field as the filter condition?

As shown in the following figure, `trx_dte_par` in the `select count (*) from trxfintrx2012 a where trx_dte_par='2016-6-30'` statement is a date field. However, no search result is found when the filter condition is `where trx_dte_par='2016-6-30'`. Search results are found only when the filter condition is `where trx_dte_par='2016-06-30'`.



Figure 26-13 Example

```

0: jdbc:hive2://ha-cluster/default> select count(*)
0: jdbc:hive2://ha-cluster/default> from TRXFINTRX2012 a
0: jdbc:hive2://ha-cluster/default> where trx_dte_par = '2016-6-30';
+-----+
| _c0 |
+-----+
| 0 |
+-----+
1 row selected (0.498 seconds)
0: jdbc:hive2://ha-cluster/default> select count(*)
0: jdbc:hive2://ha-cluster/default> from TRXFINTRX2012 a
0: jdbc:hive2://ha-cluster/default> where trx_dte_par = '2016-06-30';
+-----+
| _c0 |
+-----+
| 8520808 |
+-----+
1 row selected (15.788 seconds)

```

## Answer

If a data string of the date type is present in Spark SQL statements, the Spark SQL will search the matching character string without checking the date format. In this case, if the date format in the SQL statement is incorrect, the query will fail. For example, if the data format is yyyy-mm-dd, then no search results matching '2016-6-30' will be found.

### 26.8.2.17 Why Does the "--hivevar" Option I Specified in the Command for Starting spark-beeline Fail to Take Effect?

#### Question

Why does the `--hivevar` option I specified in the command for starting spark-beeline fail to take effect?

In the V100R002C60 version, if I use the `--hivevar <VAR_NAME>=<var_value>` option to define a variable in the command for starting spark-beeline, no error is reported in spark-beeline. However, if the variable `<VAR_NAME>` is used in SQL, the variable cannot be parsed and the `<VAR_NAME>` exception is reported.

For example:

1. Run the following command to start the spark-beeline:  
`spark-beeline --hivevar <VAR_NAME>=<var_value>`
2. After spark-beeline is started successfully, I run the SQL statements `DROP TABLE ${VAR_NAME}` in spark-beeline. The `VAR_NAME` exception occurs.

## Answer

In the V100R002C60 version, the `--hivevar <VAR_NAME>=<var_value>` feature of Hive is not supported in Spark because multi-session management function is added. Therefore, the `--hivevar` option in the command for starting spark-beeline is invalid.

## 26.8.2.18 Why Does the "Permission denied" Exception Occur When I Create a Temporary Table or View in Spark-beeline?

### Question

In normal mode, when I create a temporary table or view in spark-beeline, the error message "Permission denied" is displayed, indicating that I have no permissions on the HDFS directory. The error log information is as follows:

```
org.apache.hadoop.security.AccessControlException Permission denied: user=root, access=EXECUTE,
inode="/tmp/spark/sparkhive-scratch/omm/e579a76f-43ed-4014-8a54-1072c07ceeff/_tmp_space.db/
52db1561-60b0-4e7d-8a25-c2eaa44850a9":omm:hadoop:drwx-----
```

### Answer

In normal mode, if you run the spark-beeline command as a non-omm user, **root** user for example, without specifying the **-n** parameter, your account is still the root user. After spark-beeline is started, a new HDFS directory is created by JDBCServer. In the current version of DataSight, the user that starts the JDBCServer is **omm**. In versions earlier than DataSight V100R002C30, the user is **root**. Therefore, the owner of the HDFS directory is **omm** and the group is **hadoop**. The HDFS directory is used when you create a temporary table or view in spark-beeline and the user **root** is a common user in HDFS and has no permissions on the directory of user **omm**. As a result, the "Permission denied" exception occurs.

In normal mode, only user **omm** can create a temporary table or view. To solve this problem, you can specify the **-n omm** option for user **omm** when starting spark-beeline. In this way, you have the permissions to perform operations on the HDFS directory.

## 26.8.2.19 Why Is the "Code of method ... grows beyond 64 KB" Error Message Displayed When I Run Complex SQL Statements?

### Question

When I run a complex SQL statement, for example, SQL statements with multiple layers of nesting statements and a single layer statement contains a large number of logic clauses such as case when, an error message indicating that the code of a certain method exceeds 64 KB is displayed. The log is as follows:

```
java.util.concurrent.ExecutionException: java.lang.Exception: failed to compile:
org.codehaus.janino.JaninoRuntimeException: Code of method "(Lorg/apache/spark/sql/catalyst/expressions/
GeneratedClass$SpecificUnsafeProjection;Lorg/apache/spark/sql/catalyst/InternalRow;)V" of class
"org.apache.spark.sql.catalyst.expressions.GeneratedClass$SpecificUnsafeProjection" grows beyond 64 KB
```

### Answer

If Project Tungsten is enabled, Spark will use codegen method to generate Java code for part of execution plan. However, each function in Java code to be compiled by JDK must be less than 64 KB. If complex SQL statements are run, the function in the Java code generated by codegen may exceed 64 KB, causing compilation failure.

To solve the problem, go to the **spark-defaults.conf** file on the client and set the **spark.sql.codegen.wholeStage** parameter to **false** to disable Project Tungsten.

## 26.8.2.20 Why Is Memory Insufficient if 10 Terabytes of TPCDS Test Suites Are Consecutively Run in Beeline/JDBCServer Mode?

### Question

When the driver memory is set to 10 GB and the 10 TB TPCDS test suites are continuously run in Beeline/JDBCServer mode, SQL statements fail to be executed due to insufficient driver memory. Why?

### Answer

By default, 1000 UI data records of jobs and stages are reserved in the memory.

The function of overflowing UI data to disks has been added to optimize large clusters. The overflow condition is that the size of UI data in each stage reaches the minimum threshold 5 MB. If the number of tasks in each stage is small, the size of UI data in the stage may not reach the threshold. As a result, the UI data in the stage is cached in the memory until the number of UI data records reaches the upper limit (1000 by default). Only then the old UI data is cleared from the memory.

Therefore, before the old UI data is cleared, the UI data occupies a large amount of memory. As a result, the driver memory is insufficient when 10 terabytes of TPCDS test suites are executed.

Workaround:

- Set **spark.ui.retainedJobs** and **spark.ui.retainedStages** based on service requirements to specify the number of UI data records of jobs and stages to be reserved. For details, see [Table 26-15](#) in [Common Parameters](#).
- If a large amount of UI data of jobs and stages needs to be reserved, increase the memory of the driver by setting the **spark.driver.memory** parameter. For details, see [Table 26-12](#) in [Common Parameters](#).

## 26.8.2.21 Why Functions Cannot Be Used When Different JDBCServer Are Connected?

### Question

Scenario 1:

I run the **add jar** command to create permanent functions. When Beeline connects to different JDBCServer or the JDBCServer is restarted, I need to run the **add jar** command again.

Figure 26-14 Error information in Scenario 1

```

0: jdbc:hive2://192.168.91.247:23040/default> create function a1 as '
-----+-----+
| result |
-----+-----+
NO rows selected (0.222 seconds)
0: jdbc:hive2://192.168.91.247:23040/default> SELECT test.a1(array(1, 2, 3), array(2));
-----+-----+
| _c0 |
-----+-----+
| true |
-----+-----+
1 row selected (8.282 seconds)
0: jdbc:hive2://192.168.91.247:23040/default> closing: 0: jdbc:hive2://192.168.91.247:24002,192.168.154.81:24002,192.168.8.27:24002;serviceDiscoveryMode=zooKeeper;auth-conf;auth=KERBEROS;principal=spark/hadoop,hadoop.com@HADOOP.COM;
100-106-121-140:/opt/hadoopclient # ./spark-beeline
it's running the fi spark-beeline, it calls /opt/hadoopclient/spark/spark/bin/beeline
and helps to connect to the JDBCServer automatically
connecting to jdbc:hive2://192.168.91.247:24002,192.168.154.81:24002;serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=sparkthriftserver;sas
doop,hadoop.com@HADOOP.COM;
2017-06-15 08:17:55:495 | WARN | Thread-2 | TGT refresh thread time adjusted from : Thu Jun 15 05:59:42 GMT+08:00 2017 to : Thu Jun 15 08:18:55 GMT+08:00 2017
fresh interval (60 seconds) from now, | org.apache.zookeeper.Login$.run(Login.java:177)
2017-06-15 08:17:56:743 | WARN | main | Unable to load native-hadoop library for your platform... using builtin-java classes where applicable | org.apache.hadoop
java:62)
2017-06-15 08:17:56:773 | WARN | TGT Renewer for sparkuser@HADOOP.COM | Exception encountered while running the renewal command. Aborting renew thread. ExitCo
requested option while renewing credentials
| org.apache.hadoop.security.UserGroupInformation$.run(UserGroupInformation.java:946)
connected to: spark sql (version)
Driver: Hive JDBC (version 1.2.1.spark)
Transaction isolation: TRANSACTION_REPEATABLE_READ
beeline version 1.2.1.spark by Apache Hive
[INFO] unable to bind key for unsupported operation: backward-delete-word
[INFO] unable to bind key for unsupported operation: backward-delete-word
[INFO] unable to bind key for unsupported operation: down-history
[INFO] unable to bind key for unsupported operation: up-history
[INFO] unable to bind key for unsupported operation: up-history
[INFO] unable to bind key for unsupported operation: down-history
[INFO] unable to bind key for unsupported operation: up-history
[INFO] unable to bind key for unsupported operation: up-history
[INFO] unable to bind key for unsupported operation: down-history
[INFO] unable to bind key for unsupported operation: up-history
[INFO] unable to bind key for unsupported operation: down-history
0: jdbc:hive2://192.168.8.27:23040/default> SELECT test.a1(array(1, 2, 3), array(2));
Error: org.apache.spark.sql.AnalysisException: unable to load udf class (state=,code=0)
0: jdbc:hive2://192.168.8.27:23040/default> set role admin;
-----+-----+
| key | value |
-----+-----+
| role admin |
-----+-----+
1 row selected (0.465 seconds)
0: jdbc:hive2://192.168.8.27:23040/default> add jar /home/smartcare-udf-0.0.1-SNAPSHOT.jar;
-----+-----+
| result |
-----+-----+
| 0 |
-----+-----+

```

Scenario 2:

The functions can be queried by running the **show functions** command, but these functions cannot be used. This is because JAR files in the corresponding path do not exist on the connected JDBC node. After the JAR files are added, the query succeeds.

Figure 26-15 Error information in scenario 2

```

-----+-----+
| function |
-----+-----+
| stddev_pop |
| stddev_samp |
| str_to_map |
| string |
| struct |
| substr |
| substring |
| substring_index |
| sum |
| tan |
| tanh |
| test.a1 |
| timestamp |
| tinyint |
| to_date |
| to_unix_timestamp |
| to_utc_timestamp |
| translate |
| trim |
| trunc |
| ucase |
| unbase64 |
| unhex |
| unix_timestamp |
| upper |
| var_pop |
| var_samp |
| variance |
| weekofyear |
| when |
| window |
| xpath |
0: jdbc:hive2://192.168.8.27:22550/default> use test;
-----+-----+
| Result |
-----+-----+
No rows selected (0.038 seconds)
0: jdbc:hive2://192.168.8.27:22550/default> SELECT test.a1(array(1, 2, 3), array(2));
Error: org.apache.spark.sql.AnalysisException: undefined function: 'test.a1'. This function is neither a registered temporary function nor a permanent
7 (state=,code=0)
0: jdbc:hive2://192.168.8.27:22550/default> show functions;
-----+-----+
| function |
-----+-----+

```

Answer

Scenario 1:

The **addjar** statement loads the jar only to the jarClassLoader of the currently connected JDBCServer. Different JDBCServer do not share the jarClassLoader.

After JDBCServer restarts, new jarClassLoader is created. So the **addjar** statement needs to be run again.

You can add a JAR file in either of the following ways: Add a JAR file when starting Spark SQL, for example, by running **spark-sql --jars /opt/test/two\_udfs.jar**. Add a JAR file after Spark SQL is started, for example, by running **add jar /opt/test/two\_udfs.jar**. The path specified by **add jar** can be a local or an HDFS path.

Scenario 2:

The **show functions** command obtains all functions in the current database from the external catalog. When a function is used in SQL, JDBCServer loads the JAR file corresponding to the function.

If the JAR file does not exist, the function cannot be used. In this case, run the **add jar** command again.

## 26.8.2.22 Why Does an Exception Occur When I Drop Functions Created Using the Add Jar Statement?

### Question

- Question 1

Why can I successfully drop functions without the drop function permission? The specific operations are as follows:

- a. On FusionInsight Manager, I added user **user1** and granted the user the admin permission.

```
set role admin;add jar /home/smartcare-udf-0.0.1-SNAPSHOT.jar;create database db4;use db4;create function f11 as 'com..smartcare.dac.hive.udf.UDFArrayGreaterEqual';create function f12 as 'com..smartcare.dac.hive.udf.UDFArrayGreaterEqual';
```

- b. Then I canceled the admin permission:

```
drop functiondb4.f11;
```

The result shows that the drop operation is successful.

**Figure 26-16** Command output of the drop function command

```
source /opt/${clientPath}/bigdata_env;/opt/${clientPath}/Spark2x/spark/bin/beeline -u 'jdbc:hive2://10.90.46.60:24002,10.90.46.61:24002,10.90.46.62:24002;serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=sparkthriftserver2x;saslQop=auth-conf;auth=KERBEROS;principal=spark2x/hadoop.hadoop.com@HADOOP.COM;' -e "drop function db4.f11;"
```

- Question 2

After the preceding operation, I ran the **show function** command and the command output indicates that the function still exists. The specific operations are as follows:

- a. On FusionInsight Manager, I added user **user1** and granted the user the admin permission. Then I ran the following commands in spark-beeline:

```
set role admin;create database db2;use db2;add jar /home/smartcare-udf-0.0.1-SNAPSHOT.jar;create function f11 as 'com..smartcare.dac.hive.udf.UDFArrayGreaterEqual';create function f12 as 'com..smartcare.dac.hive.udf.UDFArrayGreaterEqual';
```

- b. I exited and re-logged in to spark-beeline and ran the following command:

```
set role admin;use db2;drop function db2.f11;
```

- c. I exited and re-logged in to spark-beeline and ran the following command:

```
use db2;show functions;
```

The result shows that the dropped function still exists.

**Figure 26-17** Output of the show functions command

```
| datediff |
| day |
| dayofmonth |
| dayofyear |
| db2. f11 |
| db2. f12 |
| decimal |
| decode |
```

## Answer

- Root cause:  
The root cause is that, in multi-active instance or multi-tenant mode, the function created using the **add jar** command in spark-beeline is invisible to JDBCServer instances. When you run the **drop function** command, if the JDBCServer instance connected by the session is not the JDBCServer instance that creates the function, the function cannot be found in the session. In addition, the default value of **hive.exec.drop.ignorenonexistent** in Hive is **true**. Therefore, if the function does not exist, no exception is reported when the function deletion operation is performed. As a result, no exception is reported when you delete a non-existing function and perform drop operation even though you do not have the drop function permission. When you run the **show function** command after restarting the session and connecting to the JDBCServer instance that creates the function, the command output indicates that the function still exists. This problem is inherited from the Hive community.
- Solution:  
Before you run the **drop function** command, run the **add jar** command. In this way, the drop operation succeeds only when you are granted the drop function permission, and the dropped function will not be displayed if you run the **show function** command.

### 26.8.2.23 Why Does Spark2x Have No Access to DataSource Tables Created by Spark1.5?

#### Question

When Spark2x accesses the DataSource table created by Spark1.5, a message is displayed indicating that schema information cannot be obtained. As a result, the table cannot be accessed. Why?

#### Answer

- Cause analysis:  
This is because the formats of the DataSource table information stored in Spark2x and Spark1.5 are inconsistent. Spark 1.5 divides schema information

into multiple parts and uses **path.park.0** as the key for storage. Spark 1.5 reads information from each part and reassembles the information into complete one. Spark2x directly uses the corresponding key to obtain the corresponding information. In this case, when Spark2x reads the DataSource table created by Spark1.5, the information corresponding to the key cannot be read. As a result, the DataSource table information fails to be parsed.

When processing Hive tables, Spark2x and Spark1.5 use the same storage mode. Therefore, Spark2x can directly read tables created by Spark1.5.

- Workaround:

In Spark2x, create a foreign table to point to the actual data in the Spark1.5 table. In this way, the DataSource table created by Spark1.5 can be read in Spark2x. In addition, after Spark1.5 updates data, Spark2x can detect the change. The reverse is also true. In this way, Spark2x can access the DataSource table created by Spark1.5.

## 26.8.2.24 Why Does Spark-beeline Fail to Run and Error Message "Failed to create ThriftService instance" Is Displayed?

### Question

Why does "Failed to create ThriftService instance" occur when spark beeline fails to run?

Beeline logs are as follows:

```
Error: Failed to create ThriftService instance (state=,code=0)
Beeline version 1.2.1.spark by Apache Hive
[INFO] Unable to bind key for unsupported operation: backward-delete-word
[INFO] Unable to bind key for unsupported operation: backward-delete-word
[INFO] Unable to bind key for unsupported operation: down-history
[INFO] Unable to bind key for unsupported operation: up-history
[INFO] Unable to bind key for unsupported operation: up-history
[INFO] Unable to bind key for unsupported operation: down-history
[INFO] Unable to bind key for unsupported operation: up-history
[INFO] Unable to bind key for unsupported operation: down-history
[INFO] Unable to bind key for unsupported operation: up-history
[INFO] Unable to bind key for unsupported operation: down-history
[INFO] Unable to bind key for unsupported operation: up-history
[INFO] Unable to bind key for unsupported operation: down-history
[INFO] Unable to bind key for unsupported operation: up-history
[INFO] Unable to bind key for unsupported operation: down-history
beeline>
```

In addition, the "Timed out waiting for client to connect" error log is generated on the JDBCServer. The details are as follows:

```
2017-07-12 17:35:11,284 | INFO | [main] | Will try to open client transport with JDBC Uri:
jdbc:hive2://192.168.101.97:23040/default;principal=spark/hadoop.<System domain name>@<System
domain name>;healthcheck=true;saslQop=auth-conf;auth=KERBEROS;user.principal=spark/hadoop.<System
domain name>@<System domain name>;user.keytab=${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/
FusionInsight-Spark-*/keytab/spark/JDBCServer/spark.keytab |
org.apache.hive.jdbc.HiveConnection.openTransport(HiveConnection.java:317)
2017-07-12 17:35:11,326 | INFO | [HiveServer2-Handler-Pool: Thread-92] | Client protocol version:
HIVE_CLI_SERVICE_PROTOCOL_V8 |
org.apache.proxy.service.ThriftCLIProxyService.OpenSession(ThriftCLIProxyService.java:554)
2017-07-12 17:35:49,790 | ERROR | [HiveServer2-Handler-Pool: Thread-113] | Timed out waiting for client
to connect.
Possible reasons include network issues, errors in remote driver or the cluster has no available resources, etc.
Please check YARN or Spark driver's logs for further information. |
org.apache.proxy.service.client.SparkClientImpl.<init>(SparkClientImpl.java:90)
java.util.concurrent.ExecutionException: java.util.concurrent.TimeoutException: Timed out waiting for
client connection.
at io.netty.util.concurrent.AbstractFuture.get(AbstractFuture.java:37)
```

```
at org.apache.proxy.service.client.SparkClientImpl.<init>(SparkClientImpl.java:87)
at org.apache.proxy.service.client.SparkClientFactory.createClient(SparkClientFactory.java:79)
at org.apache.proxy.service.SparkClientManager.createSparkClient(SparkClientManager.java:145)
at org.apache.proxy.service.SparkClientManager.createThriftServerInstance(SparkClientManager.java:160)
at org.apache.proxy.service.ThriftServiceManager.getOrCreateThriftServer(ThriftServiceManager.java:182)
at org.apache.proxy.service.ThriftCLIProxyService.OpenSession(ThriftCLIProxyService.java:596)
at org.apache.hive.service.cli.thrift.TCLIService$Processor$OpenSession.getResult(TCLIService.java:1257)
at org.apache.hive.service.cli.thrift.TCLIService$Processor$OpenSession.getResult(TCLIService.java:1242)
at org.apache.thrift.ProcessFunction.process(ProcessFunction.java:39)
at org.apache.thrift.TBaseProcessor.process(TBaseProcessor.java:39)
at org.apache.hadoop.hive.thrift.HadoopThriftAuthBridge$Server
$TUGIAssumingProcessor.process(HadoopThriftAuthBridge.java:696)
at org.apache.thrift.server.TThreadPoolServer$WorkerProcess.run(TThreadPoolServer.java:286)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
at java.lang.Thread.run(Thread.java:748)
Caused by: java.util.concurrent.TimeoutException: Timed out waiting for client connection.
```

## Answer

This problem occurs when the network is unstable. When a timed-out exception occurs in beeline, Spark does not attempt to reconnect to beeline.

### Solution

Restart spark-beeline for reconnection.

## 26.8.2.25 Why Cannot I Query Newly Inserted Data in an ORC Hive Table Using Spark SQL?

### Question

Why cannot I query newly inserted data in an ORC Hive table using Spark SQL? This problem occurs in the following scenarios:

- For partitioned tables and non-partitioned tables, after data is inserted on the Hive client, the latest inserted data cannot be queried using Spark SQL.
- After data is inserted into a partitioned table using Spark SQL, if the partition information remains unchanged, the newly inserted data cannot be queried using Spark SQL.

### Answer

To improve Spark performance, ORC metadata is cached. When the ORC table is updated by Hive or another means, the cached metadata remains unchanged, resulting in Spark SQL failing to query the newly inserted data.

For an ORC Hive partition table, if the partition information remains unchanged after data is inserted, the cached metadata is not updated. As a result, the newly inserted data cannot be queried by Spark SQL.

### Solution

1. To solve the query problem, update metadata before starting a Spark SQL query.

```
REFRESH TABLE table_name;
```

*table\_name* indicates the name of the table to be updated. The table must exist. Otherwise, an error is reported.



When the query statement is executed, the latest inserted data can be obtained.

- Run the following command to disable Spark optimization when using Spark:  
**set spark.sql.hive.convertMetastoreOrc=false;**

## 26.8.3 Spark Streaming

### 26.8.3.1 Same DAG Log Is Recorded Twice for a Streaming Task

#### Question

I use Spark Streaming to run the following command:

```
spark-submit -master yarn-client --conf spark.logLineage=true --jars $SPARK_HOME/jars/streamingClient/kafka-clients-0.8.2.1.jar,$SPARK_HOME/jars/streamingClient/kafka_2.11-0.8.2.1.jar,$SPARK_HOME/jars/streamingClient/spark-streaming-kafka-0-8_2.11-2.1.0.jar --class com..bigdata.spark.examples.FemaleInfoCollectionPrint /opt/female/SparkStreamingJavaExample-1.0.jar <checkpoint> <batchTime> <windowTime> <topics> <brokers>
```

When there is no Kafka data input, the directed acyclic graph (DAG) of RDD displayed in the log is printed twice in a batch. The log is as follows:

```

Time: 1491447950000 ms

17/04/06 11:06:00 INFO SparkContext: RDD's recursive dependencies:
(2) MapPartitionsRDD[49] at filter at FemaleInfoCollectionPrint.java:111 []
| MapPartitionsRDD[48] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 []
| CoGroupedRDD[47] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 []
| MapPartitionsRDD[38] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 []
| CachedPartitions: 2; MemorySize: 8.0 B; ExternalBlockStoreSize: 0.0 B; DiskSize: 0.0 B
| ReliableCheckpointRDD[40] at print at FemaleInfoCollectionPrint.java:123 []
| ShuffledRDD[36] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 []
| CachedPartitions: 2; MemorySize: 8.0 B; ExternalBlockStoreSize: 0.0 B; DiskSize: 0.0 B
+- (5) MapPartitionsRDD[35] at map at FemaleInfoCollectionPrint.java:81 []
| MapPartitionsRDD[34] at filter at FemaleInfoCollectionPrint.java:81 []
| MapPartitionsRDD[33] at map at FemaleInfoCollectionPrint.java:72 []
| MapPartitionsRDD[32] at map at FemaleInfoCollectionPrint.java:63 []
| KafkaRDD[31] at createDirectStream at FemaleInfoCollectionPrint.java:63 []
| ShuffledRDD[46] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 []
+- (5) MapPartitionsRDD[45] at map at FemaleInfoCollectionPrint.java:81 []
| MapPartitionsRDD[44] at filter at FemaleInfoCollectionPrint.java:81 []
| MapPartitionsRDD[43] at map at FemaleInfoCollectionPrint.java:72 []
| MapPartitionsRDD[42] at map at FemaleInfoCollectionPrint.java:63 []
| KafkaRDD[41] at createDirectStream at FemaleInfoCollectionPrint.java:63 []
17/04/06 11:06:00 INFO SparkContext: RDD's recursive dependencies: (2) MapPartitionsRDD[48] at
reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 [Memory Serialized 1x Replicated]
| CachedPartitions: 1; MemorySize: 4.0 B; ExternalBlockStoreSize: 0.0 B; DiskSize: 0.0 B
| CoGroupedRDD[47] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 [Memory
Serialized 1x Replicated]
| MapPartitionsRDD[38] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 [Memory
Serialized 1x Replicated]
| CachedPartitions: 2; MemorySize: 8.0 B; ExternalBlockStoreSize: 0.0 B; DiskSize: 0.0 B
| ReliableCheckpointRDD[40] at print at FemaleInfoCollectionPrint.java:123 [Memory Serialized 1x
Replicated]
| ShuffledRDD[36] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 [Memory Serialized
1x Replicated]
| CachedPartitions: 2; MemorySize: 8.0 B; ExternalBlockStoreSize: 0.0 B; DiskSize: 0.0 B
+- (5) MapPartitionsRDD[35] at map at FemaleInfoCollectionPrint.java:81 [Memory Serialized 1x
Replicated]
| MapPartitionsRDD[34] at filter at FemaleInfoCollectionPrint.java:81 [Memory Serialized 1x Replicated]
| MapPartitionsRDD[33] at map at FemaleInfoCollectionPrint.java:72 [Memory Serialized 1x Replicated]
```

```

| MapPartitionsRDD[32] at map at FemaleInfoCollectionPrint.java:63 [Memory Serialized 1x Replicated]
| KafkaRDD[31] at createDirectStream at FemaleInfoCollectionPrint.java:63 [Memory Serialized 1x
Replicated]
| ShuffledRDD[46] at reduceByKeyAndWindow at FemaleInfoCollectionPrint.java:98 [Memory Serialized
1x Replicated]
| CachedPartitions: 1; MemorySize: 4.0 B; ExternalBlockStoreSize: 0.0 B; DiskSize: 0.0 B
+-(5) MapPartitionsRDD[45] at map at FemaleInfoCollectionPrint.java:81 [Memory Serialized 1x
Replicated]
| MapPartitionsRDD[44] at filter at FemaleInfoCollectionPrint.java:81 [Memory Serialized 1x Replicated]
| MapPartitionsRDD[43] at map at FemaleInfoCollectionPrint.java:72 [Memory Serialized 1x Replicated]
| MapPartitionsRDD[42] at map at FemaleInfoCollectionPrint.java:63 [Memory Serialized 1x Replicated]
| KafkaRDD[41] at createDirectStream at FemaleInfoCollectionPrint.java:63 [Memory Serialized 1x
Replicated]

Time: 1491447960000 ms

```

## Answer

In this program, the print operator in DStream is used to display the result. This operator calls the take operator in RDD to implement bottom-layer calculation.

The take operator triggers calculation for multiple times by partition.

In this problem, due to the shuffle operation, the take operator has two partitions by default. Spark Streaming first calculates the first partition. But because there is no data input, the number of obtained results is less than 10. Then the second calculation is triggered. Therefore, the DAG of RDD is printed twice.

To resolve this issue, change the print operator to foreach(collect) in the code.

### 26.8.3.2 What Can I Do If Spark Streaming Tasks Are Blocked?

#### Question

After a Spark Streaming task is run and data is input, no processing result is displayed. Open the web page to view the Spark job execution status. The following figure shows that two jobs are waiting to be executed but cannot be executed successfully.

**Figure 26-18 Active Jobs**

Active Jobs (2)

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total
3	<a href="#">print at test2StreamFromKafka.scala:31</a>	2015/05/25 18:28:55	63.7 h	0/3
2	<a href="#">start at test2StreamFromKafka.scala:34</a>	2015/05/25 18:28:55	63.7 h	0/1

Check the completed jobs. Only two jobs are found, indicating that Spark Streaming does not trigger data computing tasks. (By default, Spark Streaming has two jobs that attempt to run. See the figure below.)

**Figure 26-19 Completed Jobs**

Completed Jobs (2)

Job Id	Description	Submitted	Duration	Stages: Succeeded/Total
1	<a href="#">print at test2StreamFromKafka.scala:31</a>	2015/05/25 18:28:55	0.7 s	2/2 (1 skipped)
0	<a href="#">start at test2StreamFromKafka.scala:34</a>	2015/05/25 18:28:54	1 s	2/2

## Answer

After fault locating, it is found that the number of computing cores of Spark Streaming is less than the number of receivers. As a result, after some receivers are started, no resources are available to run computing tasks. Therefore, the first task keeps waiting and subsequent tasks keep queuing. [Figure 26-18](#) is an example of two queuing tasks.

To address this problem, it is advised to check whether the number of Spark cores is greater than the number of receivers when two tasks are queuing.

### NOTE

Receiver is a permanent Spark job in Spark Streaming. It is common for Spark, but its life cycle is the same as that of a Spark Streaming task and occupies one computing core.

Pay attention to the relationship between the number of cores and the number of receivers in scenarios where default configurations are often used, such as debugging and testing.

### 26.8.3.3 What Should I Pay Attention to When Optimizing Spark Streaming Task Parameters?

#### Question

When Spark Streaming tasks are running, the data processing performance does not improve significantly as the number of executors increases. What should I pay attention to if I perform parameter optimization?

#### Answer

When the number of executor cores is 1, comply with the following rules to optimize Spark Streaming running parameters:

- The Spark task processing speed is related to the number of partitions in Kafka. When the number of partitions is less than the specified number of executors, the number of actually used executors is the same as the number of partitions, and other executors will be idle. Therefore, the number of executors must be less than or equal to the number of partitions.
- When data skew occurs on different partitions of Kafka, the executor corresponding to the partition with a large amount of data touches the glass ceiling of data processing. Therefore, when the Producer program is executed, data is sent to each partition on average to improve the processing speed.
- When partition data is evenly distributed, increasing the number of partitions and executors will improve the Spark processing speed. (When the number of partitions is the same as that of executors, the processing speed is the fastest.)
- When partition data is evenly distributed, ensure that the number of partitions is an integer multiple of the number of executors for proper allocation of resources.

### 26.8.3.4 Why Does the Spark Streaming Application Fail to Be Submitted After the Token Validity Period Expires?

#### Question

Change the validity period of the Kerberos ticket and HDFS token to 5 minutes, set **dfs.namenode.delegation.token.renew-interval** to a value less than 60 seconds, and submit the Spark Streaming application. If the token expires, the error message below is displayed, and the application exits. Why?

```
token (HDFS_DELEGATION_TOKEN token 17410 for spark2x) is expired
```

#### Answer

- Possible causes:

The credential refresh thread of the ApplicationMaster process uploads the updated credential file to the HDFS based on the *token renew period multiplied by 0.75*.

In the executor process, the credential refresh thread obtains the updated credential file from the HDFS based on the time ratio of the *token renewal period multiplied by 0.8* to update the token in UserGroupInformation, preventing the token from being invalid.

When the credential refresh thread of the executor process detects that the current time is later than the credential file update time (*token renew period x 0.8*), it waits for 1 minute and then obtains the latest credential file from the HDFS to ensure that the AM has stored the updated credential file in the HDFS.

When the value of **dfs.namenode.delegation.token.renew-interval** is less than 60 seconds, the started executor detects that the current time is later than the time when the credential file is updated. One minute later, the executor obtains the latest credential file from the HDFS. However, the token is already invalid, and the task fails to be executed. Then, other executor processes retry within 1 minute. The task also fails to run on other executors. As a result, the executors that fail to run are added to the blacklist. If no executors are available, the application exits.

- Solution:

In the Spark application scenario, set **dfs.namenode.delegation.token.renew-interval** to a value greater than 80 seconds. For details about the **dfs.namenode.delegation.token.renew-interval** parameter, see [Table 26-84](#).

**Table 26-84** Parameter description

Parameter	Description	Default Value
dfs.namenode.delegation.token.renew-interval	This parameter is a server parameter. It specifies the maximum lifetime to renew a token. Unit: milliseconds.	86400000

## 26.8.3.5 Why does Spark Streaming Application Fail to Restart from Checkpoint When It Creates an Input Stream Without Output Logic?

### Question

Spark Streaming application creates one input stream without output logic. The application fails to restart from checkpoint and an error will be shown like below:

```
17/04/24 10:13:57 ERROR Utils: Exception encountered
java.lang.NullPointerException
at org.apache.spark.streaming.dstream.DStreamCheckpointData$$anonfun$writeObject$1.applymcVsp(DStreamCheckpointData.scala:125)
at org.apache.spark.streaming.dstream.DStreamCheckpointData$$anonfun$writeObject$1.apply(DStreamCheckpointData.scala:123)
at org.apache.spark.streaming.dstream.DStreamCheckpointData$$anonfun$writeObject$1.apply(DStreamCheckpointData.scala:123)
at org.apache.spark.util.Utils$.tryOrIOException(Utils.scala:1195)
at
org.apache.spark.streaming.dstream.DStreamCheckpointData.writeObject(DStreamCheckpointData.scala:123)
)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:498)
at java.io.ObjectStreamClass.invokeWriteObject(ObjectStreamClass.java:1028)
at java.io.ObjectOutputStream.writeSerialData(ObjectOutputStream.java:1496)
at java.io.ObjectOutputStream.writeOrdinaryObject(ObjectOutputStream.java:1432)
at java.io.ObjectOutputStream.writeObject0(ObjectOutputStream.java:1178)
at java.io.ObjectOutputStream.defaultWriteFields(ObjectOutputStream.java:1548)
at java.io.ObjectOutputStream.defaultWriteObject(ObjectOutputStream.java:441)
at org.apache.spark.streaming.dstream.DStream$$anonfun$writeObject$1.applymcVsp(DStream.scala:515)
at org.apache.spark.streaming.dstream.DStream$$anonfun$writeObject$1.apply(DStream.scala:510)
at org.apache.spark.streaming.dstream.DStream$$anonfun$writeObject$1.apply(DStream.scala:510)
at org.apache.spark.util.Utils$.tryOrIOException(Utils.scala:1195)
at org.apache.spark.streaming.dstream.DStream.writeObject(DStream.scala:510)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:498)
at java.io.ObjectStreamClass.invokeWriteObject(ObjectStreamClass.java:1028)
at java.io.ObjectOutputStream.writeSerialData(ObjectOutputStream.java:1496)
at java.io.ObjectOutputStream.writeOrdinaryObject(ObjectOutputStream.java:1432)
at java.io.ObjectOutputStream.writeObject0(ObjectOutputStream.java:1178)
at java.io.ObjectOutputStream.writeArray(ObjectOutputStream.java:1378)
at java.io.ObjectOutputStream.writeObject0(ObjectOutputStream.java:1174)
at java.io.ObjectOutputStream.defaultWriteFields(ObjectOutputStream.java:1548)
at java.io.ObjectOutputStream.writeSerialData(ObjectOutputStream.java:1509)
at java.io.ObjectOutputStream.writeOrdinaryObject(ObjectOutputStream.java:1432)
at java.io.ObjectOutputStream.writeObject0(ObjectOutputStream.java:1178)
at java.io.ObjectOutputStream.defaultWriteFields(ObjectOutputStream.java:1548)
at java.io.ObjectOutputStream.defaultWriteObject(ObjectOutputStream.java:441)
at org.apache.spark.streaming.DStreamGraph$$anonfun$writeObject$1.applymcVsp(DStreamGraph.scala:191)
at org.apache.spark.streaming.DStreamGraph$$anonfun$writeObject$1.apply(DStreamGraph.scala:186)
at org.apache.spark.streaming.DStreamGraph$$anonfun$writeObject$1.apply(DStreamGraph.scala:186)
at org.apache.spark.util.Utils$.tryOrIOException(Utils.scala:1195)
at org.apache.spark.streaming.DStreamGraph.writeObject(DStreamGraph.scala:186)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:498)
at java.io.ObjectStreamClass.invokeWriteObject(ObjectStreamClass.java:1028)
at java.io.ObjectOutputStream.writeSerialData(ObjectOutputStream.java:1496)
at java.io.ObjectOutputStream.writeOrdinaryObject(ObjectOutputStream.java:1432)
at java.io.ObjectOutputStream.writeObject0(ObjectOutputStream.java:1178)
at java.io.ObjectOutputStream.defaultWriteFields(ObjectOutputStream.java:1548)
at java.io.ObjectOutputStream.writeSerialData(ObjectOutputStream.java:1509)
```

```

at java.io.ObjectOutputStream.writeOrdinaryObject(ObjectOutputStream.java:1432)
at java.io.ObjectOutputStream.writeObject0(ObjectOutputStream.java:1178)
at java.io.ObjectOutputStream.writeObject(ObjectOutputStream.java:348)
at org.apache.spark.streaming.Checkpoint$$anonfun$serialize$1.applymcVsp(Checkpoint.scala:142)
at org.apache.spark.streaming.Checkpoint$$anonfun$serialize$1.apply(Checkpoint.scala:142)
at org.apache.spark.streaming.Checkpoint$$anonfun$serialize$1.apply(Checkpoint.scala:142)
at org.apache.spark.util.Utils$.tryWithSafeFinally(Utils.scala:1230)
at org.apache.spark.streaming.Checkpoint$.serialize(Checkpoint.scala:143)
at org.apache.spark.streaming.StreamingContext.validate(StreamingContext.scala:566)
at org.apache.spark.streaming.StreamingContext.liftedTree1$1(StreamingContext.scala:612)
at org.apache.spark.streaming.StreamingContext.start(StreamingContext.scala:611)
at com.spark.test.kafka08LifoTwoInkfk$.main(kafka08LifoTwoInkfk.scala:21)
at com.spark.test.kafka08LifoTwoInkfk.main(kafka08LifoTwoInkfk.scala)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:498)
at org.apache.spark.deploy.SparkSubmit$.org$apache$spark$deploy$SparkSubmit$
runMain(SparkSubmit.scala:772)
at org.apache.spark.deploy.SparkSubmit$.doRunMain$1(SparkSubmit.scala:183)
at org.apache.spark.deploy.SparkSubmit$.submit(SparkSubmit.scala:208)
at org.apache.spark.deploy.SparkSubmit$.main(SparkSubmit.scala:123)
at org.apache.spark.deploy.SparkSubmit.main(SparkSubmit.scala)

```

## Answer

When Streaming Context starts, DStream checkpoint object of application should be serialized with application set to checkpoint and Dstream context will be used during this serialization.

Dstream.context is the Dstream which Streaming Context relies on to check reversely from output Stream, set the context one by one. If Spark Streaming application creates one input stream which does not have output logic, there will be no context set for the input stream. 'NullPointerException' will be reported during serialization.

Solution: If there is no input logic for the output stream in the application, delete the input stream in the code or add the relevant output logic for that input stream.

### 26.8.3.6 Why Is the Input Size Corresponding to Batch Time on the Web UI Set to 0 Records When Kafka Is Restarted During Spark Streaming Running?

#### Question

When the Kafka is restarted during the execution of the Spark Streaming application, the application cannot obtain the topic offset from the Kafka. As a result, the job fails to be generated. As shown in [Figure 26-20, 2017/05/11 10:57:00-2017/05/11 10:58:00](#) indicates the Kafka restart time. After the restart is successful at 10:58:00 on May,11,2017, the value of **Input Size** is **0 records**.

**Figure 26-20** On the Web UI, the **input size** corresponding to the **batch time** is **0 records**.

Completed Batches (last 9 out of 9)

Batch Time	Input Size	Scheduling Delay (?)	Processing Time (?)	Total Delay (?)	Output Ops: Succeeded/Total
2017/05/11 10:58:50	18 records	0 ms	0.4 s	0.4 s	1/1
2017/05/11 10:58:40	20 records	4 s	0.3 s	4 s	1/1
2017/05/11 10:58:30	20 records	14 s	0.5 s	14 s	1/1
2017/05/11 10:58:20	20 records	23 s	0.4 s	24 s	1/1
2017/05/11 10:58:10	20 records	33 s	0.5 s	33 s	1/1
2017/05/11 10:58:00	0 records	6 ms	43 s	43 s	1/1
2017/05/11 10:57:00	19 records	1 ms	0.9 s	0.9 s	1/1
2017/05/11 10:56:50	20 records	1 ms	0.6 s	0.6 s	1/1
2017/05/11 10:56:40	28 records	13 ms	5 s	5 s	1/1

## Answer

After Kafka is restarted, the application supplements the missing RDD between 10:57:00 on May 11, 2017 and 10:58:00 on May 11, 2017 based on the batch time. Although the number of read data records displayed on the UI is **0**, the missing data is processed in the supplemented RDD. Therefore, no data loss occurs.

The data processing mechanism during the Kafka restart period is as follows:

The Spark Streaming application uses the **state** function (for example, **updateStateByKey**). After Kafka is restarted, the Spark Streaming application generates a batch task at 10:58:00 on May 11, 2017. The missing RDD between 10:57:00 on May 11, 2017 and 10:58:00 on May 11, 2017 is supplemented based on the batch time (data that is not read in Kafka before Kafka restart, which belongs to the batch before 10:57:00 on May 11, 2017).

## 26.8.4 Why the Job Information Obtained from the restful Interface of an Ended Spark Application Is Incorrect?

### Question

The job information obtained from the restful interface of an ended Spark application is incorrect: the value of **numActiveTasks** is negative, as shown in [Figure 26-21](#):

Figure 26-21 job information

```
[{
 "jobId" : 0,
 "name" : "reduce at SparkPi.scala:36",
 "submissionTime" : "2016-05-28T09:35:34.415GMT",
 "completionTime" : "2016-05-28T09:35:35.686GMT",
 "stageIds" : [0],
 "status" : "SUCCEEDED",
 "numTasks" : 2,
 "numActiveTasks" : -1,
 "numCompletedTasks" : 2,
 "numSkippedTasks" : 2,
 "numFailedTasks" : 0,
 "numActiveStages" : 0,
 "numCompletedStages" : 1,
 "numSkippedStages" : 0,
 "numFailedStages" : 0
}]
```

 NOTE

numActiveTasks indicates the number of active tasks.

## Answer

The job information can be obtained in either of the following methods:

- Set **spark.history.briefInfo.gather=true** and then view the brief JobHistory information.
- Visit the JobHistory2x page of Spark (URL: <https://IP:port/api/v1/<appid>/jobs/>).

The value of **numActiveTasks** in the job information is calculated from the difference between the number of SparkListenerTaskStart events and the number of SparkListenerTaskEnd events in the **eventLog** file. If some events are not recorded in the **eventLog** file, the job information obtained from the restful interface is incorrect.

## 26.8.5 Why Cannot I Switch from the Yarn Web UI to the Spark Web UI?

### Question

In FusionInsight, the Spark application is run in yarn-client mode on the client. The following error occurs during the switch from the Yarn web UI to the application web UI:

Error Occurred.

Problem accessing /proxy/application\_ /

Powered by Jetty://

The YARN ResourceManager log shows the following information:



```
2016-07-21 16:35:27,099 | INFO | Socket Reader #1 for port 8032 | Auth successful for mapred/
hadoop.<System domain name>@<System domain name> (auth:KERBEROS) | Server.java:1388
2016-07-21 16:35:27,105 | INFO | 1526016381@qtp-1178290888-1015 | admin is accessing unchecked
http://10.120.169.53:23011 which is the app master GUI of
application_1468986660719_0045 owned by spark | WebAppProxyServlet.java:393
2016-07-21 16:36:02,843 | INFO | Socket Reader #1 for port 8032 | Auth successful for hive/
hadoop.<System domain name>@<System domain name> (auth:KERBEROS) | Server.java:1388
2016-07-21 16:36:02,851 | INFO | Socket Reader #1 for port 8032 | Auth successful for hive/
hadoop.<System domain name>@<System domain name> (auth:KERBEROS) | Server.java:1388
2016-07-21 16:36:12,163 | WARN | 1526016381@qtp-1178290888-1015 | /proxy/
application_1468986660719_0045/: java.net.ConnectException: Connection timed out |
Slf4jLog.java:76
2016-07-21 16:37:03,918 | INFO | Socket Reader #1 for port 8032 | Auth successful for hive/
hadoop.<System domain name>@<System domain name> (auth:KERBEROS) | Server.java:1388
2016-07-21 16:37:03,926 | INFO | Socket Reader #1 for port 8032 | Auth successful for hive/
hadoop.<System domain name>@<System domain name> (auth:KERBEROS) | Server.java:1388
2016-07-21 16:37:11,956 | INFO | AsyncDispatcher event handler | Updating application attempt
appattempt_1468986660719_0045_000001 with final state: FINISHING,
and exit status: -1000 | RMAAppAttemptImpl.java:1253
```

## Answer

On FusionInsight Manager, the IP address of the Yarn service is in the 192 network segment.

In Yarn logs, the IP address of Spark web UI read by Yarn is `http://10.120.169.53:23011`, which is in the 10 network segment. The IP addresses in the 192 network segment cannot communicate with those in the 10 network segment. As a result, the Spark web UI fails to be accessed.

Solution:

Log in to the client whose IP address is **10.120.169.53** and change the IP address in the `/etc/hosts` file to the IP address in the 192 network segment. Run the Spark application again. The Spark web UI is displayed.

## 26.8.6 What Can I Do If an Error Occurs when I Access the Application Page Because the Application Cached by HistoryServer Is Recycled?

### Question

An error occurs when I access a Spark application page on the HistoryServer page.

Check the HistoryServer logs. The "FileNotFoundException" exception is found. The related logs are as follows:

```
2016-11-22 23:58:03,694 | WARN | [qtp55429210-232] | /history/application_1479662594976_0001/stages/
stage/ | org.sparkproject.jetty.servlet.ServletHandler.doHandle(ServletHandler.java:628)
java.io.FileNotFoundException: ${BIGDATA_HOME}/tmp/spark/jobHistoryTemp/
blockmgr-5f1f6aca-2303-4290-9845-88fa94d78480/09/temp_shuffle_11f82aaf-e226-46dc-
b1f0-002751557694 (No such file or directory)
```

### Answer

If a Spark application with a large number of tasks is run on the HistoryServer page, the memory overflows to disk and files with the `temp_shuffle` prefix are generated.

By default, HistoryServer caches 50 Spark applications (determined by the `spark.history.retainedApplications` configuration item). When the number of

Spark applications in the memory exceeds 50, HistoryServer reclaims the first cached Spark application and clears the corresponding **temp\_shuffle** file.

When a user is viewing Spark applications to be recycled, the **temp\_shuffle** file may not be found. As a result, the current page cannot be accessed.

If the preceding problem occurs, use either of the following methods to solve the problem:

- Access the HistoryServer page of the Spark application again. The correct page information is displayed.
- If more than 50 Spark applications need to be accessed at the same time, increase the value of **spark.history.retainedApplications**.

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Spark2x** > **Configuration**, and click **All Configurations**. In the navigation tree on the left, choose **JobHistory2x** > **GUI**, and set parameters.

**Table 26-85** Parameter description

Parameter	Description	Default Value
spark.history.retainedApplications	Number of Spark applications cached by HistoryServer. When the number of applications to be cached exceeds the value of this parameter, HistoryServer reclaims the first cached Spark application.	50

## 26.8.7 Why Is not an Application Displayed When I Run the Application with the Empty Part File?

### Question

When I run an application with an empty part file in HDFS with the log grouping function enabled, why is not the application displayed on the homepage of JobHistory?

### Answer

On the JobHistory page, information about applications is updated only with changed sizes of part files in HDFS. If a file is read for the first time, its size is compared with 0. The file is read only when the file size is greater than 0.

When the log grouping function is enabled, if the application you run does not have jobs in running status, the part file is empty. As a result, JobHistory does not read the part file and the application information is not displayed on the JobHistory page. However, if the size of part file is changed later, the application will be displayed on JobHistory.

## 26.8.8 Why Does Spark2x Fail to Export a Table with the Same Field Name?

### Question

The following code fails to be executed on spark-shell of Spark2x:

```
val acctId = List(("49562", "Amal", "Derry"), ("00000", "Fred", "Xanadu"))
val rddLeft = sc.makeRDD(acctId)
val dfLeft = rddLeft.toDF("Id", "Name", "City")
//dfLeft.show
val acctCustId = List(("Amal", "49562", "CO"), ("Dave", "99999", "ZZ"))
val rddRight = sc.makeRDD(acctCustId)
val dfRight = rddRight.toDF("Name", "CustId", "State")
//dfRight.show
val dfJoin = dfLeft.join(dfRight, dfLeft("Id") === dfRight("CustId"), "outer")
dfJoin.show
dfJoin.repartition(1).write.format("com.databricks.spark.csv").option("delimiter", "\t").option("header", "true").option("treatEmptyValuesAsNulls", "true").option("nullValue", "").save("/tmp/outputDir")
```

### Answer

In Spark2x, the duplicate field name of the **join** statement is checked. You need to modify the code to ensure that no duplicate field exists in the saved data.

## 26.8.9 Why JRE fatal error after running Spark application multiple times?

### Question

Why JRE fatal error after running Spark application multiple times?

### Answer

When you run Spark application multiple times, JRE fatal error occurs and this is due to the problem with the Linux Kernel.

To resolve this issue, upgrade the **kernel version to 4.13.9-2.ge7d7106-default**.

## 26.8.10 Native Spark2x UI Fails to Be Accessed or Is Incorrectly Displayed when Internet Explorer Is Used for Access

### Question

When Internet Explorer 9, 10, or 11 is used to access the native Spark2x UI, the access fails or the page is incorrectly displayed.

### Symptom

Internet Explorer fails to access the native Spark2x UI.



Turn on TLS 1.0, TLS 1.1, and TLS 1.2 in Advanced settings and try connecting to

## Cause

Some versions of Internet Explorer 9, 10, and 11 fail to process SSL handshakes.

## Solution

Use Google Chrome 71 or later for access.

## 26.8.11 How Does Spark2x Access External Cluster Components?

### Question

There are two clusters, cluster 1 and cluster 2. How do I use Spark2x in cluster 1 to access HDFS, Hive, HBase, and Kafka components in cluster 2?

### Answer

1. Components in two clusters can access each other. However, there are the following restrictions:
  - Only one Hive MetaStore can be accessed. Specifically, Hive MetaStore in cluster 1 and Hive MetaStore in cluster 2 cannot be accessed at the same time.
  - User systems in different clusters are not synchronized. When users access components in another cluster, user permission is determined by the user configuration of the peer cluster. For example, if user A of cluster 1 does not have the permissions to access the HBase meta table in cluster 1 but user A of cluster 2 can access the HBase meta table in cluster 2, user A of cluster 1 can access the HBase meta table in cluster 2.
  - To enable components in a security cluster to communicate with each other across Manager, you need to configure mutual trust.
2. The following describes how to access Hive, HBase, and Kafka components in cluster 2 as user A.

#### NOTE

The following operations are based on the scenario where a user uses the FusionInsight client to submit the Spark2x application. If the user uses the configuration file directory, the user needs to modify the corresponding file in the configuration directory of the application and upload the configuration file to the executor.

When the HDFS and HBase clients access the server, **hostname** is used to configure the server address. Therefore, the hosts configuration of all nodes to be accessed must be saved in the **/etc/hosts** file on the client. You can add the host of the peer cluster node to the **/etc/hosts** file of the client node in advance.

- Access Hive metastore: Replace the **hive-site.xml** file in the **conf** directory of the Spark2x client in cluster 1 with the **hive-site.xml** file in the **conf** directory of the Spark2x client in cluster 2.

After the preceding operations are performed, you can use Spark SQL to access Hive MetaStore. To access Hive table data, you need to perform the operations in **Access HDFS of two clusters at the same time**: and set **nameservice** of the peer cluster to **LOCATION**.

- Access HBase of the peer cluster.
  - i. Configure the IP addresses and host names of all ZooKeeper nodes and HBase nodes in cluster 2 in the `/etc/hosts` file on the client node of cluster 1.
  - ii. Replace the `hbase-site.xml` file in the `conf` directory of the Spark2x client in cluster 1 with the `hbase-site.xml` file in the `conf` directory of the Spark2x client in cluster 2.
- Access Kafka: Set the address of the Kafka Broker to be accessed to the Kafka Broker address in cluster 2.
- Access HDFS of two clusters at the same time:
  - Two tokens with the same NameService cannot be obtained at the same time. Therefore, the NameServices of the HDFS in two clusters must be different. For example, one is `hacluster`, and the other is `test`.

- 1) Obtain the following configurations from the `hdfs-site.xml` file of cluster2 and add them to the `hdfs-site.xml` file in the `conf` directory of the Spark2x client in cluster1:

**dfs.nameservices.mappings, dfs.nameservices, dfs.namenode.rpc-address.test.\*, dfs.ha.namenodes.test, and dfs.client.failover.proxy.provider.test**

The following is an example:

```
<property>
<name>dfs.nameservices.mappings</name>
<value>[{"name":"hacluster","roleInstances":["14","15"]},
{"name":"test","roleInstances":["16","17"]}]</value>
</property>
<property>
<name>dfs.nameservices</name>
<value>hacluster,test</value>
</property>
<property>
<name>dfs.namenode.rpc-address.test.16</name>
<value>192.168.0.1:8020</value>
</property>
<property>
<name>dfs.namenode.rpc-address.test.17</name>
<value>192.168.0.2:8020</value>
</property>
<property>
<name>dfs.ha.namenodes.test</name>
<value>16,17</value>
</property>
<property>
<name>dfs.client.failover.proxy.provider.test</name>
<value>org.apache.hadoop.hdfs.server.namenode.ha.ConfiguredFailoverProxyProvider
</value>
</property>
```

- 2) Modify `spark.yarn.extra.hadoopFileSystems = hdfs://test` and `spark.hadoop.hdfs.externalToken.enable = true` in the `spark-defaults.conf` configuration file under the `conf` directory on the Spark client of cluster 1.
 

```
spark.yarn.extra.hadoopFileSystems = hdfs://test
spark.hadoop.hdfs.externalToken.enable = true
```
- 3) In the application submission command, add the `--keytab` and `--principal` parameters and set them to the user who submits the task in cluster1.

- 4) Use the Spark client of cluster1 to submit the application. Then, the two HDFS services can be accessed at the same time.
- Access HBase of two clusters at the same time:
    - i. Modify **spark.hadoop.hbase.externalToken.enable = true** in the **spark-defaults.conf** configuration file under the **conf** directory on the Spark client of cluster 1.

```
spark.hadoop.hbase.externalToken.enable = true
```
    - ii. When accessing HBase, you need to use the configuration file of the corresponding cluster to create a **Configuration** object for creating a **Connection** object.
    - iii. In an MRS cluster, tokens of multiple HBase services can be obtained at the same time to solve the problem that the executor cannot access HBase. The method is as follows:

Assume that you need to access HBase of the current cluster and HBase of cluster2. Save the **hbase-site.xml** file of cluster2 in a compressed package named **external\_hbase\_conf\*\*\***, and use **--archives** to specify the compressed package when submitting the command.

## 26.8.12 Why Does the Foreign Table Query Fail When Multiple Foreign Tables Are Created in the Same Directory?

### Question

Assume there is a data file path named **/test\_data\_path**. User A creates a foreign table named **tableA** for the directory, and user B creates a foreign table named **tableB** for the directory. When user B performs the insert operation on **tableB**, user A fails to query data using **tableA** and the error "Permission denied" is displayed.

### Answer

After user B performs the insert operation on **tableB**, a new data file is generated in the foreign table path and the file belongs to user B. When user A queries data using **tableA**, all files in the foreign table directory are read. In this case, the query fails because user A does not have the read permissions on the file generated by user B.

This problem also occurs in other scenarios. For example, the **inset overwrite** operation will also duplicate other table files in this directory.

Due to the Spark SQL implementation mechanism, check restrictions in this scenario will lead to inconsistency and performance deterioration. Therefore, no restriction is added in this scenario, and this method is not recommended.

## 26.8.13 Why Is the Native Page of an Application in Spark2x JobHistory Displayed Incorrectly?

### Question

Submit a Spark application that contains millions of tasks in a single job. After the application is complete, when I access the native page of the application in

JobHistory, the browser waits for a long time before the page is displayed. If the native page cannot be displayed within 10 minutes, "Proxy Error" is displayed.

**Figure 26-22** Example error information

### Proxy Error

```
The proxy server received an invalid response from an upstream server.
The proxy server could not handle the request GET /Spark2x/JobHistory2x/77/history/application [redacted] /1/jobs/
Reason: Error reading from remote server
```

## Answer

When you switch to the native page of an application in JobHistory, JobHistory needs to replay the event logs of the application. If the application contains a large number of event logs, the replay takes a long time and the browser waits for a long time.

When the browser accesses the native page in JobHistory, the httpd proxy is required. The timeout interval of the proxy is 10 minutes. Therefore, if JobHistory cannot parse and return the event logs within 10 minutes, httpd returns the "Proxy Error" message to the browser.

## Solution

The local disk caching function is enabled for JobHistory. When an application is accessed, the parsing result of the event logs of the application is cached to the local disk. When the application is accessed for the second time, the response speed is significantly improved. In this case, you only need to wait for a while and access the original link again.

## 26.8.14 Why Do I Fail to Create a Table in the Specified Location on OBS After Logging to spark-beeline?

### Question

When the OBS ECS/BMS image cluster is connected, after spark-beeline is logged in, an error is reported when a location is specified to create a table on OBS.

**Figure 26-23** Error message

```
de-master2qCKJ:22550/> create database sparkdb location 'obs://800mrs/sparktest/sparkdb';

0.626 seconds)
de-master2qCKJ:22550/> use sparkdb;

0.072 seconds)
de-master2qCKJ:22550/> create table orc (id int,name string) using orc;
Exception: Configuration problem with provider path. (state=,code=0)
```

## Answer

The permission on the `ssl.jceks` file in HDFS is insufficient. As a result, the table fails to be created.

```
Caused by: org.apache.hadoop.security.AccessControlException: Permission denied: user=root, access=READ, inode="/user/spark2x/jars/8.0.2/ssl.jceks":spark2x:hadoop:-rw-----
at org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.check(FSPermissionChecker.java:410)
at org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkPermission(FSPermissionChecker.java:264)
at com.hadoop.adapter.hdfs.plugin.HWAccessControlEnforce.checkPermission(.....java:54)
at org.apache.hadoop.hdfs.server.namenode.FSPermissionChecker.checkPermission(FSPermissionChecker.java:194)
at org.apache.hadoop.hdfs.server.namenode.FSDirectory.checkPermission(FSDirectory.java:1857)
at org.apache.hadoop.hdfs.server.namenode.FSDirectory.checkPathAccess(FSDirectory.java:1841)
at org.apache.hadoop.hdfs.server.namenode.FSDirectory.checkPathAccess(FSDirectory.java:1891)
at org.apache.hadoop.hdfs.server.namenode.FSDirStatAndListingOp.getLocations(FSDirStatAndListingOp.java:175)
at org.apache.hadoop.hdfs.server.namenode.FSNamesystem.getBlockLocations(FSNamesystem.java:1950)
at org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.getBlockLocations(NameNodeRpcServer.java:762)
at org.apache.hadoop.hdfs.protocolPB.ClientNameNodeProtocolServerSideTranslatorPB.getBlockLocations(ClientNameNodeProtocolServerSideTranslatorPB.java:445)
at org.apache.hadoop.hdfs.protocol.proto.ClientNameNodeProtocolProtosClientNameNodeProtocol$2.callBlockingMethod(ClientNameNodeProtocolProtos.java)
at org.apache.hadoop.ipc.ProtobufRpcEngine2ProtoBufRpcInvoker.call(ProtobufRpcEngine.java:528)
at org.apache.hadoop.ipc.RPCServer.call(RPC.java:1036)
at org.apache.hadoop.ipc.Server$RpcCall.run(Server.java:985)
at org.apache.hadoop.ipc.Server$RpcCall.run(Server.java:913)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:422)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1737)
at org.apache.hadoop.ipc.Server$Handler.run(Server.java:1287)
```

## Solution

1. Log in to the node where Spark2x resides as user **omm** and run the following command:  
**vi \${BIGDATA\_HOME}/FusionInsight\_Spark2x\_8.1.0.1/install/FusionInsight-Spark2x-\*/spark/sbin/fake\_prestart.sh**
2. Change **eval "\${hdfsCmd}" -chmod 600 "\${InnerHdfsDir}"/ssl.jceks >> "\${PRESTART\_LOG}" 2>&1** to **eval "\${hdfsCmd}" -chmod 644 "\${InnerHdfsDir}"/ssl.jceks >> "\${PRESTART\_LOG}" 2>&1**.
3. Restart the SparkResource instance.

## 26.8.15 Spark Shuffle Exception Handling

### Question

In some scenarios, the following exception occurs in the Spark shuffle phase:

```
2021-06-18 02:53:08.304 INFO [shuffle-server-6-1] | 01659741:Unmatched MACs | java.security.sasl.unwrap(DigestMD5Base.java:148)
2021-06-18 02:53:08.368 WARN [shuffle-server-6-1] | Exception in connection from /0000000000000000 | org.apache.spark.network.server.TransportChannelHandler.exceptionCaught(TransportChannelHandler.java:29)
io.netty.handler.codec.DecoderException: java.security.sasl.SaslException: DIGEST-MD5: Out of order sequencing of messages from server. Got: 16 Expected: 14
at io.netty.handler.codec.MessageToMessageDecoder.channelRead(MessageToMessageDecoder.java:98)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:365)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelRead(AbstractChannelHandlerContext.java:379)
at org.apache.spark.network.util.TransportFrameDecoder.channelRead(TransportFrameDecoder.java:102)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:365)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.DefaultChannelPipeline$HeadContext.channelRead(DefaultChannelPipeline.java:1410)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:365)
at io.netty.channel.DefaultChannelPipeline.fireChannelRead(DefaultChannelPipeline.java:910)
at io.netty.channel.nio.NioEventLoop.processSelectedKey(NioEventLoop.java:163)
at io.netty.channel.nio.NioEventLoop.processSelectedKey(NioEventLoop.java:714)
at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:500)
at io.netty.util.concurrent.SingleThreadEventExecutor$5.run(SingleThreadEventExecutor.java:989)
at io.netty.util.concurrent.FastThreadLocalRunnable.run(FastThreadLocalRunnable.java:30)
at java.lang.Thread.run(Thread.java:748)
Caused by: java.security.sasl.SaslException: DIGEST-MD5: Out of order sequencing of messages from server. Got: 16 Expected: 14
at com.sun.security.sasl.digest.DigestMD5Base$DigestPrivacy.unwrap(DigestMD5Base.java:1489)
at org.apache.spark.network.sasl.SparkSaslServer.unwrap(SparkSaslServer.java:149)
at com.sun.security.sasl.digest.DigestMD5Base.unwrap(DigestMD5Base.java:213)
at org.apache.spark.network.sasl.SaslEncryptionHandler.decode(SaslEncryption.java:126)
at org.apache.spark.network.sasl.SaslEncryptionHandler.decode(SaslEncryption.java:101)
at io.netty.handler.codec.MessageToMessageDecoder.channelRead(MessageToMessageDecoder.java:88)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:365)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelRead(AbstractChannelHandlerContext.java:379)
at org.apache.spark.network.util.TransportFrameDecoder.channelRead(TransportFrameDecoder.java:102)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:365)
at io.netty.channel.AbstractChannelHandlerContext.fireChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.DefaultChannelPipeline$HeadContext.channelRead(DefaultChannelPipeline.java:1410)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:379)
at io.netty.channel.AbstractChannelHandlerContext.invokeChannelRead(AbstractChannelHandlerContext.java:365)
at io.netty.channel.DefaultChannelPipeline.fireChannelRead(DefaultChannelPipeline.java:910)
at io.netty.channel.nio.NioEventLoop.processSelectedKey(NioEventLoop.java:163)
at io.netty.channel.nio.NioEventLoop.processSelectedKey(NioEventLoop.java:714)
at io.netty.channel.nio.NioEventLoop.run(NioEventLoop.java:500)
at io.netty.util.concurrent.SingleThreadEventExecutor$5.run(SingleThreadEventExecutor.java:989)
at io.netty.util.concurrent.FastThreadLocalRunnable.run(FastThreadLocalRunnable.java:30)
at java.lang.Thread.run(Thread.java:748)
```

## Solution

For JDBC:

Log in to FusionInsight Manager, change the value of the JDBCServer parameter **spark.authenticate.enableSaslEncryption** to **false**, and restart the corresponding instance.

For client jobs:

When the client submits the application, change the value of **spark.authenticate.enableSaslEncryption** in the **spark-defaults.conf** file to **false**.



# 27 Using Sqoop

---

## 27.1 Using Sqoop from Scratch

Sqoop is an open-source tool for transferring data between Hadoop (Hive) and traditional databases (such as MySQL and PostgreSQL). It can transfer data from a relational database (such as MySQL, Oracle, and PostgreSQL) to HDFS of Hadoop and the other way around.

### Prerequisites

- You have selected the Sqoop component when creating a cluster of MRS 3.1.0 or later.
- You have installed the client. For details, see [Installing a Client \(MRS 3.x or Later\)](#). For example, the installation directory of the client is `/opt/client`. The client directory in the following operations is an example. Change it to the actual installation directory.
- The corresponding driver package (for example, MySQL driver package `mysql-connector-java-5.1.47.jar`) exists in the client directory `/Sqoop/sqoop/lib`. If it does not exist, download it by referring to [3 in Adapting Sqoop 1.4.7 to MRS 3.x Clusters](#).

### Exporting Data From HDFS to MySQL Using the `sqoop export` Command

**Step 1** Log in to the node where the client is located.

**Step 2** Run the following command to initialize environment variables:

```
source /opt/client/bigdata_env
```

**Step 3** Run the following command to operate the Sqoop client:

```
sqoop export --connect jdbc:mysql://10.100.231.134:3306/test --username root --password xxxxxx --table component13 -export-dir hdfs://hacluster/user/hive/warehouse/component_test3 --fields-terminated-by ',' -m 1
```

**Table 27-1** Parameter description

Parameter	Description
-direct	Imports data to a relational database using a database import tool, for example, mysqlimport of MySQL, more efficient than the JDBC connection mode.
-export-dir <dir>	Specifies the source directory for storing data in the HDFS.
-m or -num-mappers <n>	Starts <i>n</i> (4 by default) maps to import data concurrently. The value cannot be greater than the maximum number of maps in a cluster.
-table <table-name>	Specifies the relational database table to be imported.
-update-key <col-name>	Specifies the column used for updating the existing data in a relational database.
-update-mode <mode>	Specifies how updates are performed. The value can be <b>updateonly</b> or <b>allowinsert</b> . This parameter is used only when the relational data table does not contain the data record to be imported. For example, if the HDFS data to be imported to the destination table contains a data record <b>id=1</b> and the table contains an existing data record <b>id=2</b> , the update will fail.
-input-null-string <null-string>	This parameter is optional. If it is not specified, <b>null</b> will be used.
-input-null-non-string <null-string>	This parameter is optional. If it is not specified, <b>null</b> will be used.
-staging-table <staging-table-name>	Creates a table with the same data structure as the destination table for storing data before it is imported to the destination table.  This parameter ensures the transaction security when data is imported to a relational database table. Due to multiple transactions during an import, this parameter can prevent other transactions from being affected when one transaction fails. For example, the imported data is incorrect or duplicate records exist.
-clear-staging-table	Clears data in the staging table before data is imported if the staging-table is not empty.

----End

## Importing Data from MySQL to Hive Using the sqoop import Command

**Step 1** Log in to the node where the client is located.

**Step 2** Run the following command to initialize environment variables:

```
source /opt/client/bigdata_env
```

**Step 3** Run the following command to operate the Sqoop client:

```
sqoop import --connect jdbc:mysql://10.100.231.134:3306/test --username root --password xxxxxx --table component --hive-import --hive-table component_test2 --delete-target-dir --fields-terminated-by "," -m 1 --as-textfile
```

**Table 27-2** Parameter description

Parameter	Description
-append	Appends data to an existing dataset in the HDFS. Once this parameter is used, Sqoop imports data to a temporary directory, renames the temporary file where the data is stored, and moves the file to a formal directory to avoid duplicate file names in the directory.
-as-avrodatafile	Imports data to a data file in the Avro format.
-as-sequencefile	Imports data to a sequence file.
-as-textfile	Import data to a text file. After the text file is generated, you can run SQL statements in Hive to query the result.
-boundary-query <statement>	Specifies the SQL statement for performing boundary query. Before importing data, use a SQL statement to obtain a result set and import the data in the result set. The data format can be <b>-boundary-query 'select id,creationdate from person where id = 3'</b> (indicating a data record whose ID is 3) or <b>select min(&lt;split-by&gt;), max(&lt;split-by&gt;) from &lt;table name&gt;</b> . The fields to be queried cannot contain fields whose data type is string. Otherwise, the error message "java.sql.SQLException: Invalid value for getLong()" is displayed.
-columns<col,col,col...>	Specifies the fields to be imported. The format is - <b>Column id,Username</b> .
-direct	Imports data to a relational database using a database import tool, for example, mysqlimport of MySQL, more efficient than the JDBC connection mode.
-direct-split-size	Splits the imported streams by byte. Especially when data is imported from PostgreSQL using the direct mode, a file that reaches the specified size can be divided into several independent files.
-inline-lob-limit	Sets the maximum value of an inline LOB.
-m or -num-mappers	Starts <i>n</i> (4 by default) maps to import data concurrently. The value cannot be greater than the maximum number of maps in a cluster.

Parameter	Description
-query, -e<statement>	Imports data from the query result. To use this parameter, you must specify the <b>-target-dir</b> and <b>-hive-table</b> parameters and use the query statement containing the WHERE clause as well as \$CONDITIONS. Example: <b>-query'select * from person where \$CONDITIONS' -target-dir /user/hive/warehouse/person -hive-table person</b>
-split-by<column-name>	Specifies the column of a table used to split work units. Generally, the column name is followed by the primary key ID.
-table <table-name>	Specifies the relational database table from which data is obtained.
-target-dir <dir>	Specifies the HDFS path.
-warehouse-dir <dir>	Specifies the directory for storing data to be imported. This parameter is applicable when data is imported to HDFS but cannot be used when you import data to Hive directories. This parameter cannot be used together with <b>-target-dir</b> .
-where	Specifies the WHERE 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	Specifies the Hadoop compression codec. GZIP is used by default.
-null-string <null-string>	Specifies the string to be interpreted as <b>NULL</b> for string columns.
-null-non-string<null-string>	Specifies the string to be interpreted as null for non-string columns. If this parameter is not specified, <b>NULL</b> will be used.
-check-column (col)	Specifies the column for checking incremental data import, for example, <b>id</b> .
-incremental (mode) append or last modified	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> .

Parameter	Description
-last-value (value)	Specifies the maximum value (greater than the specified value) of the column after the last import. This parameter can be set as required.

----End

## Sqoop Usage Example

- Importing data from MySQL to HDFS using the **sqoop import** command  
**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 data from OBS to MySQL using the **sqoop export** command  
**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 data from MySQL to OBS using the **sqoop import** command  
**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 data from MySQL to OBS tables outside Hive  
**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**

## 27.2 Adapting Sqoop 1.4.7 to MRS 3.x Clusters

Sqoop is a tool designed for efficiently transmitting a large amount of data between Apache Hadoop and structured databases (such as relational databases). Customers need to use Sqoop to migrate data in MRS. However, MRS of an earlier version does not provide Sqoop. This section describes how to install and use Sqoop. In MRS 3.1.0 or later, you can select the Sqoop component during cluster creation.

### Prerequisites

The MRS client and the JDK environment have been installed.

```
2021-04-08 10:03:55,018 INFO metastore.HiveMetaStore
[root@node-master1fKEj bin]# echo $JAVA_HOME
/opt/Bigdata/client/JDK/jdk1.8.0_242
```

## Procedure

**Step 1** **Download** the open-source `sqoop-1.4.7.bin__hadoop-2.6.0.tar.gz` package.

**Step 2** Save the downloaded package to the `/opt/Bigdata/client` directory on the node where the MRS client is installed and decompress it.

```
tar zxvf sqoop-1.4.7.bin__hadoop-2.6.0.tar.gz
```


**Step 3** Download the MySQL JDBC driver `mysql-connector-java-xxx.jar` from the MySQL official website. For details about how to select the MySQL JDBC driver, see the following table.

**Table 27-3** Version information

JDBC Driver Version	MySQL Version
Connector/J 5.1	MySQL 4.1, MySQL 5.0, MySQL 5.1, and MySQL 6.0 alpha
Connector/J 5.0	MySQL 4.1, MySQL 5.0 servers, and distributed transaction (XA)
Connector/J 3.1	MySQL 4.1, MySQL 5.0 servers, and MySQL 5.0 except distributed transaction (XA)
Connector/J 3.0	MySQL 3.x and MySQL 4.1


**Step 4** Put the MySQL driver package in the `/opt/Bigdata/client/sqoop-1.4.7.bin__hadoop-2.6.0/lib` directory of Sqoop and modify the owner group and permission of the JAR package. For details, see the owner group and permission of `omm:wheel` and `755` in [Figure 27-1](#).

**Figure 27-1** Owner group and permission of the MySQL driver package



**Step 5** Replace the JAR package in the `lib` directory of Sqoop with that starting with `jackson` in the `lib` directory of Hive on the MRS client, for example, `/opt/Bigdata/client/Hive/Beeline/lib`.

**Figure 27-2** JAR package starting with `jackson`



**Step 6** Copy the `jline` package from the `/opt/Bigdata/client/Hive/Beeline/lib` directory of the MRS Hive client to the `lib` directory of Sqoop.

**Step 7** Run the `vim $JAVA_HOME/jre/lib/security/java.policy` command to add the following configuration:

**permission javax.management.MBeanTrustPermission "register";**

**Step 8** Run the following commands to go to the **conf** directory of the Sqoop and add the configuration items of variables:


```
cd /opt/Bigdata/client/sqoop-1.4.7.bin__hadoop-2.6.0/conf
```

```
cp sqoop-env-template.sh sqoop-env.sh
```

**Step 9** Run the **vim sqoop-env.sh** command to set the environment variables of Sqoop. Change the Hadoop and Hive directories as required.

```
export HADOOP_COMMON_HOME=/opt/Bigdata/client/HDFS/hadoop
export HADOOP_MAPRED_HOME=/opt/Bigdata/client/HDFS/hadoop
export HIVE_HOME=/opt/Bigdata/MRS_1.9.X/install/FusionInsight-Hive-3.1.0/hive (Enter the actual path.)
export HIVE_CONF_DIR=/opt/Bigdata/client/Hive/config
export HCAT_HOME=/opt/Bigdata/client/Hive/HCatalog
```

**Figure 27-3** Setting environment variables of Sqoop



**Step 10** Build the sqoop script. For example:

```
/opt/Bigdata/FusionInsight_Current/1_19_SqoopClient/install/FusionInsight-Sqoop-1.4.7/bin/sqoop import
--connect jdbc:mysql://192.168.0.183:3306/test
--driver com.mysql.jdbc.Driver
--username 'root'
--password 'xxx'
--query "SELECT id, name FROM tbtest WHERE \$CONDITIONS"
--hcatalog-database default
--hcatalog-table test
--num-mappers 1
```

----End

## 27.3 Common Sqoop Commands and Parameters

### Common Sqoop commands

**Table 27-4** Common Sqoop commands

Command	Description
import	Imports data to a cluster.
export	Exports data of a cluster.
codegen	Obtains data from a table in the database to generate a Java file and compress the file.
create-hive-table	Creates a Hive table.

Command	Description
eval	Executes a SQL statement and view the result.
import-all-tables	Imports all tables in a database to HDFS.
job	Generates a Sqoop job.
list-databases	Lists database names.
list-tables	List table names.
merge	Merges data in different HDFS directories and saves the data to a specified directory.
metastore	Starts the metadata database to record the metadata of a Sqoop job.
help	Prints help information.
version	Prints the version information.

## Common Parameters

**Table 27-5** Common parameters

Category	Parameter	Description
Parameters for database connection	--connect	Specifies the URL for connecting to a relational database.
	--connection-manager	Specifies the connection manager class.
	--driver jdbc	Specifies the driver package for database connection.
	--help	Prints help information.
	--password	Specifies the password for connecting to a database.
	--username	Specifies the username for connecting to a database.
	--verbose	Prints detailed information on the console.
import parameters	--fields-terminated-by	Specifies the field delimiter, which must be the same as that in a Hive table or HDFS file.
	--lines-terminated-by	Specifies the line delimiter, which must be the same as that in a Hive table or HDFS file.



Category	Parameter	Description
	--mysql-delimiters	Specifies the default delimiter settings of MySQL.
export parameters	--input-fields-terminated-by	Specifies the field delimiter.
	--input-lines-terminated-by	Specifies the line delimiter.
Hive parameters	--hive-delims-replacement	Replaces characters such as <code>\r</code> and <code>\n</code> in data with user-defined characters.
	--hive-drop-import-delims	Removes characters such as <code>\r</code> and <code>\n</code> when data is imported to Hive.
	--map-column-hive	Specifies the data type of fields during the generation of a Hive table.
	--hive-partition-key	Creates a partition.
	--hive-partition-value	Imports data to a specified partition of a database.
	--hive-home	Specifies the installation directory for Hive.
	--hive-import	Specifies that data is imported from a relational database to Hive.
	--hive-overwrite	Overwrites existing Hive data.
	--create-hive-table	Creates a Hive table. The default value is <b>false</b> . A destination table will be created if it does not exist.
	--hive-table	Specifies a Hive table to which data is to be imported.
	--table	Specifies the relational database table.
	--columns	Specifies the fields of a relational data table to be imported.
--query	Specifies the query statement for importing the query result.	
HCatalog parameters	--hcatalog-database	Specifies a Hive database and imports data to it using HCatalog.
	--hcatalog-table	Specifies a Hive table and imports data to it using HCatalog.
Others	-m or --num-mappers	Specifies the number of map tasks used by a Sqoop job.

Category	Parameter	Description
	--split-by	Specifies the column based on which Sqoop splits work units. This parameter is used together with <b>-m</b> .
	--target-dir	Specifies the temporary directory of HDFS.
	--null-string string	Specifies the string to be written for a null value for string columns.
	--null-non-string	Specifies the string to be written for a null value for non-string columns.
	--check-column	Specifies the column for determining incremental data import.
	--incremental 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	Specifies the last value of the check column from the previous import.
	--input-null-string	Specifies the string to be interpreted as <b>NULL</b> for string columns.
	--input-null-non-string	Specifies the string to be interpreted as null for non-string columns. If this parameter is not specified, <b>NULL</b> will be used.

## 27.4 Common Issues About Sqoop

### 27.4.1 What Should I Do If Class QueryProvider Is Unavailable?

#### Question

What should I do if the QueryProvider class is unavailable?

```
2021-04-06 15:57:10,756 INFO manager.SqlManager: Using default fetchSize of 1000
2021-04-06 15:57:10,756 INFO tool.CodeGenTool: Beginning code generation
Apr 06, 2021 3:57:10 PM java.util.logging.LogManager$RootLogger log
SEVERE: Error loading factory org.apache.calcite.jdbc.CalciteJdbc41Factory
java.lang.NoClassDefFoundError: org/apache/calcite/linq4j/QueryProvider
 at java.lang.ClassLoader.defineClass1(Native Method)
 at java.lang.ClassLoader.defineClass(ClassLoader.java:757)
 at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:142)
 at java.net.URLClassLoader.defineClass(URLClassLoader.java:468)
 at java.net.URLClassLoader.access$100(URLClassLoader.java:74)
 at java.net.URLClassLoader$1.run(URLClassLoader.java:369)
 at java.net.URLClassLoader$1.run(URLClassLoader.java:363)
 at java.security.AccessController.doPrivileged(Native Method)
 at java.net.URLClassLoader.findClass(URLClassLoader.java:362)
 at java.lang.ClassLoader.loadClass(ClassLoader.java:419)
 at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:352)
 at java.lang.ClassLoader.loadClass(ClassLoader.java:352)
 at java.lang.ClassLoader.defineClass1(Native Method)
 at java.lang.ClassLoader.defineClass(ClassLoader.java:757)
 at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:142)
 at java.net.URLClassLoader.defineClass(URLClassLoader.java:468)
 at java.net.URLClassLoader.access$100(URLClassLoader.java:74)
 at java.net.URLClassLoader$1.run(URLClassLoader.java:369)
 at java.net.URLClassLoader$1.run(URLClassLoader.java:363)
 at java.security.AccessController.doPrivileged(Native Method)
 at java.net.URLClassLoader.findClass(URLClassLoader.java:362)
 at java.lang.ClassLoader.loadClass(ClassLoader.java:419)
 at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:352)
 at java.lang.ClassLoader.loadClass(ClassLoader.java:352)
 at java.lang.ClassLoader.defineClass1(Native Method)
 at java.lang.ClassLoader.defineClass(ClassLoader.java:757)
```

### Answer

Search for the MRS client directory and save the following JAR packages to the **lib** directory of Sqoop.

```
-rwxr-xr-x. 1 omm wheel 4813045 Apr 6 15:56 calcite-core-1.19.0.jar
-rwxr-xr-x. 1 omm wheel 459944 Apr 6 16:01 calcite-linq4j-1.19.0.jar
```

## 27.4.2 How Do I Do If PostgreSQL or GaussDB Fails to Connect?

### Question

How do I do if PostgreSQL or GaussDB fails to connect?

```
at org.apache.sqoop.Sqoop.runTool(Sqoop.java:243)
at org.apache.sqoop.Sqoop.main(Sqoop.java:222)
2021-09-06 09:31:27,638 ERROR org.apache.sqoop.Sqoop: Got exception running Sqoop: java.lang.RuntimeException: org.postgresql.util.PSQLException: The authentication type 12 is not supported. Check that you have configured the pg_hba.conf file to include the client's IP address or subnet, and that it is using an authentication scheme supported by the driver.
java.lang.RuntimeException: org.postgresql.util.PSQLException: The authentication type 12 is not supported. Check that you have configured the pg_hba.conf file to include the client's IP address or subnet, and that it is using an authentication scheme supported by the driver.
 at org.apache.sqoop.manager.catalogquerymanager.ListTables(CatalogQueryManager.java:118)
 at org.apache.sqoop.tool.ListTablesTool.run(ListTablesTool.java:49)
 at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:76)
 at org.apache.sqoop.Sqoop.runSqoop(Sqoop.java:183)
 at org.apache.sqoop.Sqoop.runTool(Sqoop.java:244)
 at org.apache.sqoop.Sqoop.runTool(Sqoop.java:243)
 at org.apache.sqoop.Sqoop.main(Sqoop.java:222)
Caused by: org.postgresql.util.PSQLException: The authentication type 12 is not supported. Check that you have configured the pg_hba.conf file to include the client's IP address or subnet, and that it is using an authentication scheme supported by the driver.
 at org.postgresql.core.v3.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:504)
 at org.postgresql.core.v3.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:173)
 at org.postgresql.core.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:64)
 at org.postgresql.jdbc2.AbstractJdbc2Connection.<init>(AbstractJdbc2Connection.java:136)
 at org.postgresql.jdbc3.AbstractJdbc3Connection.<init>(AbstractJdbc3Connection.java:29)
 at org.postgresql.jdbc3g.AbstractJdbc3gConnection.<init>(AbstractJdbc3gConnection.java:21)
 at org.postgresql.jdbc4.AbstractJdbc4Connection.<init>(AbstractJdbc4Connection.java:31)
 at org.postgresql.jdbc4.jdbcConnection.<init>(AbstractJdbc4Connection.java:24)
 at org.postgresql.Driver.makeConnection(Driver.java:397)
 at org.postgresql.Driver.connect(Driver.java:267)
 at java.sql.DriverManager.getConnection(DriverManager.java:664)
 at java.sql.DriverManager.getConnection(DriverManager.java:247)
 at org.apache.sqoop.manager.SqlManager.makeConnection(SqlManager.java:984)
 at org.apache.sqoop.manager.GenericJdbcManager.getConnection(GenericJdbcManager.java:59)
 at org.apache.sqoop.manager.catalogquerymanager.ListTables(CatalogQueryManager.java:102)
 ... 3 more
[omh@node-master1PM1:110]#
```

### Answer

- Scenario 1: (**import** scenarios) Run the **sqoop import** command to extract the open source PostgreSQL to MRS HDFS or Hive.

- Symptom:
  - The **sqoop** command can be executed to query PostgreSQL tables, but an error is reported when the **sqoop import** command is executed.
    - The authentication type 12 is not supported. Check that you have configured the `pg_hba.conf` file to include the client's IP address or subnet, and that it
    - The authentication type 5 is not supported. Check that you have configured the `pg_hba.conf` file to include the client's IP address or subnet, and that it
- Root cause:
  - If the authentication type is 5, the root cause is as follows: When the **sqoop import** command is executed, a MapReduce job is started. The PostgreSQL driver package **gsjdbc4-\*.jar** exists in the MRS Hadoop installation directory `/opt/Bigdata/FusionInsight_HD_*/1*_DataNode/install/hadoop/share/hadoop/common/lib`, which is incompatible with the open source PostgreSQL service. As a result, an error is reported.
  - If the authentication type is 12, the root cause is as follows: The **pg\_hba.conf** file of the database is incorrectly configured.
- Solution:
  - If the authentication type is 5, the solution is as follows: Move the driver package **gsjdbc4-\*.jar** to the **tmp** directory on each MRS core node.
 

```
mv /opt/Bigdata/FusionInsight_HD_*/1*_DataNode/install/hadoop/share/hadoop/common/lib/gsjdbc4-*.jar /tmp
```
  - If the authentication type is 12, the solution is as follows: Modify the **pg\_hba.conf** file of the database by changing the value of **ADDRESS** to the IP address of the node where Sqoop resides.

```
TYPE DATABASE USER ADDRESS METHOD
"local" is for Unix domain socket connections only
local all all trust
IPv4 local connections:
host all all 127.0.0.1/32 trust
host all all 0.0.0.0/0 md5
IPv6 local connections:
host all all ::1/128 trust
#host all all 0.0.0.0/0 password
Allow replication connections from localhost, by a user with the
replication privilege.
local replication postgres trust
host replication postgres 127.0.0.1/32 trust
host replication postgres ::1/128 trust
```

- Scenario 2: (**export** scenarios) Run the **sqoop export** command to extract the open source PostgreSQL to MRS HDFS or Hive.

- Symptom:  
The **sqoop** command can be executed to query PostgreSQL tables, but the error message "The authentication type 5 is not supported." is displayed when the **sqoop export** command is executed. Check that you have configured the `pg_hba.conf` file to include the client's IP address or subnet, and that it
- Root cause:  
When the **sqoop export** command is executed, a MapReduce job is started. The PostgreSQL driver package **gsjdbc4-\*.jar** exists in the MRS Hadoop installation directory `/opt/Bigdata/FusionInsight_HD_*/1_*_DataNode/install/hadoop/share/hadoop/common/lib`, which is incompatible with the open-source PostgreSQL service. As a result, an error is reported.
- Solution:
  1. Move the driver package **gsjdbc4-\*.jar** to the `tmp` directory on each MRS core node.  

```
mv /opt/Bigdata/FusionInsight_HD_*/1_*_DataNode/install/hadoop/share/hadoop/common/lib/gsjdbc4-*.jar /tmp
```
  2. Delete `/opt/Bigdata/client/Hive/Beeline/lib/gsjdbc4-*.jars`.

### 27.4.3 What Should I Do If Data Failed to Be Synchronized to a Hive Table on the OBS Using hive-table?

#### Question

What should I do if data failed to be synchronized to a Hive table on the OBS using hive-table?

```
2021-09-03 16:28:11,611 ERROR tools.DistCp: XAttrs not supported on at least one file system:
org.apache.hadoop.tools.CopyListing$XAttrsNotSupportedException: XAttrs not supported for file system:
obs://fdd-fs
 at org.apache.hadoop.tools.util.DistCpUtils.checkFileSystemXAttrSupport(DistCpUtils.java:555)
 at org.apache.hadoop.tools.DistCp.configureOutputFormat(DistCp.java:341)
 at org.apache.hadoop.tools.DistCp.createJob(DistCp.java:308)
 at org.apache.hadoop.tools.DistCp.createAndSubmitJob(DistCp.java:218)
 at org.apache.hadoop.tools.DistCp.execute(DistCp.java:197)
 at org.apache.hadoop.tools.DistCp.run(DistCp.java:155)
```

#### Answer

Change `-hive-table` to `-hcatalog-table`.

### 27.4.4 What Should I Do If Data Failed to Be Synchronized to an ORC or Parquet Table Using hive-table?

#### Question

What should I do if data failed to be synchronized to the ORC or parquet table using hive-table and error message that contains the **kite-sdk** package name is displayed?

## Answer

Change `-hive-table` to `-hcatalog-table`.

## 27.4.5 What Should I Do If Data Failed to Be Synchronized Using `hive-table`?

### Question

What should I do if data failed to be synchronized using `hive-table`?

```
at org.apache.hadoop.hive.ql.metadata.Hive.registerAllFunctionsOnce(Native.java:200) [Native-Exec:0.1.0-UMW*E1*312LWU1*0RAFDJWU1]
.. 41 more
14:41:42.891 [E1438c-07bb-43fd-91d9-910a348f6e91 main] ERROR org.apache.hadoop.hive.metastore.ObjectStore - Version information not found in metastore. The process will exit.
14:41:42.892 [E1438c-07bb-43fd-91d9-910a348f6e91 main] ERROR org.apache.hadoop.hive.metastore.RetryingHMSHandler - ExitSecurityException
at org.apache.sqoop.util.SubprocessSecurityManager.checkExit(SubprocessSecurityManager.java:83)
at java.lang.Runtime.exit(Runtime.java:107)
at java.lang.System.exit(System.java:973)
at org.apache.hadoop.hive.metastore.ObjectStore.checkSchema(ObjectStore.java:9655)
at org.apache.hadoop.hive.metastore.ObjectStore.verifySchema(ObjectStore.java:9631)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:498)
at org.apache.hadoop.hive.metastore.RawStoreProxy.invoke(RawStoreProxy.java:97)
at com.sun.proxy.$Proxy37.verifySchema(Unknown Source)
at org.apache.hadoop.hive.metastore.HiveMetaStore$HMSHandler.getMSForConf(HiveMetaStore.java:903)
at org.apache.hadoop.hive.metastore.HiveMetaStore$HMSHandler.getMS(HiveMetaStore.java:895)
at org.apache.hadoop.hive.metastore.HiveMetaStore$HMSHandler.createDefaultDB(HiveMetaStore.java:978)
at org.apache.hadoop.hive.metastore.HiveMetaStore$HMSHandler.init(HiveMetaStore.java:565)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:498)
at org.apache.hadoop.hive.metastore.RetryingHMSHandler.invokeInternal(RetryingHMSHandler.java:148)
at org.apache.hadoop.hive.metastore.RetryingHMSHandler.invoke(RetryingHMSHandler.java:109)
at org.apache.hadoop.hive.metastore.RetryingHMSHandler.<init>(RetryingHMSHandler.java:81)
at org.apache.hadoop.hive.metastore.RetryingHMSHandler.getProxy(RetryingHMSHandler.java:94)
at org.apache.hadoop.hive.metastore.HiveMetaStore.newRetryingHMSHandler(HiveMetaStore.java:9683)
at org.apache.hadoop.hive.metastore.HiveMetaStoreClient.<init>(HiveMetaStoreClient.java:185)
at org.apache.hadoop.hive.ql.metadata.SessionHiveMetaStoreClient.<init>(SessionHiveMetaStoreClient.java:96)
at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method)
at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:62)
at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:45)
at java.lang.reflect.Constructor.newInstance(Constructor.java:423)
at org.apache.hadoop.hive.metastore.util.JavaUtils.newInstance(JavaUtils.java:84)
at org.apache.hadoop.hive.metastore.RetryingMetaStoreClient.<init>(RetryingMetaStoreClient.java:97)
.....
```

## Answer

Add the following content to the `hive-site.xml` file.

```
<property>
<name>hive.metastore.schema.verification</name>
<value>false</value>
</property>
```

## 27.4.6 What Should I Do If Data Failed to Be Synchronized to a Hive Parquet Table Using HCatalog?

### Question

When the partition fields in a Hive parquet table are not of the string type, data in the table can be synchronized only using HCatalog. What should I do if the following error message is displayed during data synchronization?

```

2021-09-28 12:12:17.623 INFO common.HCatUtil: mapreduce.lib.hcatoutput.hive.conf is set. Applying configuration differences.
2021-09-28 12:12:17.629 INFO common.HiveClientCache: Initializing cache: eviction-timeout=120 initial-capacity=50 maximum-capacity=50
2021-09-28 12:12:17.648 INFO metastore.HiveMetaStoreClient: Trying to connect to metastore with URI thrift://node-master4yp0w.a0dbfe45-7b6c-4386-83
68f775c0dd.com:9083
2021-09-28 12:12:17.649 INFO metastore.HiveMetaStoreClient: Opened a connection to metastore, current connections: 2
2021-09-28 12:12:17.651 INFO metastore.HiveMetaStoreClient: Connected to metastore.
2021-09-28 12:12:17.651 INFO metastore.RetryingMetaStoreClient: RetryingMetaStoreClient proxy=class org.apache.hive.hcatalog.common.HiveClientCache
eableHiveMetaStoreClient ugi=poseidon (auth:SIMPLE) retries=1 delay=1 lifetime=0
2021-09-28 12:12:17.875 WARN conf.HiveConf: HiveConf of name hive.http.filter.initializers does not exist
2021-09-28 12:12:17.876 WARN conf.HiveConf: HiveConf of name hive.server2.authentication.ldap.url.port does not exist
2021-09-28 12:12:17.877 INFO conf.HiveConf: current conf hive.parquet.time.zone.isLocal=true
2021-09-28 12:12:18.056 INFO hcat.SqoopHCatUtilities: Setting hCatInputFormat filter to days='20210928'
2021-09-28 12:12:18.072 WARN conf.HiveConf: HiveConf of name hive.http.filter.initializers does not exist
2021-09-28 12:12:18.072 WARN conf.HiveConf: HiveConf of name hive.server2.authentication.ldap.url.port does not exist
2021-09-28 12:12:18.073 INFO conf.HiveConf: current conf hive.parquet.time.zone.isLocal=true
2021-09-28 12:12:18.073 INFO common.HCatUtil: mapreduce.lib.hcatoutput.hive.conf is set. Applying configuration differences.
2021-09-28 12:12:18.108 ERROR tool.ImportTool: Import failed: java.io.IOException: MetaException(message:Filtering is supported only on partition h
f type string)
 at org.apache.hive.hcatalog.mapreduce.HCatInputFormat.setFilter(HCatInputFormat.java:120)
 at org.apache.sqoop.mapreduce.hcat.SqoopHCatUtilities.configureHCat(SqoopHCatUtilities.java:391)
 at org.apache.sqoop.mapreduce.hcat.SqoopHCatUtilities.configureImportOutputFormat(SqoopHCatUtilities.java:850)
 at org.apache.sqoop.mapreduce.ImportJobBase.configureOutputFormat(ImportJobBase.java:102)
 at org.apache.sqoop.mapreduce.ImportJobBase.runImport(ImportJobBase.java:263)
 at org.apache.sqoop.manager.SqlManager.importQuery(SqlManager.java:748)
 at org.apache.sqoop.tool.ImportTool.importTable(ImportTool.java:522)
 at org.apache.sqoop.tool.ImportTool.run(ImportTool.java:628)
 at org.apache.sqoop.Sqoop.run(Sqoop.java:147)
 at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:76)
 at org.apache.sqoop.Sqoop.runSqoop(Sqoop.java:103)
 at org.apache.sqoop.Sqoop.runTool(Sqoop.java:234)

```

## Answer

1. Delete the restricted code in the **SqoopHCatUtilities** class of Sqoop.
2. Change the value of the **hive.metastore.integral.jdo.pushdown** parameter in the **hive-site.xml** file on the Hive client to **true**.

## 27.4.7 What Should I Do If the Data Type of Fields timestamp and data Is Incorrect During Data Synchronization Between Hive and MySQL?

### Question

What should I do if the data type of fields timestamp and data is incorrect during data synchronization between Hive and MySQL?

```

2021-10-20 21:16:34.034 | INFO | main | current conf hive.parquet.time.zone.isLocal=true | HiveConf.java:550
2021-10-20 21:16:34.034 | INFO | Thread-19 | Auto-progress thread is finished. keepGoing=false | ProgressThread.java:159
2021-10-20 21:16:34.034 | WARN | main | Exception running child: java.lang.ClassCastException: org.apache.hadoop.hive.common.type.Timestamp cannot be cast to java.sql.Timestamp
 at org.apache.sqoop.mapreduce.hcat.SqoopHCatExportHelper.convertToSqoop(SqoopHCatExportHelper.java:203)
 at org.apache.sqoop.mapreduce.hcat.SqoopHCatExportHelper.convertToSqoopRecord(SqoopHCatExportHelper.java:139)
 at org.apache.sqoop.mapreduce.hcat.SqoopHCatExportMapper.map(SqoopHCatExportMapper.java:36)
 at org.apache.sqoop.mapreduce.hcat.SqoopHCatExportMapper.map(SqoopHCatExportMapper.java:35)
 at org.apache.hadoop.mapreduce.Mapper.run(Mapper.java:146)
 at org.apache.hadoop.mapreduce.AutoProgressMapper.run(AutoProgressMapper.java:64)
 at org.apache.hadoop.mapred.MapTask.runNewMapper(MapTask.java:799)
 at org.apache.hadoop.mapred.MapTask.run(MapTask.java:347)
 at org.apache.hadoop.mapred.YarnChild$1.run(YarnChild.java:183)
 at java.security.AccessController.doPrivileged(Native Method)
 at javax.security.auth.Subject.doAs(Subject.java:422)
 at org.apache.hadoop.mapred.security.UserGroupInformation.doAs(UserGroupInformation.java:1761)
 at org.apache.hadoop.mapred.YarnChild.main(YarnChild.java:177)
 at YarnChild.java:199

```

## Answer

- Forcibly convert the data type of the timestamp field in the Sqoop source package to be the same as that in Hive.
- Change the data type of the timestamp field in Hive to String.

# 28 Using Storm

---

## 28.1 Using Storm from Scratch

You can submit and delete Storm topologies on the MRS cluster client.

### Prerequisites

The MRS cluster client has been installed, for example, in the `/opt/hadoopclient` directory. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite.

### Procedure

**Step 1** Prepare the client based on service requirements. Log in to the node where the client is installed.

**Step 2** Run the following command to switch to the client directory, for example, `/opt/hadoopclient`:

```
cd /opt/hadoopclient
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** For clusters with Kerberos authentication enabled, run the following command to authenticate the user. For clusters with Kerberos authentication disabled, skip this step.

```
kinit Storm user
```

**Step 5** Run the following command to submit the Storm topology:

```
storm jar Path of the topology package Class name of the topology Main method
Topology name
```

If the following information is displayed, the topology is submitted successfully.

```
Finished submitting topology: topo1
```



**Step 6** Run the following command to query Storm topologies. For clusters with Kerberos authentication enabled, only users in the **stormadmin** or **storm** group can query all topologies.

```
storm list
```

**Step 7** Run the following command to delete a Storm topology.

```
storm kill Topology name
```

```
----End
```

## 28.2 Using the Storm Client

### Scenario

This section describes how to use the Storm client in an O&M scenario or service scenario.

### Prerequisites

- You have installed the client. For example, the installation directory is **/opt/hadoopclient**.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users have downloaded the keytab file. A human-machine user must change the password upon the first login. (Not involved in normal mode)

### Procedure

**Step 1** Prepare the client based on service requirements. Log in to the node where the client is installed.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/hadoopclient
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** If multiple Storm instances are installed, run the following command to load the environment variables of a specific instance when running the Storm command to submit the topology. Otherwise, skip this step. The following command uses the instance Storm-2 as an example.

```
source Storm-2/component_env
```

**Step 5** Run the following command to perform user authentication (skip this step in normal mode):

```
kinit Component service user
```

**Step 6** Run the following command to perform operations on the client:

For example, run the following command:

- **cql**
- **storm**

 **NOTE**

A Storm client cannot be connected to secure and non-secure ZooKeepers at the same time.

----End

## 28.3 Submitting Storm Topologies on the Client

### Scenario

You can submit Storm topologies on the cluster client to continuously process stream data. For clusters with Kerberos authentication enabled, users who submit topologies must be members of the **stormadmin** or **storm** group.

### Prerequisites

The client has been updated.

### Procedure

- Step 1** Prepare the client based on service requirements. Log in to the node where the client is installed.
- Step 2** Run the following command to set the permissions on the topology JAR file:  
For example, run the following command to change the permissions on **/opt/storm/topology.jar**:  
**chmod 600 /opt/storm/topology.jar**
- 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** If multiple Storm instances are installed, run the following command to load the environment variables of a specific instance when running the Storm command to submit the topology. Otherwise, skip this step. The following command uses the instance Storm-2 as an example.  
**source Storm-2/component\_env**
- Step 6** For clusters with Kerberos authentication enabled, run the following command to authenticate the user. For clusters with Kerberos authentication disabled, skip this step.  
**kinit Storm user**
- Step 7** For versions earlier than MRS 3.x, run the following command to submit the Storm topology:

**storm jar** *Path of the topology package Class name of the topology Main method Topology name*

If the following information is displayed, the topology is submitted successfully.

```
Finished submitting topology: topo1
```

 **NOTE**

- To support sampling messages, add the **topology.debug** and **topology.eventlogger.executors** parameters.
- Data processing methods vary with topologies. The topology in the example generates characters randomly and separates character strings. To query the processing status, enable the sampling function and perform operations according to [Querying Storm Topology Logs](#).

**Step 8** Run the following command to submit a topology task for MRS 3.x or later:

**storm jar** *topology-jar-path class input parameter list*

- *topology-jar-path* indicates the path of the JAR file of the topology.
- *class* indicates the class name of the main method used by the topology.
- *Input parameter list* includes input parameters of the main method used by the topology.

If the following information is displayed, the topology is submitted successfully:

```
Finished submitting topology: topology1
```

 **NOTE**

- The login authentication user must correspond to the loaded environment variable (**component\_env**). Otherwise, an error occurs when you run the **storm** command to submit the topology task.
- After the client environment variable is loaded and the corresponding user login succeeds, the user can run the Storm command on any Storm client to submit the topology task. After the command is executed, the successfully submitted topology is still in the Storm cluster of the user.
- If the cluster domain name is changed, you need to reset the domain name before submitting the topology. Run the cql statement.

**Step 9** Run the following command to query Storm topologies. For clusters with Kerberos authentication enabled, only users in the **stormadmin** or **storm** group can query all topologies.

**storm list**

----End

## 28.4 Accessing the Storm Web UI

### Scenario

The Storm web UI provides a graphical interface for using Storm.

The following information can be queried on the Storm web UI:

- Storm cluster summary

- Nimbus summary
- Topology summary
- Supervisor summary
- Nimbus configurations

## Prerequisites

- The password of user **admin** has been obtained. The password of user **admin** is specified by you during the cluster creation.
- If a user other than **admin** is used to access the Storm web UI, the user must be added to the **storm** or **stormadmin** user group.

## Procedure

**Step 1** Access the component management page.

- For versions earlier than MRS 3.x, click the cluster name to go to the cluster details page and choose **Components**.

### NOTE

- If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)
- For MRS 3.x or later, log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Choose **Cluster** > *Name of the desired cluster* > **Services**.

**Step 2** Log in to the Storm WebUI.

- For versions earlier than MRS 3.x: Choose **Storm**. On the **Storm Summary** area, click any UI link on the right side of **Storm Web UI** to open the Storm web UI.

### NOTE

- When accessing the Storm web UI for the first time, you must add the address to the trusted site list.
- For MRS 3.x or later, choose **Storm** > **Overview**. In the **Basic Information** area, click any UI link on the right side of **Storm Web UI** to open the Storm web UI.

----End

## Related Tasks

- Click a topology name to view details, status, Spouts information, Bolts information, and configuration information of the topology.
- In the **Topology actions** area, click **Activate**, **Deactivate**, **Rebalance**, **Kill**, **Debug**, **Stop Debug**, and **Change Log Level** to activate, deactivate, redeploy, delete, debug, and stop debugging the topology, and modify the log levels, respectively. You need to set the waiting time for the redeployment and deletion operations. The unit is second.
- In the **Topology Visualization** area, click **Show Visualization** to visualize a topology. After the topology is visualized, the WebUI displays the topology structure.

## 28.5 Managing Storm Topologies

### Scenario

You can manage Storm topologies on the Storm web UI. Users in the **storm** group can manage only the topology tasks submitted by themselves, while users in the **stormadmin** group can manage all topology tasks.

### Procedure

**Step 1** For details about how to access the Storm WebUI, see [Accessing the Storm Web UI](#).

**Step 2** In the **Topology summary** area, click the desired topology.

**Step 3** Use options in **Topology actions** to manage the Storm topology.

- Activating a topology  
Click **Activate** to activate the topology.
- Deactivating a topology  
Click **Deactivate** to deactivate the topology.
- Re-deploying a topology  
Click **Rebalance** and specify the wait time (in seconds) of re-deployment. Generally, if the number of nodes in a cluster changes, the topology can be re-deployed to maximize resource usage.
- Deleting a topology  
Click **Kill** and specify the wait time (in seconds) of the deletion.
- Starting or stopping sampling messages  
Click **Debug**. In the dialog box displayed, specify the percentage of the sampled data volume. For example, if the value is set to **10**, 10% of data is sampled.  
To stop sampling, click **Stop Debug**.

#### NOTE

This function is available only if the sampling function is enabled when the topology is submitted. For details about querying data processing information, see [Querying Storm Topology Logs](#).

- Modifying the topology log level  
Click **Change Log Level** to specify a new log level.

**Step 4** Displaying a topology

In the **Topology Visualization** area, click **Show Visualization** to visualize the topology.

----End

## 28.6 Querying Storm Topology Logs

### Scenario

You can query topology logs to check the execution of a Storm topology in a worker process. To query the data processing logs of a topology, enable the **Debug** function when submitting the topology. Only streaming clusters with Kerberos authentication enabled support this function. In addition, the user who queries topology logs must be the one who submits the topology or a member of the **stormadmin** group.

### Prerequisites

- The network of the working environment has been configured.
- The sampling function has been enabled for the topology.

### Querying Worker Process Logs

**Step 1** For details about how to access the Storm WebUI, see [Accessing the Storm Web UI](#).

**Step 2** In the **Topology Summary** area, click the desired topology to view details.

**Step 3** Click the desired **Spouts** or **Bolts** task. In the **Executors (All time)** area, click a port in **Port** to view detailed logs.

----End

### Querying Data Processing Logs of a Topology

**Step 1** For details about how to access the Storm WebUI, see [Accessing the Storm Web UI](#).

**Step 2** In the **Topology Summary** area, click the desired topology to view details.

**Step 3** Click **Debug**, specify the data sampling ratio, and click **OK**.

**Step 4** Click the **Spouts** or **Bolts** task. In **Component summary**, click **events** to view data processing logs.

----End

## 28.7 Storm Common Parameters

This section applies to MRS 3.x or later.

### Navigation Path

For details about how to set parameters, see [Modifying Cluster Service Configuration Parameters](#).

## Parameter Description

**Table 28-1** Parameter description

Parameter	Description	Default Value
supervisor.slots.ports	Specifies the list of ports that can run workers on the supervisor. Each worker occupies a port, and each port runs only one worker. This parameter is used to set the number of workers that can run on each server. Ports range from 1024 to 65535, and ports are separated by commas (,).	6700,6701,6702,6703
WORKER_GC_OPTS	Specifies the JVM option used for supervisor to start worker. It is recommended that you set this parameter based on memory usage of a service. For simple service processing, the recommended value is <b>-Xmx1G</b> . If window cache is used, the value of this parameter is calculated based on the following formula: Size of each record x Period x 2	-Xms1G -Xmx1G -XX:+UseG1GC -XX:+PrintGCDetails -Xloggc:artifacts/gc.log -XX:+PrintGCDateStamps -XX:+PrintGCTimeStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=1M -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=artifacts/heapdump
default.scheduler.mode	Specifies the default scheduling mode of the scheduler. Options are as follows: <ul style="list-style-type: none"> <li>● <b>AVERAGE</b>: indicates that the scheduling mechanism that uses the number of idle slots as the priority is used.</li> <li>● <b>RATE</b>: indicates that the scheduling mechanism that uses the rate of idle slots as the priority is used.</li> </ul>	AVERAGE
nimbus.thrift.threads	Set the maximum number of connection threads when the active Nimbus externally provides services. If the Storm cluster is large and the number of Supervisor instances is large, increase connection threads.	512

## 28.8 Configuring a Storm Service User Password Policy

### Scenario

This section applies to MRS 3.x or later.

After submitting a topology task, a Storm service user must ensure that the task continuously runs. During topology running, the worker process may need to restart to ensure continuous topology work. If the password of a service user is changed or the number of days that a password is used exceeds the maximum number specified in a password policy, topology running may be affected. An MRS cluster administrator must configure a separate password policy for Storm service users based on enterprise security requirements.

#### NOTE

If a separate password policy is not configured for Storm service users, an old topology can be deleted and then submitted again after a service user password is changed so that the topology can continuous run.

### Impact on the System

- After a separate password policy is configured for a Storm service user, the user is not affected by **Password Policy** on the Manager page.
- If a separate password policy is configured for a Storm service user and cross-cluster entrusted relationships are configured, a password must be reset for the Storm service user on Manager based on the password policy.

### Prerequisites

An MRS cluster administrator has understood service requirements and created a Human-Machine user, for example, **testpol**.

### Procedure

**Step 1** Log in to any node in the cluster 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 commands to export the environment variables:

```
EXECUTABLE_HOME="${CONTROLLER_HOME}/kerberos_user_specific_binay/
kerberos"
```

```
LD_LIBRARY_PATH=${EXECUTABLE_HOME}/lib:$LD_LIBRARY_PATH
```



```
PATH=${EXECUTABLE_HOME}/bin:$PATH
```

- Step 4** Run the following command and enter the Kerberos administrator password to log in to the Kerberos console:

```
kadmin -p kadmin/admin
```

 **NOTE**

For initial use, the **kadmin/admin** password must be changed for the **kadmin/admin** user.

If the following information is displayed, you have successfully logged in to the Kerberos console.

```
kadmin:
```

- Step 5** Run the following command to check details about the created **Human-Machine** user:

```
getprinc Username
```

Sample command for viewing details about the **testpol** user:

```
getprinc testpol
```

If the following information is displayed, the specified user has used the default password policy:

```
Principal: testpol@<System domain name>
```

```
.....
```

```
Policy: default
```

- Step 6** Run the following command to create a separate password policy, such as **streampol**, for the Storm service user:

```
addpol -maxlife 0day -minlife 0sec -history 1 -maxfailure 5 -
failurecountinterval 5min -lockoutduration 5min -minlength 8 -minclasses 4
streampol
```

In the command, **-maxlife** indicates the maximum validity period of a password, and **0day** indicates that a password will never expire.

- Step 7** Run the following command to view the newly created policy **streampol**:

```
getpol streampol
```

If the following information is displayed, the new policy specifies that the password will never expire:

```
Policy: streampol
```

```
Maximum password life: 0 days 00:00:00
```

```
.....
```

- Step 8** Run the following command to apply the new policy **streampol** to the **testpol** Storm user:

```
modprinc -policy streampol testpol
```

In the command, **streampol** indicates a policy name, and **testpol** indicates a username.

If the following information is displayed, the properties of the specified user have been modified:

```
Principal "testpol@<System domain name>" modified.
```

**Step 9** Run the following command to view current information about the **testpol** Storm user:

**getprinc testpol**

If the following information is displayed, the specified user has used the new password policy:

```
Principal: testpol@<System domain name>
```

```
.....
Policy: streampol
```

```
----End
```

## 28.9 Migrating Storm Services to Flink

### 28.9.1 Overview

This section applies to MRS 3.x or later.

From 0.10.0, Flink provides a set of APIs to smoothly migrate services compiled using Storm APIs to the Flink platform. This can be used in most of the service scenarios.

Flink supports the following service migration modes:

1. Complete migration of Storm services: Convert and run a complete Storm topology developed by Storm APIs.
2. Embedded migration of Storm services: Storm code is embedded in DataStream of Flink, for example, Spout/Bolt compiled using Storm APIs.

Flink provides the flink-storm package for the preceding service migration.

### 28.9.2 Completely Migrating Storm Services

#### Scenarios

This section describes how to convert and run a complete Storm topology developed using Storm API.

#### Procedure

**Step 1** Open the Storm service project, modify the POM file of the project, and add the reference of **flink-storm\_2.11**, **flink-core**, and **flink-streaming-java\_2.11**. The following figure shows an example.

```
<dependency>
 <groupId>org.apache.flink</groupId>
 <artifactId>flink-storm_2.11</artifactId>
 <version>1.4.0</version>
 <exclusions>
 <exclusion>
 <groupId>*</groupId>
 <artifactId>*</artifactId>
 </exclusion>
 </exclusions>
```

```

 </exclusions>
 </dependency>
 <dependency>
 <groupId>org.apache.flink</groupId>
 <artifactId>flink-core</artifactId>
 <version>1.4.0</version>
 <exclusions>
 <exclusion>
 <groupId>*</groupId>
 <artifactId>*</artifactId>
 </exclusion>
 </exclusions>
 </dependency>

 <dependency>
 <groupId>org.apache.flink</groupId>
 <artifactId>flink-streaming-java_2.11</artifactId>
 <version>1.4.0</version>
 <exclusions>
 <exclusion>
 <groupId>*</groupId>
 <artifactId>*</artifactId>
 </exclusion>
 </exclusions>
 </dependency>

```

 **NOTE**

If the project is not a non-Maven project, manually collect the preceding JAR packages and add them to the `classpath` environment variable of the project.

**Step 2** Modify the code for submission of the topology. The following uses WordCount as an example:

1. Keep the structure of the Storm topology unchanged, including the Spout and Bolt developed using Storm API.

```

TopologyBuilder builder = new TopologyBuilder();
builder.setSpout("spout", new RandomSentenceSpout(), 5);
builder.setBolt("split", new SplitSentenceBolt(), 8).shuffleGrouping("spout");
builder.setBolt("count", new WordCountBolt(), 12).fieldsGrouping("split", new Fields("word"));

```

2. Modify the code for submission of the topology. An example is described as follows:

```

Config conf = new Config();
conf.setNumWorkers(3);
StormSubmitter.submitTopology("word-count", conf, builder.createTopology());

```

Perform the following operations:

```

Config conf = new Config();
conf.setNumWorkers(3);
//converts Storm Config to StormConfig of Flink.
StormConfig stormConfig = new StormConfig(conf);
//Construct FlinkTopology using TopologBuilder of Storm.
FlinkTopology topology = FlinkTopology.createTopology(builder);
//Obtain the Stream execution environment.
StreamExecutionEnvironment env = topology.getExecutionEnvironment();
//Set StormConfig to the environment variable of Job to construct Bolt and Spout.
//If StormConfig is not required during the initialization of Bolt and Spout, you do not need to set this parameter.
env.getConfig().setGlobalJobParameters(stormConfig);
//Submit the topology.
topology.execute();

```

3. After the package is repacked, run the following command to submit the package:

```
flink run -class {MainClass} WordCount.jar
```

```
----End
```

## 28.9.3 Performing Embedded Service Migration

### Scenarios

This section describes how to embed Storm code in DataStream of Flink in embedded migration mode. For example, the code of Spout or Bolt compiled using Storm API is embedded.

### Procedure

- Step 1** In Flink, perform embedded conversion to Spout and Bolt in the Storm topology to convert them to Flink operators. The following is an example of the code:

```
//set up the execution environment
final StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
//get input data
final DataStream<String> text = getTextDataStream(env);
final DataStream<Tuple2<String, Integer>> counts = text
 //split up the lines in pairs (2-tuples) containing: (word,1)
 //this is done by a bolt that is wrapped accordingly
 .transform("CountBolt",
 TypeExtractor.getForObject(new Tuple2<String, Integer>("", 0)),
 new BoltWrapper<String, Tuple2<String, Integer>>(new CountBolt()))
 //group by the tuple field "0" and sum up tuple field "1"
 .keyBy(0).sum(1);
// execute program
env.execute("Streaming WordCount with bolt tokenizer");
```

- Step 2** After the modification, run the following command to submit the modification:

```
flink run -class {MainClass} WordCount.jar
```

```
----End
```

## 28.9.4 Migrating Services of External Security Components Interconnected with Storm

### Migrating Services for Interconnecting Storm with HDFS and HBase

If the Storm services use the **storm-hdfs** or **storm-hbase** plug-in package for interconnection, you need to specify the following security parameters when migrating Storm services as instructed in [Completely Migrating Storm Services](#).

```
//Initialize Storm Config.
Config conf = new Config();

//Initialize the security plug-in list.
List<String> auto_tgts = new ArrayList<String>();
//Add the AutoTGT plug-in.
auto_tgts.add("org.apache.storm.security.auth.kerberos.AutoTGT");
//Add the AutoHDFS plug-in.
//If HBase is interconnected, use auto_tgts.add("org.apache.storm.hbase.security.AutoHBase") to replace the
following:
auto_tgts.add("org.apache.storm.hdfs.common.security.AutoHDFS");

//Set security parameters.
conf.put(Config.TOPOLOGY_AUTO_CREDENTIALS, auto_tgts);
```

```
//Set the number of workers.
conf.setNumWorkers(3);

//Convert Storm Config to StormConfig of Flink.
StormConfig stormConfig = new StormConfig(conf);

//Construct FlinkTopology using TopologBuilder of Storm.
FlinkTopology topology = FlinkTopology.createTopology(builder);

//Obtain the StreamExecutionEnvironment.
StreamExecutionEnvironment env = topology.getExecutionEnvironment();

//Add StormConfig to the environment variable of Job to construct Bolt and Spout.
//If Config is not required during the initialization of Bolt and Spout, do not set this parameter.
env.getConfig().setGlobalJobParameters(stormConfig);

//Submit the topology.
topology.execute();
```

After the preceding security plug-in is configured, unnecessary logins during the initialization of HDFS Bolt and HBase Bolt are avoided because the security context has been configured in Flink.

## Migrating Services of Storm Interconnected with Other Security Components

If the plug-in packages, such as **storm-kakfa-client** and **storm-solr** are used for interconnection between Storm and other components for service migration, the previously configured security plug-ins need to be deleted.

```
List<String> auto_tgts = new ArrayList<String>();
//keytab mode
auto_tgts.add("org.apache.storm.security.auth.kerberos.AutoTGTFromKeytab");

//Write the plug-in list configured on the client to the specified config parameter.
//Mandatory in security mode
//This configuration is not required in common mode, and you can comment out the following line.
conf.put(Config.TOPOLOGY_AUTO_CREDENTIALS, auto_tgts);
```

The AutoTGTFromKeytab plug-in must be deleted during service migration. Otherwise, the login will fail when Bolt or Spout is initialized.

## 28.10 Storm Log Introduction

This section applies to MRS 3.x or later.

### Log Description

Log paths: The default paths of Storm log files are **/var/log/Bigdata/storm/Role name** (run logs) and **/var/log/Bigdata/audit/storm/Role name** (audit logs).

- Nimbus: **/var/log/Bigdata/storm/nimbus** (run logs) and **/var/log/Bigdata/audit/storm/nimbus** (audit logs)
- Supervisor: **/var/log/Bigdata/storm/supervisor** (run logs) and **/var/log/Bigdata/audit/storm/supervisor** (audit logs)
- UI: **/var/log/Bigdata/storm/ui** (run logs) and **/var/log/Bigdata/audit/storm/ui** (audit logs)
- Logviewer: **/var/log/Bigdata/storm/logviewer** (run logs) and **/var/log/Bigdata/audit/storm/logviewer** (audit logs)

Log archive rule: The automatic Storm log compression function is enabled. By default, when the size of logs exceeds 10 MB, logs are automatically compressed into a log file named in the following format: *<Original log name>.log.[ID].gz*. A maximum of 20 latest compressed files are reserved by default. You can configure the number of compressed files and the compression threshold.

Names of compressed audit log files are in the format of **audit.log.[yyyy-MM-dd].[ID].zip**. These files permanently exist.

**Table 28-2** Storm log list

Log Type	Log File Name	Description
Run log	nimbus/access.log	Nimbus user access log
	nimbus/nimbus-<PID>-gc.log	GC log of the Nimbus process
	nimbus/checkavailable.log	Nimbus availability check log
	nimbus/checkService.log	Nimbus serviceability check log
	nimbus/metrics.log	Nimbus monitoring statistics log
	nimbus/nimbus.log	Run log of the Nimbus process
	nimbus/postinstall.log	Work log after Nimbus installation
	nimbus/prestart.log	Work log before Nimbus startup
	nimbus/start.log	Work log of Nimbus startup
	nimbus/stop.log	Work log of Nimbus shutdown
	supervisor/access.log	Supervisor access log
	supervisor/metrics.log	Supervisor monitoring statistics log
	supervisor/postinstall.log	Work log after supervisor installation
	supervisor/prestart.log	Work log before supervisor startup
	supervisor/start.log	Work log of supervisor startup
	supervisor/stop.log	Work log of supervisor shutdown

Log Type	Log File Name	Description
	supervisor/supervisor.log	Run log of the supervisor process
	supervisor/supervisor- <PID>-gc.log	GC log of the supervisor process
	ui/access.log	UI access log
	ui/metric.log	UI monitoring statistics log
	ui/ui-<PID>-gc.log	GC log of the UI process
	ui/postinstall.log	Work log after UI installation
	ui/prestart.log	Work log before UI startup
	ui/start.log	Work log of UI startup
	ui/stop.log	Work log of UI shutdown
	ui/ui.log	Run log of the UI process
	logviewer/access.log	Logviewer access log
	logviewer/metric.log	Logviewer monitoring statistics log
	logviewer/logviewer- <PID>-gc.log	GC log file of the logviewer process
	logviewer/logviewer.log	Run log of the logviewer process
	logviewer/postinstall.log	Work log after logviewer installation
	logviewer/prestart.log	Work log before logviewer startup
	logviewer/start.log	Work log of logviewer startup
	logviewer/stop.log	Work log of logviewer shutdown
	supervisor/[topologyId]- worker-[Port number].log	Run log of the Worker process. One port occupies one log file. By default, the system contains five ports: 29100, 29101, 29102, 29103 and 29304.

Log Type	Log File Name	Description
	supervisor/metadata/[topologyid]-worker-[Port number].yaml	Worker log metadata file, which is used by logviewer to delete logs. This file is automatically deleted by the logviewer log deletion thread based on certain conditions.
	nimbus/cleanup.log	Cleanup log of Nimbus uninstallation
	logviewer/cleanup.log	Cleanup log of logviewer uninstallation
	ui/cleanup.log	Cleanup log of UI uninstallation
	supervisor/cleanup.log	Cleanup log of supervisor uninstallation
	leader_switch.log	Run log file that records the Storm active/standby switchover
Audit log	nimbus/audit.log	Nimbus audit log
	ui/audit.log	UI audit log
	supervisor/audit.log	Supervisor audit log
	logviewer/audit	Logviewer audit log

## Log Levels

**Table 28-3** describes the log levels supported by Storm.

Levels of run logs and audit logs are ERROR, WARN, INFO, and DEBUG from the highest to the lowest priority. Run logs of equal or higher levels are recorded. The higher the specified log level, the fewer the logs recorded.

**Table 28-3** Log levels

Level	Description
ERROR	Logs of this level record error information about system running.
WARN	Logs of this level record exception information about the current event processing.



Level	Description
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of Storm by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Save the configuration. In the displayed dialog box, click **OK** to make the configurations take effect.

----End

## Log Format

The following table lists the Storm log formats:

**Table 28-4** Log Formats

Log Type	Format	Example
Run log	%d{yyyy-MM-dd HH:mm:ss,SSS}   %-5p   [%t]   %m   %logger (%F:%L) %n	2015-03-11 23:04:00,241   INFO   [RMI TCP Connection(2646)-10.0.0.2]   The baseSleepTimeMs [1000] the maxSleepTimeMs [1000] the maxRetries [1]   backtype.storm.utils.StormBoundedExponentialBackoffRetry (StormBoundedExponentialBackoffRetry.java:46)
	<yyyy-MM-dd HH:mm:ss,SSS><HostName><RoleName><logLevel><Message>	2017-03-28 02:57:52 493 10-5-146-1 storm- INFO Nimbus start normally

Log Type	Format	Example
Audit log	<i>&lt;Username&gt;&lt;User IP address&gt;&lt;Time&gt;&lt;Operation&gt;&lt;Operation object&gt;&lt;Operation result&gt;</i>	UserName=storm/hadoop, UserIP=10.10.0.2, Time=Tue Mar 10 01:15:35 CST 2015, Operation=Kill, Resource=test, Result=Success

## 28.11 Performance Tuning

### 28.11.1 Storm Performance Tuning

#### Scenario

You can modify Storm parameters to improve Storm performance in specific service scenarios.

This section applies to MRS 3.x or later.

Modify the service configuration parameters. For details, see [Modifying Cluster Service Configuration Parameters](#).

#### Topology Tuning

This task enables you to optimize topologies to improve efficiency for Storm to process data. It is recommended that topologies be optimized in scenarios with lower reliability requirements.

**Table 28-5** Tuning parameters

Parameter	Default Value	Scenario
topology.acker.executors	null	Specifies the number of acker executors. If a service application has lower reliability requirements and certain data does not need to be processed, this parameter can be set to <b>null</b> or <b>0</b> so that you can set acker off, flow control is weakened, and message delay is not calculated. This improves performance.

Parameter	Default Value	Scenario
topology.max.spout.pending	null	Specifies the number of messages cached by spout. The parameter value takes effect only when acker is not <b>0</b> or <b>null</b> . Spout adds each message sent to downstream bolt into the pending queue. The message is removed from the queue after downstream bolt processes the message and the processing is confirmed. When the pending queue is full, spout stops sending messages. Increasing the pending value improves the message throughput of spout per second but prolongs the delay.
topology.transfer.buffer.size	32	Specifies the size of the Distruptor message queue for each worker process. It is recommended that the size be between 4 to 32. Increasing the queue size improves the throughput but may prolong the delay.
RES_CPUSET_PERCENTAGE	80	Specifies the percentage of physical CPU resources used by the supervisor role instance (including startup and management worker processes) on each node. Adjust the parameter value based on service volume requirements of the node on which the supervisor exists, to optimize CPU usage.

## JVM Tuning

If an application must occupy more memory resources to process a large volume of data and the size of worker memory is greater than 2 GB, the G1 garbage collection algorithm is recommended.

**Table 28-6** Tuning parameters

Parameter	Default Value	Scenario
WORKER_GC_OPTS	-Xms1G - Xmx1G - XX:+UseG1GC - XX:+PrintGCDetails - Xloggc:artifacts/gc.log - XX:+PrintGCDateStamps - XX:+PrintGCTimeStamps - XX:+UseGCLogFileRotation - XX:NumberOfGCLogFiles=10 - XX:GCLogFileSize=1M - XX:+HeapDumpOnOutOfMemoryError - XX:HeapDumpPath=artifacts/heapdump	If an application must occupy more memory resources to process a large volume of data and the size of worker garbage memory is greater than 2 GB, the G1 garbage collection algorithm is recommended. In this case, change the parameter value to <b>-Xms2G -Xmx5G -XX:+UseG1GC</b> .

# 29 Using Tez

## 29.1 Precautions

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

## 29.2 Common Tez Parameters

### Navigation path for setting parameters:

On Manager, choose **Cluster > Service > Tez > Configuration > All Configurations**. Enter a parameter name in the search box.

### Parameter description

Table 29-1 Parameter description

Parameter	Description	Default Value
property.tez.log.dir	TezUI log directory	/var/log/Bigdata/tez/tezui
property.tez.log.level	TezUI log level	INFO

## 29.3 Accessing TezUI

Tez displays the Tez task execution process on a GUI. You can view the task execution details on the GUI.

### Prerequisite

The TimelineServer instance of the Yarn service has been installed.

## How to Use

Log in to Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). On Manager, choose **Cluster > Services > Tez**. Click the link on the right of **Tez WebUI** in the **Basic Information** area, and go to Tez web UI. You can view the details about Tez task execution.

## 29.4 Log Overview

### Log Description

**Log path:** The default save path of Tez logs is `/var/log/Bigdata/tez/role name`.

TezUI: `/var/log/Bigdata/tez/tezui` (run logs) and `/var/log/Bigdata/audit/tez/tezui` (audit logs)

**Log archive rule:** The automatic compression and archiving function of Tez is enabled. By default, when the size of a log file exceeds 20 MB (which is adjustable), the log file is automatically compressed. The naming rule of the compressed log file is as follows: `<Original log file name>-<yyyy-mm-dd_hh-mm-ss>.[/D].log.zip`. A maximum of 20 latest compressed files are retained. The number of compressed files and compression threshold can be configured.

**Table 29-2** Tez log list

Log Type	Name	Description
Run log	tezui.out	Log file that records TezUI running environment information
	tezui.log	Run log of the TezUI process
	tezui-omm-<Date>-gc.log.<No.>	GC log of the TezUI process
	prestartDetail.log	Work logs generated before the TezUI is started
	check-serviceDetail.log	Log file that records whether the TezUI service starts successfully
	postinstallDetail.log	Work logs after the TezUI is installed
	startDetail.log	Startup log of the TezUI process
	stopDetail.log	Stop log of the TezUI process
Audit log	tezui-audit.log	TezUI audit log

## Log Level

**Table 29-3** describes the log levels supported by TezUI.

Levels of run logs are ERROR, WARN, INFO, and DEBUG from the highest to the lowest priority. Run logs of equal or higher levels are recorded. The higher the specified log level, the fewer the logs recorded.

**Table 29-3** Log levels

Level	Description
ERROR	Logs of this level record error information about system running.
WARN	Exception information about the current event processing
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Log in to Manager.
- Step 2** Choose **Cluster > Service > Tez > Configuration**.
- Step 3** Select **All Configurations**.
- Step 4** In the navigation pane, choose **TezUI > Log**.
- Step 5** Select a desired log level.
- Step 6** Click **Save**. In the dialog box that is displayed, click **OK** to save the configuration.
- Step 7** Click **Instance**, select the **TezUI** role, choose **More > Restart Instance**, enter the user password, and click **OK** in the dialog box that is displayed.
- Step 8** Wait until the instance is restarted for the configuration to take effect.

----End

## Log Format

The following table lists the Tez log formats.

**Table 29-4** Log formats

Log Type	Format	Example
Run log	<yyyy-MM-dd HH:mm:ss,SSS>  <LogLevel> <Thread that generates the log>  <Message in the log>  <Location of the log event>	2020-07-31 11:44:21,378   INFO   TezUI-health-check   Start health check   com.XXX.tez.HealthCheck.run( HealthCheck.java:30)
Audit logs	<yyyy-MM-dd HH:mm:ss,SSS>  <LogLevel> <Thread that generates the log> <User Name><User IP><Time><Operation><Re source><Result><Detail > < Location of the log event >	2018-12-24 12:16:25,319   INFO   HiveServer2-Handler- Pool: Thread-185   UserName=hive UserIP=10.153.2.204 Time=2018/12/24 12:16:25 Operation=CloseSession Result=SUCCESS Detail=   org.apache.hive.service.cli.thrif t.ThriftCLIService.logAuditEven t(ThriftCLIService.java:434)

## 29.5 Common Issues

### 29.5.1 TezUI Cannot Display Tez Task Execution Details

#### Question

After a user logs in to Manager and switches to the Tez web UI, the submitted Tez tasks are not displayed.

#### Answer

The Tez task data displayed on the Tez WebUI requires the support of TimelineServer of Yarn. Ensure that TimelineServer has been enabled and is running properly before the task is submitted.

When setting the Hive execution engine to Tez, you need to set **yarn.timeline-service.enabled** to **true**. For details, see [Switching the Hive Execution Engine to Tez](#).

### 29.5.2 Error Occurs When a User Switches to the Tez Web UI

#### Question

When a user logs in to Manager and switches to the Tez web UI, error 404 or 503 is displayed.



## HTTP ERROR 404

Problem accessing /null/applicationhistory. Reason:

Not Found

Powered by Jetty:// 9.3.20.v20170531

Adapter operation failed Å» 503: Error accessing https://:20026/Yarn/TimelineServer/57/ws/v1/timeline/TEZ\_DAG\_ID

### Answer

The Tez web UI depends on the TimelineServer instance of Yarn. Therefore, TimelineServer must be installed in advance and in the **Good** state.

## 29.5.3 Yarn Logs Cannot Be Viewed on the TezUI Page

### Question

A user logs in to the Tez web UI and clicks **Logs**, but the Yarn log page fails to be displayed and data cannot be loaded.



#### This site can't be reached

10-244-224-251's server IP address could not be found.

Try running Windows Network Diagnostics.

DNS\_PROBE\_FINISHED\_NXDOMAIN

Reload

### Answer

Currently, the hostname is used for the access to the Yarn log page from the Tez web UI. Therefore, you need to configure the mapping between the hostname and IP address on the Windows host. Perform the following steps:

Modify the **C:\Windows\System32\drivers\etc\hosts** file on the Windows host and add a line indicating the mapping between the host name and IP address, for example, **10.244.224.45 10-044-224-45**. Save the modification and access the host again.

## 29.5.4 Table Data Is Empty on the TezUI HiveQueries Page

### Question

A user logs in to Manager and switches to the Tez web UI page, but no data for the submitted task is displayed on the **Hive Queries** page.

### Answer

To display task data on the **Hive Queries** page on the Tez web UI, you need to set the following parameters:

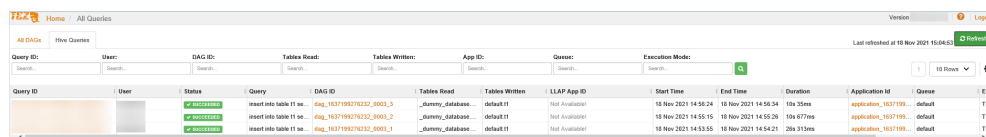
On FusionInsight Manager, choose **Cluster > Service > Hive** and click the **Configurations** tab and then **All Configurations**. In the navigation pane on the left, choose **HiveServer > Customization**. Add the following configuration to **hive-site.xml**:




Attribute	Attribute Value
hive.exec.pre.hooks	org.apache.hadoop.hive.ql.hooks.ATSHook
hive.exec.post.hooks	org.apache.hadoop.hive.ql.hooks.ATSHook
hive.exec.failure.hooks	org.apache.hadoop.hive.ql.hooks.ATSHook

### NOTE

Data display on TezUI depends on the TimelineServer instance of Yarn. If the TimelineServer instance is faulty or not started, you need to set **yarn.timeline-service.enabled** to **false** in **yarn-site.xml**. Otherwise, the Hive task fails to be executed.

After you configure the parameters and re-execute the Hive task, data can be displayed on the **Hive Queries** page. However, data of previous tasks cannot be displayed.



Query ID	User	Status	Query	DAG ID	Tables Read	Tables Written	LLAP App ID	Start Time	End Time	Duration	Application ID	Queue	Exit
			insert into table tt se...	dag_1637199792732_8003_3	_dummy_database...	default IT	Not Available	18 Nov 2021 14:56:24	18 Nov 2021 14:56:34	10s 35ms	application_1637199...	default	TEZ
			insert into table tt se...	dag_1637199792732_8003_2	_dummy_database...	default IT	Not Available	18 Nov 2021 14:55:15	18 Nov 2021 14:55:26	10s 677ms	application_1637199...	default	TEZ
			insert into table tt se...	dag_1637199792732_8003_1	_dummy_database...	default IT	Not Available	18 Nov 2021 14:53:55	18 Nov 2021 14:54:21	26s 513ms	application_1637199...	default	TEZ

# 30 Using YARN

## 30.1 Common YARN Parameters

### Allocating Queue Resources

The Yarn service provides queues for users. Users allocate system resources to each queue. After the configuration is complete, you can click **Refresh Queue** or restart the Yarn service for the configuration to take effect.

#### Navigation path for setting parameters:

For versions earlier than MRS 3.x, perform the following operations:

On the MRS console, choose **Tenants > Resource Distribution Policies**.

The following uses the **default** queue as an example. The configurations of other queues are similar. Click **Modify** to edit the parameters.

**Table 30-1** Parameter description

Parameter	Description	Default Value
Capacity	Queue resource capacity (percentage). Ensure that the capacity requirement of each queue is satisfied when the system is busy. If only a few application programs are running in a queue, the remaining resource of the queue can be shared with other queues. Note that the total capacity of all queues must be smaller than 100.	20
Maximum Capacity	Maximum queue resource usage (percentage). Due to resource sharing, the resources used by a queue may exceed its capacity. The maximum resource usage can be limited using this parameter.	100

For MRS 3.x or later, perform the following operations:

On Manager, choose **Tenant Resources > Dynamic Resource Plan > Queue Configuration**.

The following uses the **default** tenant who modifies the Superior scheduler as an example. The configurations of other queues are similar. Click **Modify** to edit the parameters.

**Table 30-2** 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	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 0 indicates that the task cannot be executed. The value ranges from -1 to 2147483647.
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. The value ranges from -1 to 2147483647.
Max Pending Apps	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. The value ranges from -1 to 2147483647.
Resource Allocation Rule	Indicates the rule for allocating resources to different tasks of a user. The rule can be FIFO or FAIR. 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.

Parameter	Description
Default Resource Label	Indicates that tasks are executed on a node with a specified resource label.
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>

## Displaying Container Logs on the Web UI

By default, the system collects container logs to HDFS. If you do not need to collect container logs to HDFS, configure the parameters in [Table 30-3](#). For details, see [Modifying Cluster Service Configuration Parameters](#).

**Table 30-3** Parameter description

Parameter	Description	Default Value
yarn.log-aggregation-enable	<p>Select whether to collect container logs to HDFS.</p> <ul style="list-style-type: none"> <li>If the parameter is set to <b>true</b>, container logs are collected to an HDFS directory. The default directory is <b>{yarn.nodemanager.remote-app-log-dir}/{user}/{thisParam}</b>. You can set the directory by setting the <b>yarn.nodemanager.remote-app-log-dir-suffix</b> parameter on the web UI.</li> <li>If this parameter is set to <b>false</b>, container logs will not be collected to HDFS.</li> </ul> <p>After changing the parameter value, restart the Yarn service for the setting to take effect.</p> <p><b>NOTE</b> The container logs that are generated before the parameter is set to <b>false</b> and the setting takes effect cannot be obtained from the web UI. You can obtain container logs from the directory specified by the <b>yarn.nodemanager.remote-app-log-dir-suffix</b> parameter before the setting takes effect.</p> <p>If you want to view the logs generated before on the web UI, you are advised to set this parameter to <b>true</b>.</p>	true

## Increasing the Number of Historical Jobs to Be Displayed on the web UI

By default, the Yarn web UI supports task list pagination. A maximum of 5,000 historical jobs can be displayed on each page, and a maximum of 10,000 historical jobs can be retained. If you need to view more jobs on the WebUI, configure parameters by referring to [Table 30-4](#). For details, see [Modifying Cluster Service Configuration Parameters](#).

**Table 30-4** Parameter description

Parameter	Description	Default Value
yarn.resourcemanager.max-completed-applications	Set the total number of historical jobs to be displayed on the web UI.	10000
yarn.resourcemanager.webapp.pagination.enable	Select whether to enable the job list background pagination function for the Yarn web UI.	true

Parameter	Description	Default Value
yarn.resourcemanager.webapp.pagination.threshold	Set the maximum number of jobs displayed on each page after the job list background pagination function of the Yarn web UI is enabled.	5000

 NOTE

- If a large number of historical jobs are displayed, the performance will be affected and the time for opening the Yarn web UI will be increased. Therefore, you are advised to enable the background pagination function and modify the **yarn.resourcemanager.max-completed-applications** parameter according to the actual hardware performance.
- After changing the parameter value, restart the Yarn service for the setting to take effect.

## 30.2 Creating Yarn Roles

### Scenario

Create and configure a YARN role. The Yarn role can be assigned with Yarn administrator permission and manage Yarn queue resources.

 NOTE

If the current component uses Ranger for permission control, you need to configure permission management policies based on Ranger. Refer to [Adding a Ranger Access Permission Policy for Yarn](#) for clusters of MRS 3.x or later.

### Prerequisites

- The MRS cluster administrator has understood service requirements.
- You have logged in to Manager.

### Procedure

For versions earlier than MRS 3.x, perform the following operations:

**Step 1** Choose **System > Manage Role > Create Role**.

**Step 2** Click **Create Role** and fill in **Role Name** and **Description**.

**Step 3** Set permissions. For details, see [Table 30-5](#).

Yarn permissions:

- **Cluster Admin Operations:** Yarn administrator permissions.
- **Scheduler Queue:** queue resources management .

**Table 30-5** Setting a role

Task	Operation
Setting the Yarn administrator permission	In the <b>Permission</b> table, click <b>Yarn</b> and select <b>Cluster Admin Operations</b> . <b>NOTE</b> The Yarn service needs to be restarted to set the Yarn administrator permission so that the saved role configuration can take effect.
Setting the permission for a user to submit tasks in a specified Yarn queue	1. In the <b>Permission</b> table, choose <b>Yarn &gt; Scheduler Queue</b> . 2. In the <b>Permission</b> column of the specified queue, select <b>Submit</b> .
Setting the permission for a user to manage tasks in a specified Yarn queue	1. In the <b>Permission</b> table, choose <b>Yarn &gt; Scheduler Queue</b> . 2. In the <b>Permission</b> column of the specified queue, select <b>Admin</b> .

If the Yarn role contains the **Submit** or **Manage** permission of a parent queue, the sub-queue inherits the permission by default, that is, the **Submit** or **Manage** permission is automatically added for the sub-queue. Permissions inherited by sub-queues will not be displayed as selected in the **Configure Resource Permission** table.

If you select only the **Submit** permission of a parent queue when setting the Yarn role, you need to manually specify the queue name when submitting tasks as a user with the permission of this role. Otherwise, when the parent queue has multiple sub-queues, the system does not automatically determine the queue to which the task is submitted and therefore submits the task to the **default** queue.

**Step 4** Click **OK**.

----End

For MRS 3.x or later, perform the following operations:

**Step 1** Choose System > Permission > Role.

**Step 2** Click **Create Role** and set a role name and enter description.

**Step 3** Refer [Table 30-6](#) to configure resource permissions for roles.

Yarn permissions:

- Cluster management: Yarn administrator permissions.
- Queue scheduling: queue resource management.



**Table 30-6** Setting a role

Task	Operation
Setting the Yarn administrator permission	In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> > <b>Yarn</b> > <b>Cluster Management</b> . <b>NOTE</b> The Yarn service needs to be restarted to set the Yarn administrator permission so that the saved role configuration can take effect.
Setting the permission for a user to submit tasks in a specified Yarn queue	1. In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> > <b>Yarn</b> > <b>Scheduling Queue</b> > <b>root</b> . 2. In the <b>Permission</b> column of the specified queue, select <b>Submit</b> .
Setting the permission for a user to manage tasks in a specified Yarn queue	1. In the <b>Configure Resource Permission</b> table, choose <i>Name of the desired cluster</i> > <b>Yarn</b> > <b>Scheduling Queue</b> > <b>root</b> . 2. In the <b>Permission</b> column of the specified queue, select <b>Manage</b> .

If the Yarn role contains the **Submit** or **Manage** permission of a parent queue, the sub-queue inherits the permission by default, that is, the **Submit** or **Manage** permission is automatically added for the sub-queue. Permissions inherited by sub-queues will not be displayed as selected in the **Configure Resource Permission** table.

If you select only the **Submit** permission of a parent queue when setting the Yarn role, you need to manually specify the queue name when submitting tasks as a user with the permission of this role. Otherwise, when the parent queue has multiple sub-queues, the system does not automatically determine the queue to which the task is submitted and therefore submits the task to the **default** queue.

**Step 4** Click **OK**.

----End

## 30.3 Using the YARN Client

### Scenario

This section guides users to use a Yarn client in an O&M or service scenario.

### Prerequisites

- The client has been installed.  
For example, the installation directory is **/opt/client**. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite.

- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login. In common mode, you do not need to download the keytab file or change the password.

## Using the Yarn Client

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** If the cluster is in security mode, run the following command to authenticate the user. In normal mode, user authentication is not required.

```
kinit Component service user
```

**Step 5** Run the Yarn command. The following provides an example:

```
yarn application -list
```

```
----End
```

## Client-related FAQs

1. What Do I Do When the Yarn Client Exits Abnormally and Error Message "java.lang.OutOfMemoryError" Is Displayed After the Yarn Client Command Is Run?

This problem occurs because the memory required for running the Yarn client exceeds the upper limit (128 MB by default) set on the Yarn client. For clusters of MRS 3.x or later: You can modify **CLIENT\_GC\_OPTS** in *<Client installation path>/HDFS/component\_env* to change the memory upper limit of the Yarn client. For example, if you want to set the maximum memory to 1 GB, run the following command:

```
export CLIENT_GC_OPTS="-Xmx1G"
```

For clusters earlier than MRS 3.x: You can modify **GC\_OPTS\_YARN** in *<Client installation path >/HDFS/component\_env* to change the memory upper limit of the Yarn client. For example, if you want to set the maximum memory to 1 GB, run the following command:

```
export GC_OPTS_YARN="-Xmx1G"
```

After the modification, run the following command to make the modification take effect:

```
source <Client installation path>/bigdata_env
```

2. How Can I Set the Log Level When the Yarn Client Is Running?

By default, the logs generated during the running of the Yarn client are printed to the console. The default log level is INFO. To enable the DEBUG log level for fault locating, run the following command to export an environment variable:

```
export YARN_ROOT_LOGGER=DEBUG,console
```

Then run the Yarn Shell command to print DEBUG logs.

If you want to print INFO logs again, run the following command:

```
export YARN_ROOT_LOGGER=INFO,console
```

## 30.4 Configuring Resources for a NodeManager Role Instance

### Scenario

If the hardware resources (such as the number of CPU cores and memory size) of the nodes for deploying NodeManagers are different but the NodeManager available hardware resources are set to the same value, the resources may be wasted or the status may be abnormal. You need to change the hardware resource configuration for each NodeManager to ensure that the hardware resources can be fully utilized.

### Impact on the System

NodeManager role instances must be restarted for the new configuration to take effect, and the role instances are unavailable during restart.

### Prerequisites

- For versions earlier than MRS 3.x: You have logged in to the MRS management console.
- Clusters of MRS 3.x or later: You have logged in to Manager.

### Procedure

For versions earlier than MRS 3.x, perform the following operations:

- Step 1** Choose **Clusters > Active Clusters**, and click a cluster name. Choose **Components > Yarn > Instances**.
- Step 2** Click **NodeManager** in the **Role** column and go to the **Instance Configuration** tab page. Select **All** from the **Basic** drop-down list, and search for the required parameters.
- Step 3** Enter **yarn.nodemanager.resource.cpu-vcores** in the search box, and set the number of vCPUs that can be used by NodeManager on the current node. You are advised to set this parameter to 1.5 to 2 times the number of actual logical CPUs on the node. Enter **yarn.nodemanager.resource.memory-mb** in the search box, and set the physical memory size that can be used by NodeManager on the current node. You are advised to set this parameter to 75% to 90% of the actual physical memory size of the node.

 NOTE

Enter `yarn.scheduler.maximum-allocation-vcores` in the search box, and set the maximum number of available CPUs in a container. Enter `yarn.scheduler.maximum-allocation-mb` in the search box, and set the maximum available memory of a container. The instance level cannot be changed. The parameter values need to be changed in the configuration of the Yarn service, and the Yarn service needs to be restarted for the changes to take effect.

**Step 4** Click **Save Configuration**, select **Restart the affected services or instances**, and click **OK** to restart the NodeManager role instance.

**Operation succeeded** is displayed. Click **Finish**. The NodeManager role instance is started successfully.

----End

For MRS 3.x or later, perform the following operations:

**Step 1** Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Instance**.

**Step 2** Click the role instance name corresponding to the node where NodeManager is deployed, switch to **Instance Configuration**, and select **All Configurations**.

**Step 3** Enter `yarn.nodemanager.resource.cpu-vcores` in the search box, and set the number of vCPUs that can be used by NodeManager on the current node. You are advised to set this parameter to 1.5 to 2 times the number of actual logical CPUs on the node. Enter `yarn.nodemanager.resource.memory-mb` in the search box, and set the physical memory size that can be used by NodeManager on the current node. You are advised to set this parameter to 75% of the actual physical memory size of the node.

 NOTE

Enter `yarn.scheduler.maximum-allocation-vcores` in the search box, and set the maximum number of available CPUs in a container. Enter `yarn.scheduler.maximum-allocation-mb` in the search box, and set the maximum available memory of a container. The instance level cannot be changed. The parameter values need to be changed in the configuration of the Yarn service, and the Yarn service needs to be restarted for the changes to take effect.

**Step 4** Click **Save**, and then click **OK**. to restart the NodeManager role instance.

A message is displayed, indicating that the operation is successful. Click **Finish**. The NodeManager role instance is started successfully.

----End

## 30.5 Changing NodeManager Storage Directories

### Scenario

If the storage directories defined by YARN NodeManager are incorrect or the YARN storage plan changes, the MRS cluster administrator needs to modify the NodeManager storage directories on FusionInsight Manager to ensure smooth YARN running. The storage directories of NodeManager include the local storage directory `yarn.nodemanager.local-dirs` and log directory `yarn.nodemanager.log-dirs`. Changing the ZooKeeper storage directory includes the following scenarios:

- Change the storage directory of the NodeManager role. In this way, the storage directories of all NodeManager instances are changed.
- Change the storage directory of a single NodeManager instance. In this way, only the storage directory of this instance is changed, and the storage directories of other instances remain the same.

## Impact on the System

- The cluster needs to be stopped and restarted during the process of changing the storage directory of the NodeManager role, and the cluster cannot provide services before started.
- The NodeManager instance needs to be stopped and restarted during the process of changing the storage directory of the instance, and the instance at this node cannot provide services before it is started.
- The directory for storing service parameter configurations must also be updated.
- After the storage directories of NodeManager are changed, you need to download and install the client again.

## Prerequisites

- New disks have been prepared and installed on each data node, and the disks are formatted.
- New directories have been planned for storing data in the original directories.
- The MRS cluster administrator user **admin** has been prepared.

## Procedure

For versions earlier than MRS 3.x, perform the following operations:

### Step 1 Check the environment.

1. Log in to the MRS console. In the left navigation pane, choose **Clusters > Active Clusters**, and click a cluster name. Choose **Components** and check whether health status of Yarn is **Good**.
  - If yes, go to [Step 1.3](#).
  - If no, the Yarn status is unhealthy. Go to [Step 1.2](#).
2. Rectify the Yarn fault. No further action is required.
3. Determine whether to change the storage directory of the NodeManager role or that of a single NodeManager instance:
  - To change the storage directory of the NodeManager role, go to [Step 2](#).
  - To change the storage directory of a single NodeManager instance, go to [Step 3](#).

### Step 2 Change the storage directory of the NodeManager role.

1. Choose **Clusters > Active Clusters**, and click a cluster name. Choose **Components > Yarn > Stop** to stop the Yarn service.
2. Log in to the ECS server and go to each node where Yarn is installed as user **root**. Perform the following operations:

- a. Create a target directory.  
For example, to create the target directory `${BIGDATA_DATA_HOME}/data2`, run the following command:  
**mkdir `${BIGDATA_DATA_HOME}/data2`**
  - b. Mount the target directory to the new disk.  
For example, mount `${BIGDATA_DATA_HOME}/data2` to the new disk.
  - c. Modify permissions on the new directory.  
For example, to modify permissions on the `${BIGDATA_DATA_HOME}/data2` directory, run the following commands:  
**chmod 750 `${BIGDATA_DATA_HOME}/data2` -R and chown omm:wheel `${BIGDATA_DATA_HOME}/data2` -R**
3. On the MRS console, choose **Clusters > Active Clusters** and click a cluster name. Choose **Components > Yarn > Instances**. Select the NodeManager instance of the corresponding host. Choose **Instance Configuration > All Configurations**.  
Change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to the new target directory.  
For example, change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to `/srv/BigData/data2/nm/containerlogs`.
  4. Click **Save Configuration**, select **Restart the affected services or instances**, and click **OK** Restart the Yarn service.  
Click **Finish** when the system displays "Operation successful". Yarn is successfully started. No further action is required.

**Step 3** Change the storage directory of a single NodeManager instance.

1. Choose **Clusters > Active Clusters**, and click a cluster name. Choose **Components > Yarn > Instances**. Select the NodeManager instance whose storage directory needs to be modified, and choose **More > Stop Instance**.
2. Log in to the ECS and go to the NodeManager node as user `root`. Perform the following operations:
  - a. Create a target directory.  
For example, to create the target directory `${BIGDATA_DATA_HOME}/data2`, run the following command:  
**mkdir `${BIGDATA_DATA_HOME}/data2`**
  - b. Mount the target directory to the new disk.  
For example, mount `${BIGDATA_DATA_HOME}/data2` to the new disk.
  - c. Modify permissions on the new directory.  
For example, to modify permissions on the `${BIGDATA_DATA_HOME}/data2` directory, run the following commands:  
**chmod 750 `${BIGDATA_DATA_HOME}/data2` -R and chown omm:wheel `${BIGDATA_DATA_HOME}/data2` -R**
3. On the MRS console, click the specified NodeManager instance and switch to the **Instance Configuration** tab page.  
Change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to the new target directory.

For example, change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to `/srv/BigData/data2/nm/containerlogs`.

4. Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK** to restart the NodeManager instance.

Click **Finish** when the system displays "Operation successful". The NodeManager instance is successfully started.

----End

For MRS 3.x or later, perform the following operations:

#### Step 1 Check the environment.

1. Log in to Manager, choose **Cluster** > *Name of the desired cluster* > **Service** to check whether **Running Status** of Yarn is **Normal**.
  - If yes, go to **1.c**.
  - If no, the Yarn status is unhealthy. In this case, go to **1.b**.
2. Rectify faults of Yarn. No further action is required.
3. Determine whether to change the storage directory of the NodeManager role or that of a single NodeManager instance:
  - To change the storage directory of the NodeManager role, go to **2**.
  - To change the storage directory of a single NodeManager instance, go to **3**.

#### Step 2 Change the storage directory of the NodeManager role.

1. Choose **Cluster** > *Name of the desired cluster* > **Service** > **Yarn** > **Stop** to stop the Yarn service.
2. Log in to each data node where the Yarn service is installed as user **root** and perform the following operations:
  - a. Create a target directory.  
For example, to create the target directory `${BIGDATA_DATA_HOME}/data2`, run the following command:  
**mkdir `${BIGDATA_DATA_HOME}/data2`**
  - b. Mount the target directory to the new disk.  
For example, mount `${BIGDATA_DATA_HOME}/data2` to the new disk.
  - c. Modify permissions on the new directory.  
For example, to modify permissions on the `${BIGDATA_DATA_HOME}/data2` directory, run the following commands:  
**chmod 750 `${BIGDATA_DATA_HOME}/data2` -R** and **chown omm:wheel `${BIGDATA_DATA_HOME}/data2` -R**
3. On the Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Instance**. Select the NodeManager instance of the corresponding host, click **Instance Configuration**, and select **All Configurations**.

Change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to the new target directory.

For example, change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to `/srv/BigData/data2/nm/containerlogs`.

4. Click **Save**, and then click **OK**. Restart the Yarn service.  
Click **Finish** when the system displays "Operation successful". Yarn is successfully started. No further action is required.

**Step 3** Change the storage directory of a single NodeManager instance.

1. Choose **Cluster** > *Name of the desired cluster* > **Service** > **Yarn** > **Instance**, select the NodeManager instance whose storage directory needs to be modified, and choose **More** > **Stop**.
2. Log in to the NodeManager node as user **root**, and perform the following operations:
  - a. Create a target directory.  
For example, to create the target directory `${BIGDATA_DATA_HOME}/data2`, run the following command:  
**mkdir `${BIGDATA_DATA_HOME}/data2`**
  - b. Mount the target directory to the new disk.  
For example, mount `${BIGDATA_DATA_HOME}/data2` to the new disk.
  - c. Modify permissions on the new directory.  
For example, to modify permissions on the `${BIGDATA_DATA_HOME}/data2` directory, run the following commands:  
**chmod 750 `${BIGDATA_DATA_HOME}/data2` -R** and **chown omm:wheel `${BIGDATA_DATA_HOME}/data2` -R**
3. On Manager, click the specified NodeManager instance, and switch to the **Instance Configuration** page.  
Change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to the new target directory.  
For example, change the value of `yarn.nodemanager.local-dirs` or `yarn.nodemanager.log-dirs` to `/srv/BigData/data2/nm/containerlogs`.
4. Click **Save**, and then click **OK** to restart the NodeManager instance.  
Click **Finish** when the system displays "Operation successful". The NodeManager instance is successfully started.

----End

## 30.6 Configuring Strict Permission Control for Yarn

### Scenario

In the multi-tenant scenario in security mode, a cluster can be used by multiple users, and tasks of multiple users can be submitted and executed. Users are invisible to each other. A permission control mechanism is required to prevent task information of users from being obtained by other users.

For example, if user B logs in to the system and views the application list when the application submitted by user A is running, user B should not be able to view the application information of user A.



## Configuration Description

- Viewing Yarn configuration parameters

Go to the **All Configurations** page of Yarn and enter a parameter name list in [Table 30-7](#) in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-7** Parameter description

Parameter	Description	Default Value
yarn.acl.enable	Whether to enable Yarn permission control	true
yarn.webapp.filter-entity-list-by-user	Whether to enable the strict view function. After this function is enabled, a login user can view only the content that the user has the permission to view. To enable this function, set <b>yarn.acl.enable</b> to <b>true</b> . <b>NOTE</b> This parameter applies to clusters of MRS 3.x or later.	true

- Viewing MapReduce configuration parameters

Go to the **All Configurations** page of MapReduce and enter a parameter name in [Table 30-8](#) in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-8** Parameter description

Parameter	Description	Default Value
mapreduce.cluster.acls.enabled	Whether to enable permission control of MapReduce JobHistoryServer This parameter is a client parameter and takes effect after permission control is enabled on the JobHistoryServer server.	true

Parameter	Description	Default Value
yarn.webapp.filter-entity-list-by-user	<p>Whether to enable the strict view of MapReduce JobHistoryServer. After the strict view is enabled, a login user can view only the content that the user has the permission to view. This parameter is a server parameter of JobHistoryServer. It indicates that permission control is enabled for JHS. However, whether to control a specific application is determined by the client parameter <b>mapreduce.cluster.acls.enabled</b>.</p> <p><b>NOTE</b> This parameter applies to clusters of MRS 3.x or later.</p>	true

#### NOTICE

The preceding configurations affect the RESTful API and Shell command results. After the preceding configurations are enabled, the return results of RESTful API calls and shell commands contain only the information that the user has the permission to view.

If **yarn.acl.enable** or **mapreduce.cluster.acls.enabled** is set to **false**, the Yarn or MapReduce permission verification function is disabled. In this case, any user can submit tasks and view task information on Yarn or MapReduce, which poses security risks. Exercise caution when performing this operation.

## 30.7 Configuring Container Log Aggregation

### Scenario

Yarn provides the container log aggregation function to collect logs generated by containers on each node to HDFS to release local disk space. You can collect logs in either of the following ways:

- After the application is complete, collect container logs to HDFS at a time.
- During application running, periodically collect log segments generated by containers and save them to HDFS.

### Configuration Description

**Navigation path for setting parameters:**

Go to the **All Configurations** tab page of YARN, enter the parameters listed in [Table 30-9](#) in the search box, modify the parameters by referring to [Modifying](#)

[Cluster Service Configuration Parameters](#), and save the configuration. On the **Dashboard** tab page, choose **More > Synchronize Configuration**. After the synchronization is complete, restart the YARN service.

The `yarn.nodemanager.remote-app-log-dir-suffix` parameter must be configured on the Yarn client. The configurations on the ResourceManager, NodeManager, and JobHistory nodes must be the same as those on the Yarn client.

The periodic log collection function applies only to MapReduce applications, for which rolling output of log files must be configured. [Table 30-11](#) describes the configurations in the `mapred-site.xml` configuration file on the MapReduce client node.

**Table 30-9** Parameter description

Parameter	Description	Default Value
yarn.log-aggregation-enable	<p>Whether to enable container log aggregation</p> <ul style="list-style-type: none"> <li>If this parameter is set to <b>true</b>, logs are collected to the HDFS directory.</li> <li>If this parameter is set to <b>false</b>, the function is disabled, and logs are not collected to HDFS.</li> </ul> <p>After changing the parameter value, restart the Yarn service for the setting to take effect.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>The container logs that are generated before the parameter is set to <b>false</b> and the setting takes effect cannot be obtained from the web UI.</li> <li>If you need to view the logs generated before on the web UI, you are advised to set this parameter to <b>true</b>.</li> </ul>	true

Parameter	Description	Default Value
yarn.nodemanager.log-aggregation.rolling-monitoring-interval-seconds	<p>Interval for NodeManager to periodically collect logs</p> <ul style="list-style-type: none"><li>• If this parameter is set to <b>-1</b> or <b>0</b>, periodic log collection is disabled. Logs are collected at a time after application running is complete.</li><li>• The minimum collection interval can be set to 3,600 seconds. If this parameter is set to a value greater than 0 and less than 3,600, the collection interval is 3,600 seconds.</li></ul> <p>Interval for NodeManager to wake up and upload logs. If this parameter is set to <b>-1</b> or <b>0</b>, rolling monitoring is disabled and logs are aggregated when the application task is complete. The value must be greater than or equal to <b>-1</b>.</p>	-1

Parameter	Description	Default Value
<code>yarn.nodemanager.disk-health-checker.log-dirs.max-disk-utilization-per-disk-percentage</code>	<p>Maximum percentage of the Yarn disk quota that can be occupied by the container log directory on each disk. When the space occupied by the log directory exceeds the value of this parameter, the periodic log collection service is triggered to start a log collection activity beyond the period to release the local disk space. Maximum space for container logs that can be provided on each disk. If the disk space occupied by container logs exceeds this threshold, data aggregation in rolling mode is triggered.</p> <ul style="list-style-type: none"> <li>For clusters of versions earlier than MRS 3.x: The valid value range of the maximum disk quota percentage is 0 to 100. If the value is less than or equal to <b>0</b>, it is forcibly reset to <b>25</b>. If the value is greater than 100, the value is forcibly reset to <b>25</b>.</li> <li>For clusters of MRS 3.x or later: The valid value range of the maximum disk quota percentage is -1 to 100. If the value is less than -1, it is forcibly reset to <b>25</b>. If the value is greater than 100, the value is forcibly reset to <b>25</b>. If you set the value to <b>-1</b>, the disk capacity detection function for Container log directory is disabled.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>Percentage of the available disk space of the container log directory = Percentage of the available disk space of Yarn (<b>yarn.nodemanager.disk-health-checker.max-disk-utilization-per-disk-percentage</b>) x Percentage of the available disk space of the container log directory (<b>yarn.nodemanager.disk-health-checker.log-dirs.max-disk-utilization-per-disk-percentage</b>)</li> <li>Only applications with the periodic log collection function enabled can trigger log collection when the disk quota of the log directory exceeds the threshold.</li> </ul>	25

Parameter	Description	Default Value
yarn.nodemanager.remote-app-log-dir-suffix	Name of the HDFS folder in which container logs are to be stored. This parameter and <b>yarn.nodemanager.remote-app-log-dir</b> form the full path for storing container logs. That is, <b>{yarn.nodemanager.remote-app-log-dir}/{user}/{yarn.nodemanager.remote-app-log-dir-suffix}</b> . <b>NOTE</b> <i>{user}</i> indicates the username for running the task.	logs
yarn.nodemanager.log-aggregator.on-fail.retain-log-in-sec	Duration for retaining container logs on the local host after the logs fail to be collected, in second <ul style="list-style-type: none"> <li>• If this parameter is set to <b>0</b>, local logs are deleted immediately.</li> <li>• If this parameter is set to a positive number, local logs are retained for this period.</li> </ul>	604800

Go to the **All Configurations** page of MapReduce and enter a parameter name in [Table 30-10](#) in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-10** Parameter description

Parameter	Description	Default Value
yarn.log-aggregation.retain-seconds	Duration for retaining aggregated logs, in second <ul style="list-style-type: none"> <li>• If this parameter is set to <b>-1</b>, the container logs will be retained permanently in the HDFS.</li> <li>• If this parameter is set to <b>0</b> or a positive integer, container logs will be stored for such a period and deleted after the period expires.</li> </ul> <b>NOTE</b> A short period may increase load of the NameNode. Therefore, you are advised to set this parameter to a proper value.	1296000

Parameter	Description	Default Value
yarn.log-aggregation.retain-check-interval-seconds	<p>Interval for storing container logs in HDFS, in second</p> <ul style="list-style-type: none"> <li>If this parameter is set to <b>-1</b> or <b>0</b>, the interval will be one tenth of the period specified by <b>yarn.log-aggregation.retain-seconds</b>.</li> </ul> <p><b>NOTE</b> If this parameter is set to <b>-1</b> or <b>0</b>, <b>yarn.log-aggregation.retain-seconds</b> cannot be set to <b>0</b>.</p> <ul style="list-style-type: none"> <li>If this parameter is set to a positive number, container logs in HDFS will be scanned at such an interval.</li> </ul> <p><b>NOTE</b> A short interval may increase load of the NameNode. Therefore, you are advised to set this parameter to a proper value.</p>	86400

Go to the **All Configurations** page of Yarn and enter a parameter name list in [Table 30-11](#) in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-11** Configuring rolling output of MapReduce application log files

Parameter	Description	Default Value
mapreduce.task.userlog.limit.kb	Maximum size of a single task log file of the MapReduce application. When the maximum size of the log file has been reached, a new log file is generated. The value <b>0</b> indicates that the size of the log file is not limited.	51200

Parameter	Description	Default Value
yarn.app.mapreduce.task.container.log.backups	<p>Maximum number of task logs that can be retained for the MapReduce application. If this parameter is set to <b>0</b>, rolling output is disabled.</p> <p>Number of task log backup files when ContainerRollingLogAppender (CRLA) is used. By default, ContainerLogAppender (CLA) is used and container logs are not rolled back.</p> <p>When both <b>mapreduce.task.userlog.limit.kb</b> and <b>yarn.app.mapreduce.task.container.log.backups</b> are greater than 0, CRLA is enabled. The value ranges from 0 to 999.</p>	10
yarn.app.mapreduce.am.container.log.limit.kb	<p>Maximum size of a single ApplicationMaster log file of the MapReduce application, in KB. When the maximum size of the log file has been reached, a new log file is generated. The value <b>0</b> indicates that the size of a single ApplicationMaster log file is not limited.</p>	51200
yarn.app.mapreduce.am.container.log.backups	<p>Maximum number of ApplicationMaster logs that can be retained for the MapReduce application. If this parameter is set to <b>0</b>, rolling output is disabled. Number of ApplicationMaster log backup files when CRLA is used. By default, CLA is used and container logs are not rolled back.</p> <p>When both <b>yarn.app.mapreduce.am.container.log.limit.kb</b> and <b>yarn.app.mapreduce.am.container.log.backups</b> are greater than 0, CRLA is enabled for the ApplicationMaster. The value ranges from 0 to 999.</p>	20
yarn.app.mapreduce.shuffle.log.backups	<p>Maximum number of shuffle logs that can be retained for the MapReduce application. If this parameter is set to <b>0</b>, rolling output is disabled.</p> <p>When both <b>yarn.app.mapreduce.shuffle.log.limit.kb</b> and <b>yarn.app.mapreduce.shuffle.log.backups</b> are greater than 0, <b>syslog.shuffle</b> uses CRLA. The value ranges from 0 to 999.</p>	10



Parameter	Description	Default Value
yarn.app.mapreduce.shuffle.log.limit.kb	Maximum size of a single shuffle log file of the MapReduce application, in KB. When the maximum size of the log file has been reached, a new log file is generated. If this parameter is set to <b>0</b> , the size of a single shuffle log file is not limited. The value must be greater than or equal to <b>0</b> .	51200

## 30.8 Using CGroups with YARN

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

### Scenario

CGroups is a Linux kernel feature. In YARN this feature allows containers to be limited in their resource usage (example, CPU usage). Without CGroups, it is hard to limit the container CPU usage. Without CGroups, it is hard to limit the container CPU usage.

 **NOTE**

Currently, CGroups is only used for limiting the CPU usage.

### Configuration Description

CGroups is a Linux kernel feature and is enabled using LinuxContainerExecutor. For details about how to configure the LinuxContainerExecutor for security, see the official website. You can learn the file system permissions assigned to users and user groups:

 **NOTE**

- Do not modify users, user groups, and related permissions of various paths in the corresponding file system. Otherwise, functions of CGroups may become abnormal.
- If the parameter value of **yarn.nodemanager.resource.percentage-physical-cpu-limit** is too small, the number of available cores may be less than one. For example, if the parameter of a four-core node is set to 20%, the number available core is less than one. As a result, all cores will be used. The Quota mode can be used in Linux versions, for example, Cent OS, that do not support Quota mode.

The table below describes the parameter for configuring cpuset mode, that is, only configured CPUs can be used by YARN.

**Table 30-12** Parameter description

Parameter	Description	Default Value
yarn.nodemanager.linux-container-executor.cgroups.cpu-set-usage	Whether to enable the cpuset mode. If this parameter is set to <b>true</b> , the cpuset mode is enabled.	false

The table below describes the parameters for configuring the strictcpuset mode, that is, only configured CPUs can be used by containers.

**Table 30-13** Parameter description

Parameter	Description	Default Value
yarn.nodemanager.linux-container-executor.cgroups.cpu-set-usage	Whether to enable the cpuset mode. If this parameter is set to <b>true</b> , the cpuset mode is enabled.	false
yarn.nodemanager.linux-container-executor.cgroups.cpuset.strict.enabled	Whether containers use allocated CPUs. If this parameter is set to <b>true</b> , the container can use the allocated CPUs.	false

To switch from cpuset mode to quota mode, the following conditions must be met:

- Set the **yarn.nodemanager.linux-container-executor.cgroups.cpu-set-usage** parameter to **false**.
- Delete the **container** folder (if any) from the **/sys/fs/cgroup/cpuset/hadoop-yarn/** directory.
- Delete all CPUs configured in the **cpuset.cpus** file in **/sys/fs/cgroup/cpuset/hadoop-yarn/**.

## Procedure

- Step 1** Log in to Manager. Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** and select **All Configurations**.
- Step 2** In the navigation pane on the left, choose **NodeManager** > **Customization** and find the **yarn-site.xml** file.
- Step 3** Add the parameters in [Table 30-12](#) and [Table 30-13](#) as user-defined parameters.

Based on the configuration files and parameter functions, locate the row where parameter **yarn-site.xml** resides. Enter the parameter name in the **Name** column and enter the parameter value in the **Value** column.

Click + to add a customized parameter.

- Step 4** Click **Save**. In the displayed **Save Configuration** dialog box, confirm the modification and click **OK**. Click **Finish** when the system displays "Operation succeeded". The configuration is successfully saved.

After the configuration is saved, restart the Yarn service whose configuration has expired for the configuration to take effect.

----End

## 30.9 Configuring the Number of ApplicationMaster Retries

### Scenario

When resources are insufficient or ApplicationMaster fails to start, a client probably encounters running errors.

### Configuration Description

Go to the **All Configurations** page of Yarn and enter a parameter name list in [Table 30-14](#) in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-14** Parameter description

Parameter	Description	Default Value
yarn.resource manager.am.max-attempts	Number of retries of the ApplicationMaster. Increasing the number of retries can prevent ApplicationMaster startup failures caused by insufficient resources. This applies to global settings of all ApplicationMasters. Each ApplicationMaster can use an API to set an independent maximum number of retries. However, the number of retries cannot be greater than the global maximum number of retries. If the value is greater than the global maximum number of retries, the ResourceManager overwrites the value to allow at least one retry. The value must be greater than or equal to 1.	5

## 30.10 Configure the ApplicationMaster to Automatically Adjust the Allocated Memory

This section applies to clusters of MRS 3.x or later.

### Scenario

During the process of starting the configuration, when the ApplicationMaster creates a container, the allocated memory is automatically adjusted according to the total number of tasks, which makes resource utilization more flexible and improves the fault tolerance of the client application.

### Configuration Description

#### Navigation path for setting parameters:

On Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Yarn** > **Configuration**. On the displayed page, select **All Configurations** and enter **mapreduce.job.am.memory.policy**.

#### Configuration description

If the default value of the parameter is left empty. In this case, the automatic adjustment policy is not enabled. The memory of ApplicationMaster is still affected by the value of **yarn.app.mapreduce.am.resource.mb**.

The value of **mapreduce.job.am.memory.policy** consists of five items, and they are separated by colons (:) and commas (,) in the following format: **baseTaskCount:taskStep:memoryStep,minMemory:maxMemory**. The format is strictly checked when the value is entered.

**Table 30-15** Parameter description

Parameter	Description	Setting Requirement
baseTaskCount	Indicates the total number of tasks. The configuration of ApplicationMaster is valid only when the total number of tasks (on the sum of the Map and Reduce ends) is greater than or equal to the value of this parameter.	The value cannot be empty and must be greater than 0.
taskStep	Indicates the incremental step length of tasks. This parameter and <b>memoryStep</b> determine the memory adjustment amount.	The value cannot be empty and must be greater than 0.
memoryStep	Indicates the incremental memory step. The memory capacity is increased based on the value of <b>yarn.app.mapreduce.am.resource.mb</b> .	The value cannot be empty and must be greater than 0. The unit is MB.

Parameter	Description	Setting Requirement
minMemory	Indicates the lower limit of the memory that can be automatically adjusted. If the memory after the automatic adjustment is less than or equal to the value of this parameter, the value of <b>yarn.app.mapreduce.am.resource.mb</b> is used.	The value cannot be empty. It must be greater than 0 and cannot be greater than the value of <b>maxMemory</b> . Unit: MB
maxMemory	Indicates the upper limit of memory that can be automatically adjusted. If the adjusted memory exceeds the upper limit, use this value as the final value.	The value cannot be empty. It must be greater than 0 and cannot be less than the value of <b>minMemory</b> . Unit: MB

## Example Value

Configuration:

- yarn.app.mapreduce.am.resource.mb=1536
- mapreduce.job.am.memory.policy=100:10:50,1200:2000
- Total number of tasks of an application =120

The calculation process is as follows:

Memory after adjustment =  $1536 + [(120 - 100)/10] \times 50 = 1636$ . In this example, memory after adjustment 1636 is greater than the value of **minMemory 1200**, and less than the value of **maxMemory 2000**. Therefore, the ApplicationMaster memory is set to **1636 MB**.

If the value of **memStep** is changed to **250**, the calculation formula is as follows: Memory after adjustment =  $1536 + [(120 - 100) / 10] \times 250 = 2136$ . In this case, the memory after adjustment is greater than the value of **maxMemory 2000**. As a result, the value of **ApplicationMaster** is set to **2000 MB**.

### NOTE

If the memory after adjustment is lower than the value of **minMemory**, the configuration does not take effect but the value is still printed on the backend server. This value is provided as the reference for adjusting the value of **minMemory**.

## 30.11 Configuring the Access Channel Protocol

### Scenario

The value of the **yarn.http.policy** parameter must be consistent on both the server and clients. Web UIs on clients will be garbled if an inconsistency exists, for

example, the parameter value is **HTTPS\_ONLY** on the server but it is left unspecified on a client (the parameter value **HTTP\_ONLY** is applied to the client by default). Set the **yarn.http.policy** parameters on the clients and server to prevent garbled characters from being displayed on the clients.

## Procedure

**Step 1** On Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations**. On the displayed page, select **All Configurations** and enter **yarn.http.policy**.

- In security mode, set this parameter to **HTTPS\_ONLY**.
- In normal mode, set this parameter to **HTTP\_ONLY**.

**Step 2** Log in to the node where the client is installed as the client installation user.

**Step 3** Run the following command to switch to the client installation directory:

```
cd /opt/client
```

**Step 4** Run the following command to edit the **yarn-site.xml** file:

```
vi Yarn/config/yarn-site.xml
```

Change the value of **yarn.http.policy**.

In security mode, set this parameter to **HTTPS\_ONLY**.

In normal mode, set this parameter to **HTTP\_ONLY**.

**Step 5** Run the **:wq** command to save execution.

**Step 6** Restart the client for the settings to take effect.

----End

## 30.12 Configuring Memory Usage Detection

### Scenario

If memory usage of the submitted application cannot be estimated, you can modify the configuration on the server to determine whether to check the memory usage.

If the memory usage is not checked, the container occupies the memory until the memory overflows. If the memory usage exceeds the configured memory size, the corresponding container is killed.

### Configuration Description

Go to the **All Configurations** page of Yarn and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-16** Parameter description

Parameter	Description	Default Value
yarn.nodemanager.vmem-check-enabled	Whether to enable virtual memory usage detection. If the memory used by a task exceeds the allocated memory size, the task is forcibly stopped. <ul style="list-style-type: none"> <li>If the value is <b>true</b>, the virtual memory will be checked.</li> <li>If the value is <b>false</b>, the virtual memory will not be checked.</li> </ul>	For versions earlier than MRS 3.x: false For MRS 3.x or later: true
yarn.nodemanager.pmem-check-enabled	Whether to enable physical memory usage detection. If the memory used by a task exceeds the allocated memory size, the task is forcibly stopped. <ul style="list-style-type: none"> <li>If the value is <b>true</b>, the physical memory will be checked.</li> <li>If the value is <b>false</b>, the physical memory will not be checked.</li> </ul>	true

## 30.13 Configuring the Additional Scheduler WebUI

### Scenario

If the custom scheduler is set in ResourceManager, you can set the corresponding web page and other Web applications for the custom scheduler.

### Configuration Description

Go to the **All Configurations** page of Yarn and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-17** Configuring the Additional Scheduler WebUI

Parameter	Description	Default Value
hadoop.http.rmwebapp.scheduler.page.classes	Load the corresponding web page for the custom scheduler on the RM WebUI. This parameter is valid only when <b>yarn.resourcemanager.scheduler.class</b> is set to a custom scheduler.	-
yarn.http.rmwebapp.external.classes	Load the custom web application in the RM Web service.	-

## 30.14 Configuring Yarn Restart

### Scenario

The Yarn Restart feature includes ResourceManager Restart and NodeManager Restart.

- When ResourceManager Restart is enabled, the new active ResourceManager node loads the information of the previous active ResourceManager node, and takes over container status information on all NodeManager nodes to continue service running. In this way, status information can be saved by periodically executing checkpoint operations, avoiding data loss.
- When NodeManager Restart is enabled, NodeManager locally saves information about containers running on the node. After NodeManager is restarted, the container running progress on the node will not be lost by restoring the saved status information.

### Configuration Description

Go to the **All Configurations** page of Yarn and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

Configure ResourceManager Restart as follows:

**Table 30-18** Parameter description of ResourceManager Restart

Parameter	Description	Default Value
yarn.resourcemanager.recovery.enabled	Whether to enable ResourceManager to restore the status after startup. If this parameter is set to <b>true</b> , <b>yarn.resourcemanager.store.class</b> must also be set.	true
yarn.resourcemanager.store.class	State-store class used to store the application and task statuses and certificate content.	For clusters of versions earlier than MRS 3.x: <b>org.apache.hadoop.yarn.server.resourcemanager.recovery.ZKRMStateStore</b>  For clusters of MRS 3.x or later: org.apache.hadoop.yarn.server.resourcemanager.recovery.AsyncZKRMStateStore



Parameter	Description	Default Value
yarn.resourcemanager.zk-state-store.parent-path	Directory for storing ZKRMStateStore in ZooKeeper	/rmstore
yarn.resourcemanager.work-preserving-recovery.enabled	Whether to enable ResourceManager work serving. This configuration is used only for Yarn feature verification.	true
yarn.resourcemanager.state-store.async.load	Whether to apply asynchronous restoration to completed applications.	For clusters of versions earlier than MRS 3.x: <b>false</b> For MRS 3.x or later: <b>true</b>
yarn.resourcemanager.zk-state-store.num-fetch-threads	If asynchronous restoration is enabled, increasing the number of working threads can speed up the restoration of task information stored in ZooKeeper. The value must be greater than 0.	For clusters of versions earlier than MRS 3.x: <b>1</b> For MRS 3.x or later: <b>20</b>

Configure NodeManager Restart as follows:

**Table 30-19** Parameter description of NodeManager Restart

Parameter	Description	Default Value
yarn.nodemanager.recovery.enabled	Whether to enable the function of collecting logs upon a log collection failure when NodeManager is restarted and whether to restore the unfinished application	true
yarn.nodemanager.recovery.dir	Local directory used by NodeManager to store container status It applies to clusters of MRS 3.x or later.	\${SRV_HOME}/tmp/yarn-nm-recovery
yarn.nodemanager.recovery.supervised	Whether NodeManager is monitored. After this parameter is enabled, NodeManager does not clear containers after exit. NodeManager assumes that it will restart and restore containers immediately.	true

## 30.15 Configuring ApplicationMaster Work Preserving

This section applies to clusters of MRS 3.x or later.

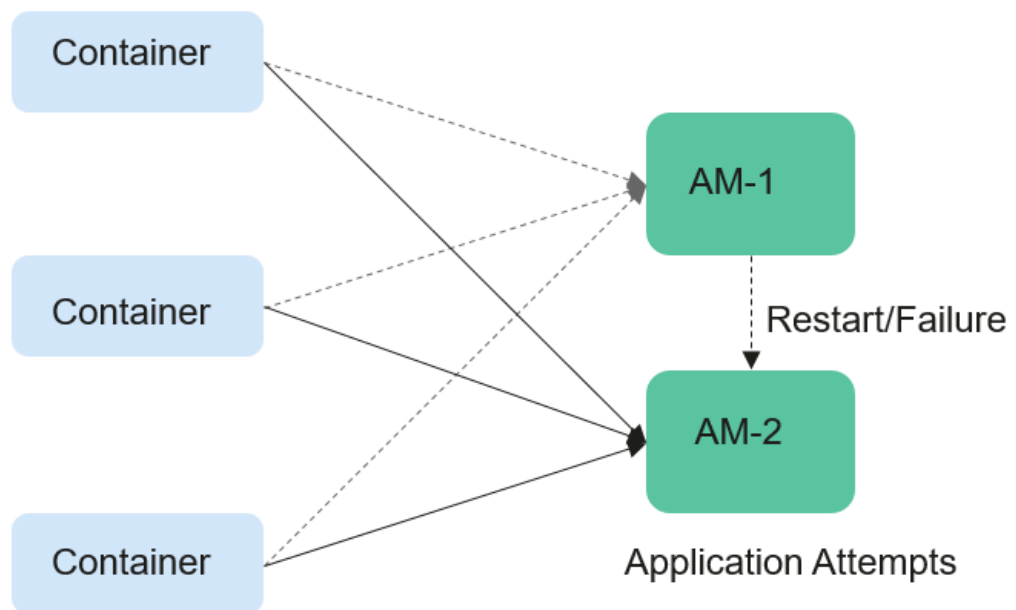
### Scenario

In YARN, ApplicationMasters run on NodeManagers just like every other container (ignoring unmanaged ApplicationMasters in this context). ApplicationMasters may break down, exit, or shut down. If an ApplicationMaster node goes down, ResourceManager kills all the containers of ApplicationAttempt, including containers running on NodeManager. ResourceManager starts a new ApplicationAttempt node on another compute node.

For different types of applications, we want to handle ApplicationMaster restart events in different ways. MapReduce applications aim to prevent task loss but allow the loss of the currently running container. However, for the long-period YARN service, users may not want the service to stop due to the ApplicationMaster fault.

YARN can retain the status of the container when a new ApplicationAttempt is started. Therefore, running jobs can continue to operate without faults.

**Figure 30-1** ApplicationMaster job preserving



### Configuration Description

Go to the **All Configurations** page of Yarn and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

Set the following parameters based on [Table 30-20](#).

**Table 30-20** Parameter description

Parameter	Description	Default Value
yarn.app.mapreduce.am.work-preserve	Whether to enable the ApplicationMaster job retention feature.	false
yarn.app.mapreduce.am.umbilical.max.retries	Maximum number of attempts to restore a running container in the ApplicationMaster job retention feature.	5
yarn.app.mapreduce.am.umbilical.retry.interval	Specifies the interval at which a running container attempts to recover in the ApplicationMaster job retention feature. Unit: millisecond	10000
yarn.resourcemanager.am.max-attempts	The number of retries of ApplicationMaster. Increasing the number of retries prevents ApplicationMaster startup failures caused by insufficient resources.  This applies to global settings of all ApplicationMasters. Each ApplicationMaster can use an API to set an independent maximum number of retries. However, the number of retries cannot be greater than the global maximum number of retries. If the value is greater than the global maximum number of retries, the ResourceManager overwrites the value. The value must be greater than or equal to 1.	2

## 30.16 Configuring the Localized Log Levels

This section applies to clusters of MRS 3.x or later.

### Scenarios

The default log level of localized container is **INFO**. You can change the log level by configuring **yarn.nodemanager.container-localizer.java.opts**.

### Configuration Description

On Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Yarn** > **Configuration**. Select **All Configurations** and set the following parameters in the configuration file **yarn-site.xml** of NodeManager to change the log level.

**Table 30-21** Parameter description

Parameter	Description	Default Value
yarn.nodemanager.container-localizer.java.opts	The additional <b>jvm</b> parameters are provided for the localized container process.	-Xmx256m -Djava.security.krb5.conf=\${KRB5_CONFIG}

The default value is **-Xmx256m -Djava.security.krb5.conf=\${KRB5\_CONFIG}** and the default log level is info. To change the localized log level of the container, add the following content:

```
-Dhadoop.root.logger=<LOG_LEVEL>,localizationCLA
```

**Example:**

To change the local log level to **DEBUG**, set the parameter as follows:

```
-Xmx256m -Dhadoop.root.logger=DEBUG,localizationCLA
```

 **NOTE**

Allowed log levels are as follows: FATAL, ERROR, WARN, INFO, DEBUG, TRACE, and ALL.

## 30.17 Configuring Users That Run Tasks

This section applies to clusters of MRS 3.x or later.

### Scenario

Currently, YARN allows the user that starts the NodeManager to run the task submitted by all other users, or the users to run the task submitted by themselves.

### Configuration Description

On Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations**. Click **All Configurations** Enter a parameter name in the search box.

**Table 30-22** Parameter description

Parameter	Description	Default Value
yarn.nodemanager.linux-container-executor.user	Indicates the user who runs a task.	The value is left blank by default. <b>NOTE</b> The value is left blank by default. The user who submits a task is the actual person who runs the task.

Parameter	Description	Default Value
yarn.nodemanager.container-executor.class	Indicates the executor who starts a task.	org.apache.hadoop.yarn.server.nodemanager.EnhancedLinuxContainerExecutor

 NOTE

- Set **yarn.nodemanager.linux-container-executor.user** to configure the user who runs the container. This parameter is left blank by default. The user who submits the task is the person who runs the container. This parameter is valid only when **yarn.nodemanager.container-executor.class** is set to **org.apache.hadoop.yarn.server.nodemanager.EnhancedLinuxContainerExecutor**.
- In non-security mode, if **yarn.nodemanager.linux-container-executor.user** is set to **omm**, **yarn.nodemanager.linux-container-executor.nonsecure-mode.local-user** must also be set to **omm**.
- For security reasons, it is advised to retain the default values of **yarn.nodemanager.linux-container-executor.user** and **yarn.nodemanager.container-executor.class**.

## 30.18 Yarn Log Overview

### Log Description

The default paths for saving Yarn logs are as follows:

- ResourceManager: **/var/log/Bigdata/yarn/rm** (run logs) and **/var/log/Bigdata/audit/yarn/rm** (audit logs)
- NodeManager: **/var/log/Bigdata/yarn/nm** (run logs) and **/var/log/Bigdata/audit/yarn/nm** (audit logs)

Log archive rule: The automatic compression and archive function is enabled for Yarn logs. By default, when the size of a log file exceeds 50 MB, the log file is automatically compressed. The naming rule of the compressed log file is as follows: *<Original log file name>-<yyyy-mm-dd\_hh-mm-ss>.[ID].log.zip*. A maximum of 100 latest compressed files are retained. The number of compressed files can be configured on Manager.

**Log archive rule:**

**Table 30-23** Yarn log list

Log Type	Log File Name	Description
Run log	hadoop-<SSH_USER>-<process_name>-<hostname>.log	Yarn component log file that records most of the logs generated when the Yarn component is running

Log Type	Log File Name	Description
	hadoop-<SSH_USER>-<process_name>-<hostname>.out	Log file that records Yarn running environment information
	<process_name>-<SSH_USER>-<DATE>-<PID>-gc.log	Garbage collection log file
	yarn-haCheck.log	ResourceManager active/standby status detection log file
	yarn-service-check.log	Log file that records the health check details of the Yarn service
	yarn-start-stop.log	Log file that records the startup and stop of the Yarn service
	yarn-prestart.log	Log file that records cluster operations before the Yarn service startup
	yarn-postinstall.log	Work log file after installation and before startup of the Yarn service
	hadoop-commission.log	Yarn service entry log file
	yarn-cleanup.log	Log file that records the cleanup operation during uninstallation of the Yarn service
	yarn-refreshqueue.log	Yarn queue refresh log file
	upgradeDetail.log	Upgrade log file
	stderr/stdin/syslog	Container log file of the applications running on the Yarn service
	yarn-application-check.log	Check log file of applications running on the Yarn service
	yarn-appsummary.log	Running result log file of applications running on the Yarn service
	yarn-switch-resourcemanager.log	Run log file that records the Yarn active/standby switchover

Log Type	Log File Name	Description
	ranger-yarn-plugin-enable.log	Log file that records the enabling of Ranger authentication for Yarn
	yarn-nodemanager-period-check.log	Periodic check log of Yarn NodeManager
	yarn-resourcemanager-period-check.log	Periodic check log of Yarn ResourceManager
	hadoop.log	Hadoop client logs
	env.log	Environment information log file before the instance is started or stopped.
Audit logs	yarn-audit-<process_name>.log ranger-plugin-audit.log	Yarn operation audit log file
	SecurityAuth.audit	Yarn security audit log file

## Log Level

**Table 30-24** describes the log levels supported by Yarn, including OFF, FATAL, ERROR, WARN, INFO, and DEBUG, from high priority to low. Logs whose levels are higher than or equal to the specified level are printed. The number of printed logs decreases as the specified log level increases.

**Table 30-24** Log levels

Level	Description
FATAL	Logs of this level record critical error information about the current event processing.
ERROR	Logs of this level record error information about the current event processing.
WARN	Logs of this level record exception information about the current event processing.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system as well as system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of the Yarn service by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Click **Save Configuration**. In the dialog box that is displayed, click **OK** to make the setting take effect.

 **NOTE**

The configurations take effect immediately without the need to restart the service.

----End

## Log Format

The following table lists the Yarn log formats.

**Table 30-25** Log formats

Log Type	Format	Example
Run log	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level>  <Thread that generates the log> <Message in the log>  <Location of the log event>	2014-09-26 14:18:59,109   INFO   main   Client environment:java.compiler= <NA>   org.apache.zookeeper.Enviro nment.logEnv(Environment. java:100)
Audit log	<yyyy-MM-dd HH:mm:ss,SSS> <Log Level>  <Thread that generates the log> <Message in the log>  <Location of the log event>	2014-09-26 14:24:43,605   INFO   main-EventThread   USER=omm OPERATION=refreshAdmin Acls TARGET=AdminService RESULT=SUCCESS   org.apache.hadoop.yarn.ser ver.resourcemanager.RMAu ditLogger\$LogLevel \$6.printLog(RMAuditLogger. java:91)

## 30.19 Yarn Performance Tuning

### 30.19.1 Preempting a Task

#### Scenario

The capacity scheduler of ResourceManager implements job preemption to simplify job running in queues and improve resource utilization. The process is as follows:



1. Assume that there are two queues (Queue A and Queue B). The capacity of Queue A is 25%, and the capacity of Queue B is 75%.
2. In the initial state, Task 1 is distributed to Queue A for processing, requiring 75% cluster resources. Task 2 is distributed to Queue B for processing, requiring 50% cluster resources.
3. Task 1 uses 25% cluster resources provided by Queue A and 50% resources from Queue B. Queue B reserves 25% cluster resources.
4. If task preemption is enabled, the resources of Task 1 will be preempted. Queue B preempts 25% cluster resources from Queue A for Task 2.
5. Task 1 will be executed when Task 2 is complete and the cluster has sufficient resources.

## Procedure

Navigation path for setting parameters:

Go to the **All Configurations** page of Yarn and enter a parameter name in the search box by referring to [Modifying Cluster Service Configuration Parameters](#).

**Table 30-26** Parameter description

Parameter	Description	Default Value
yarn.resourcemanager.scheduler.monitor.enable	Whether to start scheduler monitoring according to <b>yarn.resourcemanager.scheduler.monitor.policies</b> . If this parameter is set to <b>true</b> , scheduler monitoring is enabled based on policies specified by <b>yarn.resourcemanager.scheduler.monitor.policies</b> and task resource preemption is enabled based on the scheduler information. If this parameter is set to <b>false</b> , scheduler monitoring is disabled.	false
yarn.resourcemanager.scheduler.monitor.policies	List of the SchedulingEditPolicy class to be used with the scheduler	org.apache.hadoop.yarn.server.resourcemanager.monitor.capacity.ProportionalCapacityPreemptionPolicy

Parameter	Description	Default Value
yarn.resourcemanager.monitor.capacity.preemption.observe_only	<ul style="list-style-type: none"> <li>If this parameter is set to <b>true</b>, policies will be applied but task resource preemption will not be performed.</li> <li>If this parameter is set to <b>false</b>, policies will be applied and task resource preemption will be performed based on the policies.</li> </ul>	false
yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval	Monitoring interval, in millisecond. If this parameter is set to a larger value, capacity detection will not be performed frequently.	3000
yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill	Interval between the time when a resource preemption request is sent and the time when the container is stopped (resources are released), in millisecond. The value must be greater than or equal to <b>0</b> . By default, if ApplicationMaster does not stop the container within 15 seconds, ResourceManager will forcibly stop the container after 15 seconds.	15000
yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round	Maximum resource preemption ratio in a period. This value can be used to limit the speed at which containers are reclaimed from the cluster. After the expected total preemption value is calculated, the policy scales the preemption ratio back to this limit.	0.1
yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity	Resource preemption dead zone = Total number of resources in the cluster x Value of this configuration item + Original resources of a queue (for example, Queue A). When resources actually used by a task in Queue A exceeds the preemption dead zone, the resource beyond the preemption dead zone is preempted. The value range is 0 to 1. <b>NOTE</b> A smaller value is recommended for effective preemption.	0

Parameter	Description	Default Value
yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor	<p>Preemption percentage. Containers preempt only this percentage of the resources.</p> <p>For example, a termination factor of 0.5 will reclaim almost 95% of resources within 5 times of <b>yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill</b>, even in the absence of natural termination. That is, 5 consecutive preemptions will be performed and each time half of the target resources will be preempted. The trend is geometric convergence. The interval of each preemption is <b>yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill</b>. The value range is 0 to 1.</p>	1

## 30.19.2 Setting the Task Priority

### Scenario

The resource contention scenarios of a cluster are as follows:

1. Submit two jobs (Job 1 and Job 2) with lower priorities.
2. Some tasks of running Job 1 and Job 2 are in the running state. However, some tasks are pending due to resource deficiency because the capacity of cluster or queue resources is limited.
3. Submit a job (Job 3) with a higher priority. In this case, after the running tasks of Job 1 and Job 2 are complete, their resources will be released and then allocated to the pending tasks of Job 3.
4. After Job 3 is complete, its resources will be released and then allocated to Job 1 and Job 2.

Users can use capacity scheduler of ResourceManager to set the task priority in Yarn because the task priority is implemented by the scheduler of ResourceManager.

### Procedure

Set the **mapreduce.job.priority** parameter and use CLI or API to set the task priority.

- Through the CLI  
When submitting tasks, add the **-Dmapreduce.job.priority=<priority>** parameter.

*<priority>* can be set to any of the following values:

- VERY\_HIGH
  - HIGH
  - NORMAL
  - LOW
  - VERY\_LOW
- Through the API
- You can also set the task priority through the API.
- Set `Configuration.set("mapreduce.job.priority", <priority>)` or `Job.setPriority(JobPriority priority)`.

## 30.19.3 Optimizing Node Configuration

### Scenario

After the scheduler of a big data cluster is properly configured, you can adjust the available memory, CPU resources, and local disk of each node to optimize the performance.

The configuration items are as follows:

- Available memory
- Number of vCPUs
- Physical CPU usage
- Coordination of memory and CPU resources
- Local disk

### Procedure

For details about how to adjust parameter settings, see [Modifying Cluster Service Configuration Parameters](#).

- **Available memory**

Except the memory allocated to the OS and other services, allocate as much as possible memory to Yarn. You can adjust the following parameters to improve resource utilization.

Assume that a container uses 512 MB memory by default, then the memory usage formula is: 512 MB x Number of containers.

By default, the Map or Reduce container uses one vCPU and 1,024 MB memory, and ApplicationMaster uses 1,536 MB memory.

Parameter	Description	Default Value
yarn.nodemanager.resourcememory-mb	Physical memory that can be allocated to containers, in MB. The value must be greater than 0. You are advised to set the parameter value to 75% to 90% of the total physical memory of nodes. If the node has permanent processes of other services, reduce this parameter value to reserve sufficient resources for the processes.	MRS 3.x or later: <b>16384</b> Versions earlier than MRS 3.x: <b>8192</b>

- **Number of vCPUs**

You are advised to set this parameter to 1.5 to 2 times the number of logical CPUs. If the upper layer computing applications have low computing capability requirements, you can set the parameter to two times the number of logical CPUs.

Parameter	Description	Default Value
yarn.nodemanager.resourcememory-cpu-vcores	Number of vCPUs that can be used by Yarn on the node. The default value is <b>8</b> . You are advised to set the value to 1.5 to 2 times the number of logical CPUs.	8

- **Physical CPU usage**

You are advised to reserve appropriate CPUs for the OS and the processes, such as database and HBase, and allocate the remaining CPUs to Yarn. You can set the following parameters to adjust the physical CPU usage.

Parameter	Description	Default Value
yarn.nodemanager.resource.percentage-physical-cpu-limit	<p>Physical CPU percentage that can be used by Yarn on a node. The default value is <b>90</b>, indicating that no CPU control is implemented and Yarn can use all CPU resources. You can only view the parameter. To change the value of this parameter, set the value of RES_CPUSSET_PERCENTAGE of YARN. You are advised to set this parameter to the percentage of CPU resources that can be used by the YARN cluster.</p> <p>For example, If 20% of CPU resources are used by other services (such as HBase, HDFS, and Hive) and system processes on the node, the CPU resources can be scheduled for Yarn is <math>1 - 20\% = 80\%</math>. Therefore, you can set this parameter to <b>80</b>.</p>	90

- **Local disk**

MapReduce writes the intermediate job execution results in local disks. Therefore, configure disks as much as possible and disk space as large as possible. A simple way is to configure the same number of disks as DataNode except for the last directory.

 **NOTE**

Use commas (,) to separate multiple disks.

Parameter	Description	Default Value
yarn.nodemanager.log-dirs	<p>Directories in which logs are stored. Multiple directories can be specified.</p> <p>Storage location of container logs. The default value is % <b>{@auto.detect.datapart.nm.logs}</b>. If there is a data partition, a path list similar to <b>/srv/BigData/hadoop/data1/nm/containerlogs,/srv/BigData/hadoop/data2/nm/containerlogs</b> is generated based on the data partition. If there is no data partition, the default path <b>/srv/BigData/yarn/data1/nm/containerlogs</b> is generated. In addition to using expressions, you can enter a complete list of paths, such as <b>/srv/BigData/yarn/data1/nm/containerlogs</b> or <b>/srv/BigData/yarn/data1/nm/containerlogs,/srv/BigData/yarn/data2/nm/containerlogs</b>. In this way, data is stored in all the configured directories, which are usually on different devices. To ensure disk I/O load balancing, you are advised to provide several paths and each path corresponds to an independent disk. The localized log directory of the application exists in the relative path <b>/application_%{appid}</b>. The log directory of an independent container, that is, <b>container_{\$contid}</b>, is the subdirectory of this directory. Each container directory contains the <b>stderr</b>, <b>stdin</b>, and <b>syslog</b> files generated by the container. To add a directory, for example, <b>/srv/BigData/yarn/data2/nm/containerlogs</b>, you need to delete the files in <b>/srv/BigData/yarn/data2/nm/containerlogs</b> first. Then, assign the same read and write permissions to <b>/srv/BigData/yarn/data2/nm/containerlogs</b> as those of <b>/srv/</b></p>	<p>% {@auto.detect.datapart.nm.logs}</p>

Parameter	Description	Default Value
	<p><b>BigData/yarn/data1/nm/containerlogs</b>, and change <b>/srv/BigData/yarn/data1/nm/containerlogs</b> to <b>/srv/BigData/yarn/data1/nm/containerlogs,/srv/BigData/yarn/data2/nm/containerlogs</b>. You can add directories, but do not modify or delete existing directories. Otherwise, NodeManager data will be lost and services will be unavailable.</p> <p>Default value: % <b>{@auto.detect.datapart.nm.logs}</b> }</p> <p>Exercise caution when modifying this parameter. If the configuration is incorrect, the services are unavailable. If the value of this configuration item at the role level is changed, the value of this configuration item at all instance levels will be changed. If the value of this configuration item at the instance level is changed, the value of this configuration item of other instances remains unchanged.</p>	



Parameter	Description	Default Value
yarn.nodemanager.local-dirs	<p>Storage location of files after localization. The default value is %  <b>{@auto.detect.datapart.nm.localdir}</b>. If there is a data partition, a path list similar to <b>/srv/BigData/hadoop/data1/nm/localdir,/srv/BigData/hadoop/data2/nm/localdir</b> is generated based on the data partition. If there is no data partition, the default path <b>/srv/BigData/yarn/data1/nm/localdir</b> is generated. In addition to using expressions, you can enter a complete list of paths, such as <b>/srv/BigData/yarn/data1/nm/localdir</b> or <b>/srv/BigData/yarn/data1/nm/localdir,/srv/BigData/yarn/data2/nm/localdir</b>. In this way, data is stored in all the configured directories, which are usually on different devices. To ensure disk I/O load balancing, you are advised to provide several paths and each path corresponds to an independent disk. The localized file directory of the application is stored in the relative path <b>/usercache/%{user}/appcache/application_%{appid}</b>. The working directory of an independent container, that is, <b>container_%{contid}</b>, is the subdirectory of the directory. To add a directory, for example, <b>/srv/BigData/yarn/data2/nm/localdir</b>, you need to delete the files in <b>/srv/BigData/yarn/data2/nm/localdir</b> first. Then, assign the same read and write permissions to <b>/srv/BigData/hadoop/data2/nm/localdir</b> as those of <b>/srv/BigData/hadoop/data1/nm/localdir</b>, and change <b>/srv/BigData/yarn/data1/nm/localdir</b> to <b>/srv/BigData/yarn/data1/nm/localdir,/srv/BigData/yarn/data2/nm/localdir</b>. You can add</p>	<p>%  <b>{@auto.detect.datapart.nm.localdir}</b></p>

Parameter	Description	Default Value
	<p>directories, but do not modify or delete existing directories. Otherwise, NodeManager data will be lost and services will be unavailable.</p> <p>Default value: % <b>{@auto.detect.datapart.nm.local dir}</b></p> <p>Exercise caution when modifying this parameter. If the configuration is incorrect, the services are unavailable. If the value of this configuration item at the role level is changed, the value of this configuration item at all instance levels will be changed. If the value of this configuration item at the instance level is changed, the value of this configuration item of other instances remains unchanged.</p>	

## 30.20 Common Issues About Yarn

### 30.20.1 Why Mounted Directory for Container is Not Cleared After the Completion of the Job While Using CGroups?

#### Question

Why mounted directory for Container is not cleared after the completion of the job while using CGroups?

#### Answer

The mounted path for the Container should be cleared even if job is failed.

This happens due to the deletion timeout. Some task takes more time to complete than the deletion time.

To avoid this scenario, you can go to the **All Configurations** page of Yarn by referring to [Modifying Cluster Service Configuration Parameters](#). Search for the **yarn.nodemanager.linux-container-executor.cgroups.delete-timeout-ms** configuration item in the search box to change the deletion interval. The value is in milliseconds.

## 30.20.2 Why the Job Fails with HDFS\_DELEGATION\_TOKEN Expired Exception?

### Question

Why is the HDFS\_DELEGATION\_TOKEN expired exception reported when a job fails in security mode?

### Answer

HDFS\_DELEGATION\_TOKEN expires because the token is not updated or it is accessed after max. lifetime.

Ensure the following parameter value of max. lifetime of the token is greater than the job running time.

**dfs.namenode.delegation.token.max-lifetime=604800000** (1 week by default)

Go to the **All Configurations** page of HDFS by referring to [Modifying Cluster Service Configuration Parameters](#) and search for this parameter in the search box.

#### NOTE

You are advised to set this parameter to a value that is multiple times of the number of hours within the max. lifecycle of the token.

## 30.20.3 Why Are Local Logs Not Deleted After YARN Is Restarted?

### Question

If Yarn is restarted in either of the following scenarios, local logs will not be deleted as scheduled and will be retained permanently:

- When Yarn is restarted during task running, local logs are not deleted.
- When the task is complete and logs fail to be collected, restart Yarn before the logs are cleared as scheduled. In this case, local logs are not deleted.

### Answer

NodeManager has a restart recovery mechanism.

Go to the **All Configurations** tab page of YARN by referring to [Modifying Cluster Service Configuration Parameters](#). Set **yarn.nodemanager.recovery.enabled** of NodeManager to **true** to make the configuration take effect. The default value is **true**. In this way, redundant local logs are periodically deleted when the YARN is restarted.

## 30.20.4 Why the Task Does Not Fail Even Though AppAttempts Restarts for More Than Two Times?

### Question

Why the task does not fail even though AppAttempts restarts due to failure for more than two times?

### Answer

During the task execution process, if the **ContainerExitStatus** returns value **ABORTED**, **PREEMPTED**, **DISKS\_FAILED**, or **KILLED\_BY\_RESOURCEMANAGER**, the system will not count it as a failed attempt. Therefore, the task fails only when the AppAttempts fails actually, that is, the return value is not **ABORTED**, **PREEMPTED**, **DISKS\_FAILED**, or **KILLED\_BY\_RESOURCEMANAGER** for two times.

## 30.20.5 Why Is an Application Moved Back to the Original Queue After ResourceManager Restarts?

### Question

After I moved an application from one queue to another, why is it moved back to the original queue after ResourceManager restarts?

### Answer

This problem is caused by the constraints of the ResourceManager. If a running application is moved to another queue, information about the new queue will not be stored in the ResourceManager after the ResourceManager restarts.

Assume that a user submits a MapReduce application to the leaf queue test11. If the leaf queue test11 is deleted when the application is running, the application will go to the lost\_and\_found queue and the application stops. To start the application, the user moves the application to the leaf queue test21 and the application resumes running. If the ResourceManager restarts, the displayed submission queue is lost\_and\_found, but not test21.

If the application is not complete, the ResourceManager only stores the queue information before the application is moved. As a result, the application is moved back to the original queue. To solve this problem, move the application again after the ResourceManager is restarted to write information about the new queue to the ResourceManager.

## 30.20.6 Why Does Yarn Not Release the Blacklist Even All Nodes Are Added to the Blacklist?

### Question

Why does Yarn not release the blacklist even all nodes are added to the blacklist?

## Answer

In Yarn, when the number of application nodes added to the blacklist by ApplicationMaster (AM) reaches a certain proportion (the default value is 33% of the total number of nodes), the AM automatically releases the blacklist. In this way, all available nodes are added to the blacklist and tasks can obtain node resources.

Assume that there are 8 nodes in a cluster and they are divided into pool A and pool B by NodeLabel. There are two nodes in pool B. A user submits a task App1 to pool B, but there is not sufficient HDFS space and App1 fails to run. As a result, two nodes in pool B are added to the blacklist by the AM of App1. According to the preceding principles, 2 is less than the 33% of 8. Therefore, Yarn does not release the blacklist, and App1 cannot obtain resources and keeps running. Even if the node that is added to the blacklisted is recovered, App1 still cannot obtain resources.

The preceding principles do not apply to the resource pool scenario. Therefore, you can change the value of the client parameter **yarn.resourcemanager.am-scheduling.node-blacklisting-disable-threshold** to **(nodes number of pool / total nodes) \* 33%** to solve this problem.

## 30.20.7 Why Does the Switchover of ResourceManager Occur Continuously?

### Question

The switchover of ResourceManager occurs continuously when multiple, for example 2,000, tasks are running concurrently, causing the Yarn service unavailable.

### Answer

The cause is that the time of full GarbageCollection exceeds the interaction duration threshold between the ResourceManager and ZooKeeper duration threshold. As a result, the connection between the ResourceManager and ZooKeeper fails and the switchover of ResourceManager occurs continuously.

When there are multiple tasks, ResourceManager saves the authentication information about multiple tasks and transfers the information to NodeManagers through heartbeat, which is called heartbeat response. The lifecycle of heartbeat response is short. The default value is 1s. Normally, heartbeat response can be reclaimed during the JVM minor GarbageCollection. However, if there are multiple tasks and there are a lot of nodes, for example 5000 nodes, in the cluster, the heartbeat response of multiple nodes occupy a large amount of memory. As a result, the JVM cannot completely reclaim the heartbeat response during minor GarbageCollection. The heartbeat response failed to be reclaimed accumulate and the JVM full GarbageCollection is triggered. The JVM GarbageCollection is in a blocking mode, in other words, no jobs are performed during the GarbageCollection. Therefore, if the duration of full GarbageCollection exceeds the periodical interaction duration threshold between the ResourceManager and ZooKeeper, the switchover occurs.

Log in to FusionInsight Manager, choose **Cluster > Services > Yarn**, and click the **Configurations** tab and then **All Configurations**. In the navigation pane on the

left, choose **Yarn > Customization**, and add the **yarn.resourcemanager.zk-timeout-ms** parameter to the **yarn.yarn-site.customized.configs** file to increase the threshold of the periodic interaction duration between ResourceManager and ZooKeeper (the value range is less than or equal to 90,000 ms). In this way, the problem of continuous active/standby ResourceManager switchover can be solved.

## 30.20.8 Why Does a New Application Fail If a NodeManager Has Been in Unhealthy Status for 10 Minutes?

### Question

Why does a new application fail if a NodeManager has been in unhealthy status for 10 minutes?

### Answer

When **nodeSelectPolicy** is set to **SEQUENCE** and the first NodeManager connected to the ResourceManager is unavailable, the ResourceManager attempts to assign tasks to the same NodeManager in the period specified by **yarn.nm.liveness-monitor.expiry-interval-ms**.

You can use either of the following methods to avoid the preceding problem:

- Use another nodeSelectPolicy, for example, **RANDOM**.
- Go to the **All Configurations** page of Yarn by referring to [Modifying Cluster Service Configuration Parameters](#). Search for the following parameters in the search box and modify the following attributes in the **yarn-site.xml** file:  
**yarn.resourcemanager.am-scheduling.node-blacklisting-enabled = true;**  
**yarn.resourcemanager.am-scheduling.node-blacklisting-disable-threshold = 0.5.**

## 30.20.9 Why Does an Error Occur When I Query the ApplicationID of a Completed or Non-existing Application Using the RESTful APIs?

### Question

Why does an error occur when I query the applicationID of a completed or non-existing application using the RESTful APIs?

### Answer

The Superior scheduler only stores the applicationIDs of running applications. If you view the applicationID of a completed or non-existing application by accessing the RESTful API at **https://<SS\_REST\_SERVER>/ws/v1/sscheduler/applications/{application\_id}**, the 404 error is returned by the server. If Chrome web browser is used, the **Error Occurred** message is displayed because Chrome preferentially responds in the application/xml format. If Internet Explorer is used, the **404** error code is displayed because IE web browser preferentially responds in the application/json format.

## 30.20.10 Why May A Single NodeManager Fault Cause MapReduce Task Failures in the Superior Scheduling Mode?

### Question

In Superior scheduling mode, if a single NodeManager is faulty, why may the MapReduce tasks fail?

### Answer

In normal cases, when the attempt of a single task of an application fails on a node for three consecutive times, the AppMaster of the application adds the node to the blacklist. Then, the AppMaster instructs the scheduler not to schedule the task to the node to avoid task failure.

However, by default, if 33% nodes in the cluster are added to the blacklist, the scheduler ignores the blacklisted nodes. Therefore, the blacklist feature is prone to become invalid in small cluster scenarios. For example, there are only three nodes in the cluster. If one node is faulty, the blacklist mechanism becomes invalid. The scheduler continues to schedule the task to the node no matter how many times the attempt of the task fails on the node. As a result, the number of attempts of the task reaches the maximum (4 times by default for MapReduce). And the MapReduce tasks failed.

Workaround:

The `yarn.resourcemanager.am-scheduling.node-blacklisting-disable-threshold` parameter indicates the threshold for ignoring blacklisted nodes, in percentage. You are advised to increase the value of this parameter based on the cluster scale. For example, you are advised to set this parameter to **50%** for a three-node cluster.

#### NOTE

The framework design of the Superior scheduler is time-based asynchronous scheduling. When the NodeManager is faulty, ResourceManager cannot quickly detect that the NodeManager is faulty (10 minutes by default). Therefore, the Superior scheduler still schedules tasks to the node, causing task failures.

## 30.20.11 Why Are Applications Suspended After They Are Moved From Lost\_and\_Found Queue to Another Queue?

### Question

When a queue is deleted when there are applications running in it, these applications are moved to the "lost\_and\_found" queue. When these applications are moved back to another healthy queue, some tasks are suspended.

### Answer

If no label expression is set for the current application, the default label expression of the queue is used as label expression for new container/resource demands requested by the application. If there is no default label expression of the queue,

then **default label** is considered as the label expression for new container/resource demands requested by the application.

When application app1 is submitted to the queue Q1, **label1**, the default label expression of the queue, is used for the application's new resource requests/containers. If Q1 is deleted when app1 is running, app1 is moved to the "lost\_and\_found" queue. Because there is no label expression of the "lost\_and\_found" queue, **default label** is used as the label expression of app1's new resource requests/containers. Assume that app1 is moved to another normal queue Q2. If Q2 supports **label1** and **default label**, app1 can run properly. If Q2 does not support **label1** or **default label**, the resource request with **label1** or **default label** cannot obtain resources, causing task suspension.

To solve this problem, ensure that the queue to which the application is moved from "lost\_and\_found" queue supports label expression of the moved application.

You are not advised to delete a queue in which there are running applications.

## 30.20.12 How Do I Limit the Size of Application Diagnostic Messages Stored in the ZKstore?

### Question

How do I limit the size of application diagnostic messages stored in the ZKstore?

### Answer

In some cases, it has been observed that diagnostic messages may grow infinitely. Because diagnostic messages are stored in the ZKstore, it is not recommended that you allow diagnostic messages to grow indefinitely. Therefore, a property parameter is needed to set the maximum size of the diagnostic message.

If you need to set **yarn.app.attempt.diagnostics.limit.kc**, go to the **All Configurations** page by referring to [Modifying Cluster Service Configuration Parameters](#) and search for the following parameters in the search box:

**Table 30-27** Parameter description

Parameter	Description	Default Value
yarn.app.attempt.diagnostics.limit.kc	Data size of the diagnosis message for each application connection, in kilobytes (number of characters x 1,024). When ZooKeeper is used to store the behavior status of applications, the size of diagnosis messages needs to be limited to prevent Yarn from overloading ZooKeeper. If <b>yarn.resourcemanager.state-store.max-completed-applications</b> is set to a large value, you need to decrease the value of this property to limit the total size of stored data.	64



### 30.20.13 Why Does a MapReduce Job Fail to Run When a Non-ViewFS File System Is Configured as ViewFS?

#### Question

Why does a MapReduce job fail to run when a non-ViewFS file system is configured as ViewFS?

#### Answer

When a non-ViewFS file system is configured as a ViewFS using cluster, the user permissions on folders in the ViewFS file system are different from those of non-ViewFS folders in the default NameService. The submitted MapReduce job fails to be executed because the directory permissions are inconsistent.

When configuring the ViewFS user in the cluster, you need to check and verify the directory permissions. Before submitting a job, change the ViewFS folder permissions based on the default NameService folder permissions.

The following table lists the default permission structure of directories configured in ViewFS. If the configured directory permissions are not included in the following table, you must change the directory permissions accordingly.

**Table 30-28** Default permission structure of directories configured in ViewFS

Parameter	Description	Default Value	Default value and default permissions on the parent directory
yarn.nodemanager.remote-app-log-dir	On the default file system (usually HDFS), specify the directory to which the NM aggregates logs.	logs	777
yarn.nodemanager.remote-app-log-archive-dir	Directory for archiving logs	-	777
yarn.app.mapreduce.am.staging-dir	Staging directory used when a job is submitted	/tmp/hadoop-yarn/staging	777
mapreduce.jobhistory.intermediate-done-dir	Directory for storing historical files of MapReduce jobs	\${yarn.app.mapreduce.am.staging-dir}/history/done_intermediate	777

Parameter	Description	Default Value	Default value and default permissions on the parent directory
mapreduce.jobhistory.done-dir	Directory of historical files managed by the MR JobHistory Server.	\${yarn.app.mapreduce.am.staging-dir}/history/done	777

### 30.20.14 Why Do Reduce Tasks Fail to Run in Some OSs After the Native Task Feature is Enabled?

#### Question

After the Native Task feature is enabled, Reduce tasks fail to run in some OSs.

#### Answer

When - **Dmapreduce.job.map.output.collector.class=org.apache.hadoop.mapred.native.task.NativeMapOutputCollectorDelegator** is executed to enable the Native Task feature during the running of MapReduce tasks that contain Reduce tasks, the tasks fail to run in some OSs, and the error message "version 'GLIBCXX\_3.4.20' not found" is displayed in logs. The cause is that the GLIBCXX version of the OSs is too early. As a result, the libnativetask.so.1.0.0 library on which the feature depends cannot be loaded, leading to task failures.

Workaround:

Set **mapreduce.job.map.output.collector.class** to **org.apache.hadoop.mapred.MapTask\$MapOutputBuffer**.

# 31 Using ZooKeeper

---

## 31.1 Using ZooKeeper from Scratch

ZooKeeper is an open-source, highly reliable, and distributed consistency coordination service. ZooKeeper is designed to solve the problem that data consistency cannot be ensured for complex and error-prone distributed systems. There is no need to develop dedicated collaborative applications, which is suitable for high availability services to ensure data consistency.

### Background Information

Before using the client, you need to download and update the client configuration file on all clients except the client of the active management node.

### Procedure

For MRS 2.x or earlier, perform the following operations:

**Step 1** Download the client configuration file.

1. Log in to the MRS console. In the left navigation pane, choose **Clusters > Active Clusters**, and click the cluster to be operated.
2. Click the **Components** tab.

**Step 2** Log in to the active management node of MRS Manager.

1. On the MRS console, choose **Clusters > Active Clusters** and click a cluster name. On the **Nodes** tab, view the node names. The node whose name contains **master1** is the Master1 node, and the node whose name contains **master2** is the Master2 node.

The active and standby management nodes of MRS Manager are installed on Master nodes by default. Because Master1 and Master2 are switched over in active and standby mode, Master1 is not always the active management node of MRS Manager. Run a command in Master1 to check whether Master1 is active management node of MRS Manager. For details about the command, see [Step 2.4](#).

2. Log in to the Master1 node using the password as user **root**.

- Run the following commands to switch to user **omm**:

```
sudo su - root
su - omm
```

- Run the following command to check the active management node of MRS Manager:

```
sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh
```

In the command output, the node whose **HAActive** is **active** is the active management node, and the node whose **HAActive** is **standby** is the standby management node. In the following example, **mgtomsdat-sh-3-01-1** is the active management node, and **mgtomsdat-sh-3-01-2** is the standby management node.

Ha mode	double	NodeName	HostName	HAVersion	StartTime	HAActive
HAAllResOK		192-168-0-30	mgtomsdat-sh-3-01-1	V100R001C01	2014-11-18 23:43:02	<b>active</b>
	normal		Activated			
		192-168-0-24	mgtomsdat-sh-3-01-2	V100R001C01	2014-11-21 07:14:02	<b>standby</b>
	normal		Deactivated			

- Log in to the active management node, for example, **192-168-0-30** of MRS Manager as user **root**, and run the following command to switch to user **omm**:

```
sudo su - omm
```

- Step 3** Run the following command to switch to the client installation directory, for example, **/opt/client**:

```
cd /opt/client
```

- Step 4** Run the following command to update the client configuration for the active management node.

```
sh refreshConfig.sh /opt/client Full path of the client configuration file package
```

For example, run the following command:

```
sh refreshConfig.sh /opt/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.
```

 **NOTE**

You can also refer to method 2 in [Updating a Client \(Versions Earlier Than 3.x\)](#) to perform operations in [Step 1](#) to [Step 4](#).

- Step 5** Use the client on a Master node.

- On the active management node where the client is updated, for example, node **192-168-0-30**, run the following command to go to the client directory:

```
cd /opt/client
```

- Run the following command to configure environment variables:

```
source bigdata_env
```

- If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step:

**kinit** *MRS cluster user*

Example: **kinit zookeeperuser**.

- Run the following Zookeeper client command:

**zkCli.sh -server** *<zookeeper installation node IP>:<port>*

Example: **zkCli.sh -server node-master1DGhZ:2181**

**Step 6** Run the ZooKeeper client command.

- Create a ZNode.

```
create /test
```

- View ZNode information.

```
ls /
```

- Write data to the ZNode.

```
set /test "zookeeper test"
```

- View the data written to the ZNode.

```
get /test
```

- Delete the created ZNode.

```
delete /test
```

----End

For MRS 3.x or later, perform the following operations:

**Step 1** Download the client configuration file.

- Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#).
- Choose **Cluster** > *Name of the desired cluster* > **Dashboard** > **More** > **Download Client**.
- Download the cluster client.

Set **Select Client Type** to **Configuration Files Only**, select a platform type, and click **OK** to generate the client configuration file which is then saved in the **/tmp/FusionInsight-Client/** directory on the active management node by default.

**Step 2** Log in to the active management node of Manager.

- Log in to any node where Manager is deployed as user **root**.
- Run the following command to identify the active and standby nodes:

**sh** **`\${BIGDATA\_HOME}/om-server/om/sbin/status-oms.sh**

In the command output, the value of **HAActive** for the active management node is **active**, and that for the standby management node is **standby**. In the following example, **node-master1** is the active management node, and **node-master2** is the standby management node.

```
HAMode
double
NodeName HostName HAVersion StartTime HAActive
HAAllResOK HARunPhase
192-168-0-30 node-master1 V100R001C01 2020-05-01 23:43:02 active
normal
192-168-0-24 node-master2 V100R001C01 2020-05-01 07:14:02 standby
normal
Deactivated
```

3. Log in to the primary management node as user **root** and run the following command to switch to user **omm**:

```
sudo su - omm
```

- Step 3** Run the following command to switch to the client installation directory, for example, **/opt/client**:

```
cd /opt/client
```

- Step 4** Run the following command to update the client configuration for the active management node.

```
sh refreshConfig.sh /opt/client Full path of the client configuration file package
```

For example, run the following command:

```
sh refreshConfig.sh /opt/client /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_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.
```

- Step 5** Use the client on a Master node.

1. On the active management node where the client is updated, for example, node **192-168-0-30**, run the following command to go to the client directory:

```
cd /opt/client
```

2. Run the following command to configure environment variables:

```
source bigdata_env
```

3. If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step:

```
kinit MRS cluster user
```

Example: **kinit zookeeperuser**.

4. Run the following Zookeeper client command:

```
zkCli.sh -server <zookeeper installation node IP>:<port>
```

Example: **zkCli.sh -server node-master1DGhZ:2181**

- Step 6** Run the ZooKeeper client command.

1. Create a ZNode.

```
create /test
```

2. View ZNode information.

```
ls /
```

3. Write data to the ZNode.

```
set /test "zookeeper test"
```

4. View the data written to the ZNode.

```
get /test
```

5. Delete the created ZNode.

```
delete /test
```

----End

## 31.2 Common ZooKeeper Parameters

Navigation path for setting parameters:

Go to the **All Configurations** page of ZooKeeper by referring to [Modifying Cluster Service Configuration Parameters](#). Enter a parameter name in the search box.

**Table 31-1** Parameters

Parameter	Description	Default Value
skipACL	Specifies whether to skip the permission check of the ZooKeeper node.	no
maxClientCnxns	Specifies the maximum number of connections of ZooKeeper. It is recommended this parameter is set to a larger value in scenarios with a large number of connections.	2000
LOG_LEVEL	Specifies the log level. This parameter can be set to <b>DEBUG</b> during commissioning.	INFO
acl.compare.shortName	Specifies whether to perform ACL authentication only by principal username when the Znode ACL authentication type is SASL.	true
synclimit	Specifies the interval of synchronization between the follower and leader (unit: tick). If the leader does not respond within the specified time range, the connection cannot be established.	15
tickTime	Specifies the duration of a tick (in milliseconds). It is the basic time unit used by ZooKeeper, which defines heartbeat and timeout durations.	4000

 NOTE

The ZooKeeper internal time is determined by **ticktime** and **synclimit**. To increase the ZooKeeper internal timeout interval, increase the timeout interval for the client to connect to ZooKeeper.

## 31.3 Using a ZooKeeper Client

### Scenario

Use a ZooKeeper client in an O&M scenario or service scenario.

### Prerequisites

You have installed the client. For example, the installation directory is **/opt/client**. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite.

### Procedure

**Step 1** Log in to the node where the client is installed as the client installation user.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Run the following command to authenticate the user: (skip this step in common mode):

```
kinit Component service user
```

**Step 5** 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
```

```
----End
```

## 31.4 Configuring the ZooKeeper Permissions

### Scenario

Configure znode permission of ZooKeeper.

ZooKeeper uses an access control list (ACL) to implement znode access control. The ZooKeeper client specifies a znode ACL, and the ZooKeeper server determines whether a client that requests for a znode has related operation permission according to the ACL. ACL configuration involves the following four operations:



- Check znode ACLs in ZooKeeper.
- Add znode ACLs to ZooKeeper.
- Modify znode ACLs in ZooKeeper.
- Delete znode ACLs from ZooKeeper.

The ZooKeeper ACL permission is described as follows:

ZooKeeper supports five types of permission, create, delete, read, write, and admin. ZooKeeper permission control is of a znode level. That is, the permission configuration for a parent znode is not inherited by its child znodes. The ZooKeeper znode default permission is **world:anyone: cdrwa**. That is, any user has all permissions.

 **NOTE**

ACL has three parts:

The first part is the authentication type. For example, **world** indicates all authentication types and **sasl** indicates the kerberos authentication type.

The second part is the account. For example, anyone indicates any user.

The third part is permission. For example, **cdrwa** indicates all permissions.

In particular, because starting the client in common mode does not need authentication, ACL with **sasl** authentication type cannot be used in common mode. Authentications of **sasl** scheme in this document are performed in clusters that have the security mode enabled.

**Table 31-2** Five types of ZooKeeper ACLs

Permission Description	Permission Name	Permission Details
Create permission	create(c)	Users with this permission can create child znodes in the current znode.
Delete permission	delete(d)	Users with this permission can delete the current znode.
Read permission	read(r)	Users with this permission can obtain data of the current znode and list all the child znodes of the current znode.
Write permission	write(w)	Users with this permission can write data to the current znode and its child znodes.
Administration permission	admin(a)	Users with this permission can set permission for the current znode.

## Impact on the System

### NOTICE

Modifying ZooKeeper ACLs is a critical operation. If znode permission is modified in ZooKeeper, other users may have no permission to access the znode and some system functions are abnormal. In 3.5.6 and later versions, users must have the read permission for the **getAcl** operation.

## Prerequisites

- The ZooKeeper client has been installed in a directory, for example, **/opt/client**.
- You have obtained the username and password of an MRS cluster administrator.

## Procedure

**Start the ZooKeeper client.**

**Step 1** Log in to the server where the ZooKeeper client is installed as user **root**.

**Step 2** Run the following command to go to the client installation directory:

```
cd /opt/client
```

**Step 3** Run the following command to configure environment variables:

```
source bigdata_env
```

**Step 4** Run the following command and enter the user password to authenticate the user's identity (This step is required only for clusters in security mode, and user **userA** is provided as an example of an authorized user.):

```
kinit userA
```

**Step 5** On the ZooKeeper client, run the following command to go to the ZooKeeper command-line interface (CLI):

```
sh zkCli.sh -server ZooKeeper plane IP address of any instance:clientPort
```

The default **clientPort** is **2181**.

Example: **sh zkCli.sh -server 192.168.0.151:2181**

**Step 6** Run the **ls** command to view the znode list in ZooKeeper. For example, you can view the list of znodes in the root directory.

```
ls /
```

```
[zk: 192.168.0.151:2181(CONNECTED) 1] ls /
[hadoop-flag, hadoop-ha, test, test2, test3, test4, test5, test6, zookeeper]
```

**View the ZooKeeper znode ACL.**

**Step 7** Start the ZooKeeper client.

**Step 8** Run the **getAcl** command to view znodes. The following command can be used to view the created znode ACL named **test**:

**getAcl** /znode name

```
[zk: 192.168.0.151:2181(CONNECTED) 2] getAcl /test
'world,'anyone
: cdrwa
```

Add a ZooKeeper znode ACL.

**Step 9** Start the ZooKeeper client.

**Step 10** View the old ACL information to check whether the current account has the permission to modify the znode ACL information (a permission). If no, use kinit to switch to a user that has the permission and restart the ZooKeeper client.

**getAcl** /znode name

```
[zk: 192.168.0.151:2181(CONNECTED) 3] getAcl /test
'world,'anyone
: cdrwa
```

**Step 11** Run the **setAcl** command to add an ACL. The command for adding an ACL is as follows:

**setAcl** /test world:anyone:cdrwa,sasl: username@: <system domain name>:ACL value

For example, to create the ACL of user **admin** to the test znode, run the following command:

**setAcl** /test world:anyone:cdrwa,sasl:userA@HADOOP.COM:cdrwa

#### NOTE

When adding a new ACL, reserve the existing ones. The new and old ACLs are separated by a comma. The newly added ACL has three parts:

- The first part is the authentication type. For example, **sasl** indicates kerberos authentication.
- The second part is the account. For example, **userA@HADOOP.COM** indicates user **userA**.
- The third part is permission. For example, **cdrwa** indicates all permissions.

**Step 12** After adding the ACL, run the **getAcl** command to check whether the permission is added successfully:

**getAcl** /znode name

```
[zk: 192.168.0.151:2181(CONNECTED) 4] getAcl /test
'world,'anyone
: cdrwa
'sasl,'userA@<System domain name>
: cdrwa
```

**Modify the ZooKeeper znode ACL.**

**Step 13** Start the ZooKeeper client.

**Step 14** View the old ACL information to check whether the current account has the permission to modify the znode ACL information (a permission). If no, use kinit to switch to a user that has the permission and restart the ZooKeeper client.

**getAcl** /znode name

```
[zk: 192.168.0.151:2181(CONNECTED) 5] getAcl /test
'world,'anyone
: cdrwa
'sasl,'userA@<System domain name>
: cdrwa
```

- Step 15** Run the **setAcl** command to modify an ACL. The command for adding an ACL is as follows:

```
setAcl /test sasl:Username@<System domain name>:ACL value
```

For example, to reserve all permissions of user **userA** and delete the rw permission of user **anyone**, run the following command:

```
setAcl /test sasl:userA@HADOOP.COM:cdrwa
```

- Step 16** After modifying the ACL, run the **getAcl** command to check whether the permission is modified successfully:

```
getAcl /znode name
```

```
[zk: 192.168.0.151:2181(CONNECTED) 6] getAcl /test
'sasl,'userA@<System domain name>
: cdrwa
```

#### Delete the ZooKeeper znode ACL.

- Step 17** Start the ZooKeeper client.

- Step 18** View the old ACL information to check whether the current account has the permission to modify the znode ACL information (a permission). If no, use kinit to switch to a user that has the permission and restart the ZooKeeper client.

```
getAcl /znode name
```

```
[zk: 192.168.0.151:2181(CONNECTED) 5] getAcl /test
'world,'anyone
: rw
'sasl,'userA@<System domain name>
: cdrwa
```

- Step 19** Run the **setAcl** command to add an ACL. The command for adding an ACL is as follows:

```
setAcl /test sasl:Username@<System domain name>:ACL value
```

For example, to reserve all permissions of user **userA** and delete the rw permission of user **anyone**, run the following command:

```
setAcl /test sasl:userA@HADOOP.COM:cdrwa
```

- Step 20** After modifying the ACL, run the **getAcl** command to check whether the permission is modified successfully:

```
getAcl /znode name
```

```
[zk: 192.168.0.151:2181(CONNECTED) 6] getAcl /test
'sasl,'userA@<System domain name>
: cdrwa
```

----End

## 31.5 ZooKeeper Log Overview

### Log Description

**Log path:** `/var/log/Bigdata/zookeeper/quorumpeer` (Run log), `/var/log/Bigdata/audit/zookeeper/quorumpeer` (Audit log)

**Log archive rule:** The automatic ZooKeeper log compression function is enabled. By default, when the size of logs exceeds 30 MB, logs are automatically compressed into a log file. A maximum of 20 compressed files can be reserved. The number of compressed files can be configured on Manager.

**Table 31-3** ZooKeeper log list

Log Type	Log File Name	Description
Run logs	zookeeper-<SSH_USER>-<process_name>-<hostname>.log	ZooKeeper system log file, which records most of the logs generated when the ZooKeeper system is running.
	check-serviceDetail.log	Log that records whether the ZooKeeper service starts successfully.
	zookeeper-<SSH_USER>-<DATA>-<PID>-gc.log	ZooKeeper garbage collection log file
	instanceHealthDetail.log	Log that records the health check details of ZooKeeper instance
	zookeeper-omm-server-<hostname>.out	Log indicating that ZooKeeper unexpectedly quits
	zk-err-<zkpid>.log	ZooKeeper fatal error log
	java_pid<zkpid>.hprof	ZooKeeper memory overflow log
	funcDetail.log	ZooKeeper instance startup log
	zookeeper-period-check.log	Health check log of the ZooKeeper instance
	zookeeper-period-check-java.log	ZooKeeper quota monitoring period check log
Audit Log	zk-audit-quorumpeer.log	ZooKeeper operation audit log

## Log levels

**Table 31-4** describes the log levels supported by ZooKeeper. 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 specified level are printed. The number of printed logs decreases as the specified log level increases.

**Table 31-4** Log levels

Level	Description
FATAL	Logs of this level record fatal error information about the current event processing that may result in a system crash.
ERROR	Error information about the current event processing, which indicates that system running is abnormal.
WARN	Abnormal information about the current event processing. These abnormalities will not result in system faults.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

To modify log levels, perform the following operations:

- Step 1** Go to the **All Configurations** page of the ZooKeeper service by referring to [Modifying Cluster Service Configuration Parameters](#).
- Step 2** On the menu bar on the left, select the log menu of the target role.
- Step 3** Select a desired log level.
- Step 4** Click **Save**. In the displayed dialog box, click **OK** to make the configuration take effect.

 **NOTE**

The configurations take effect immediately without the need to restart the service.

----End

## Log Format

The following table lists the ZooKeeper log formats.

**Table 31-5** Log Format

Log Type	Component	Format	Example
Run logs	zookeeper quorumpeer	<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>	2020-01-20 16:33:43,816   INFO   main   Defaulting to majority quorums   org.apache.zookee per.server.quorum. QuorumPeerConfi g.parseProperties( QuorumPeerConfi g.java:335)
Audit logs	zookeeper quorumpeer	<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>	2020-01-20 16:33:54,313   INFO   CommitProcessor: 13   session=0xd4b067 9daea0000 ip=10.177.112.145 operation=create znode target=ZooKeeper Server znode=/zk- write-test-2 result=success   org.apache.zookee per.ZKAuditLogger \$LogLevel \$5.printLog(ZKAu ditLogger.java:70)

## 31.6 Common Issues About ZooKeeper

### 31.6.1 Why Do ZooKeeper Servers Fail to Start After Many znodes Are Created?

#### Question

After a large number of znodes are created, ZooKeeper servers in the ZooKeeper cluster become faulty and cannot be automatically recovered or restarted.

Logs of followers:

```
2016-06-23 08:00:18,763 | WARN | QuorumPeer[myid=26](plain=/10.16.9.138:2181)(secure=disabled) |
Exception when following the leader |
```

```
org.apache.zookeeper.server.quorum.Follower.followLeader(Follower.java:93)
java.net.SocketTimeoutException: Read timed out
 at java.net.SocketInputStream.socketRead0(Native Method)
 at java.net.SocketInputStream.socketRead(SocketInputStream.java:116)
 at java.net.SocketInputStream.read(SocketInputStream.java:170)
 at java.net.SocketInputStream.read(SocketInputStream.java:141)
 at java.io.BufferedInputStream.fill(BufferedInputStream.java:246)
 at java.io.BufferedInputStream.read(BufferedInputStream.java:265)
 at java.io.DataInputStream.readInt(DataInputStream.java:387)
 at org.apache.jute.BinaryInputArchive.readInt(BinaryInputArchive.java:63)
 at org.apache.zookeeper.server.quorum.QuorumPacket.deserialize(QuorumPacket.java:83)
 at org.apache.jute.BinaryInputArchive.readRecord(BinaryInputArchive.java:99)
 at org.apache.zookeeper.server.quorum.Learner.readPacket(Learner.java:156)
 at org.apache.zookeeper.server.quorum.Learner.registerWithLeader(Learner.java:276)
 at org.apache.zookeeper.server.quorum.Follower.followLeader(Follower.java:75)
 at org.apache.zookeeper.server.quorum.QuorumPeer.run(QuorumPeer.java:1094)
2016-06-23 08:00:18,764 | INFO | QuorumPeer[myid=26](plain=/10.16.9.138:2181)(secure=disabled) |
shutdown called | org.apache.zookeeper.server.quorum.Follower.shutdown(Follower.java:198)
java.lang.Exception: shutdown Follower
 at org.apache.zookeeper.server.quorum.Follower.shutdown(Follower.java:198)
 at org.apache.zookeeper.server.quorum.QuorumPeer.stopFollower(QuorumPeer.java:1141)
 at org.apache.zookeeper.server.quorum.QuorumPeer.run(QuorumPeer.java:1098)
```

### Logs of the leader:

```
2016-06-23 07:30:57,481 | WARN | QuorumPeer[myid=25](plain=/10.16.9.136:2181)(secure=disabled) |
Unexpected exception | org.apache.zookeeper.server.quorum.QuorumPeer.run(QuorumPeer.java:1108)
java.lang.InterruptedExcepion: Timeout while waiting for epoch to be acked by quorum
 at org.apache.zookeeper.server.quorum.Leader.waitForEpochAck(Leader.java:1221)
 at org.apache.zookeeper.server.quorum.Leader.lead(Leader.java:487)
 at org.apache.zookeeper.server.quorum.QuorumPeer.run(QuorumPeer.java:1105)
2016-06-23 07:30:57,482 | INFO | QuorumPeer[myid=25](plain=/10.16.9.136:2181)(secure=disabled) |
Shutdown called | org.apache.zookeeper.server.quorum.Leader.shutdown(Leader.java:623)
java.lang.Exception: shutdown Leader! reason: Forcing shutdown
 at org.apache.zookeeper.server.quorum.Leader.shutdown(Leader.java:623)
 at org.apache.zookeeper.server.quorum.QuorumPeer.stopLeader(QuorumPeer.java:1149)
 at org.apache.zookeeper.server.quorum.QuorumPeer.run(QuorumPeer.java:1110)
```

## Answer

After a large number of znodes are created, a large volume of data needs to be synchronized between the follower and leader. If the data synchronization is not complete within the specified time, all ZooKeeper servers fail to start.

Go to the **All Configurations** page of the ZooKeeper service by referring to [Modifying Cluster Service Configuration Parameters](#). To recover ZooKeeper servers, increase the values of **syncLimit** and **initLimit** in the ZooKeeper configuration file **zoo.cfg** until ZooKeeper servers are successfully started.

**Table 31-6** Parameters

Parameter	Description	Default Value
syncLimit	Interval (unit: tick) at which data is synchronized between the follower and the leader. If the leader does not respond to the follower within the specified time, the connection between the leader and follower cannot be set up.	15



Parameter	Description	Default Value
initLimit	Interval (unit: tick) within which the connection and synchronization between the follower and leader must be completed.	15

If ZooKeeper servers do not recover even after **initLimit** and **syncLimit** are set to **300** ticks, check that no other application is killing the ZooKeeper. For example, if the parameter value is **300** and the ticket duration is 2000 ms, the maximum synchronization duration is 600s (300 x 2000 ms).

There may exist the situation where an overwhelming amount of data is created in ZooKeeper and it takes long to synchronize data between the follower and the leader and to save data to the hard disk. This means that ZooKeeper needs to run for a long time. Ensure that no other monitoring application kills the ZooKeeper while ZooKeeper is running.

## 31.6.2 Why Does the ZooKeeper Server Display the java.io.IOException: Len Error Log?

### Question

After a large number of znodes are created in a parent directory, the ZooKeeper client will fail to fetch all child nodes of this parent directory in a single request.

Logs of client:

```
2017-07-11 13:17:19,610 [myid:] - WARN [New I/O worker #3:ClientCnxnSocketNetty
$ZKClientHandler@468] - Exception caught: [id: 0xb66cbb85, /10.18.97.97:49192 ->
10.18.97.97/10.18.97.97:2181] EXCEPTION: java.nio.channels.ClosedChannelException
java.nio.channels.ClosedChannelException
at org.jboss.netty.handler.ssl.SslHandler$6.run(SslHandler.java:1580)
at org.jboss.netty.channel.socket.ChannelRunnableWrapper.run(ChannelRunnableWrapper.java:40)
at org.jboss.netty.channel.socket.nio.AbstractNioWorker.executeInIoThread(AbstractNioWorker.java:71)
at org.jboss.netty.channel.socket.nio.NioWorker.executeInIoThread(NioWorker.java:36)
at org.jboss.netty.channel.socket.nio.AbstractNioWorker.executeInIoThread(AbstractNioWorker.java:57)
at org.jboss.netty.channel.socket.nio.NioWorker.executeInIoThread(NioWorker.java:36)
at org.jboss.netty.channel.socket.nio.AbstractNioChannelSink.execute(AbstractNioChannelSink.java:34)
at org.jboss.netty.handler.ssl.SslHandler.channelClosed(SslHandler.java:1566)
at org.jboss.netty.channel.Channels.fireChannelClosed(Channels.java:468)
at org.jboss.netty.channel.socket.nio.AbstractNioWorker.close(AbstractNioWorker.java:376)
at org.jboss.netty.channel.socket.nio.NioWorker.read(NioWorker.java:93)
at org.jboss.netty.channel.socket.nio.AbstractNioWorker.process(AbstractNioWorker.java:109)
at org.jboss.netty.channel.socket.nio.AbstractNioSelector.run(AbstractNioSelector.java:312)
at org.jboss.netty.channel.socket.nio.AbstractNioWorker.run(AbstractNioWorker.java:90)
at org.jboss.netty.channel.socket.nio.NioWorker.run(NioWorker.java:178)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
at java.lang.Thread.run(Thread.java:745)
```

Logs of leader:

```
2017-07-11 13:17:33,043 [myid:1] - WARN [New I/O worker #7:NettyServerCnxn@445] - Closing
connection to /10.18.101.110:39856
java.io.IOException: Len error 45
at org.apache.zookeeper.server.NettyServerCnxn.receiveMessage(NettyServerCnxn.java:438)
at org.apache.zookeeper.server.NettyServerCnxnFactory
$CnxnChannelHandler.processMessage(NettyServerCnxnFactory.java:267)
```

```

at org.apache.zookeeper.server.NettyServerCnxnFactory
$CnxnChannelHandler.messageReceived(NettyServerCnxnFactory.java:187)
at org.jboss.netty.channel.SimpleChannelHandler.handleUpstream(SimpleChannelHandler.java:88)
at org.jboss.netty.channel.DefaultChannelPipeline.sendUpstream(DefaultChannelPipeline.java:564)
at org.jboss.netty.channel.DefaultChannelPipeline.sendUpstream(DefaultChannelPipeline.java:559)
at org.jboss.netty.channel.Channels.fireMessageReceived(Channels.java:268)
at org.jboss.netty.channel.Channels.fireMessageReceived(Channels.java:255)
at org.jboss.netty.channel.socket.nio.NioWorker.read(NioWorker.java:88)
at org.jboss.netty.channel.socket.nio.AbstractNioWorker.process(AbstractNioWorker.java:109)
at org.jboss.netty.channel.socket.nio.AbstractNioSelector.run(AbstractNioSelector.java:312)
at org.jboss.netty.channel.socket.nio.AbstractNioWorker.run(AbstractNioWorker.java:90)
at org.jboss.netty.channel.socket.nio.NioWorker.run(NioWorker.java:178)
at org.jboss.netty.util.ThreadRenamingRunnable.run(ThreadRenamingRunnable.java:108)
at org.jboss.netty.util.internal.DeadLockProofWorker$1.run(DeadLockProofWorker.java:42)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)
at java.lang.Thread.run(Thread.java:745)

```

## Answer

After a large number of znodes are created in a single parent directory and the client tries to fetch all the child znodes in a single request, the server will fail to return because the results exceed the data size that can be stored in a znode.

To avoid this problem, set **jute.maxbuffer** to a larger value based on the client application.

**jute.maxbuffer** can only be set to a Java system property without the Zookeeper prefix. To set **jute.maxbuffer** to *X*, set **Djute.maxbuffer** to *X* when starting the ZooKeeper client or the service.

For example, set the parameter to 4 MB: **-Djute.maxbuffer=0x400000**.

**Table 31-7** Parameters

Parameter	Description	Default Value
jute.maxbuffer	<p>Specifies the maximum length of data that can be stored in znode. The unit is byte. Default value: 0xfffff, which is less than 1 MB.</p> <p><b>NOTE</b> If this option is changed, the system property must be set on all servers and clients, otherwise problems will arise.</p>	0xfffff

### 31.6.3 Why Four Letter Commands Don't Work With Linux netcat Command When Secure Netty Configurations Are Enabled at Zookeeper Server?

#### Question

Why four letter commands do not work with linux netcat command when secure netty configurations are enabled at Zookeeper server?

For example,

*echo stat /netcat host port*

## Answer

Linux *netcat* command does not have option to communicate Zookeeper server securely, so it cannot support Zookeeper four letter commands when secure netty configurations are enabled.

To avoid this problem, user can use below Java API to execute four letter commands.

```
org.apache.zookeeper.client.FourLetterWordMain
```

For example,

```
String[] args = new String[]{host, port, "stat"};
org.apache.zookeeper.client.FourLetterWordMain.main(args);
```

### NOTE

*netcat* command should be used only with non secure netty configuration.

## 31.6.4 How Do I Check Which ZooKeeper Instance Is a Leader?

### Question

How to check whether the role of a ZooKeeper instance is a leader or follower.

### Answer

Log in to Manager and choose **Cluster** > *Name of the desired cluster* > **Service** > **ZooKeeper** > **Instance**. On the displayed page, click the name of the quorumpeer instance. On the displayed instance details page, view the server status of the instance.

## 31.6.5 Why Cannot the Client Connect to ZooKeeper using the IBM JDK?

### Question

When the IBM JDK is used, the client fails to connect to ZooKeeper.

### Answer

The possible cause is that the **jaas.conf** file format of the IBM JDK is different from that of the common JDK.

If IBM JDK is used, use the following **jaas.conf** template. The **useKeytab** file path must start with **file://**, followed by an absolute path.

```
Client {
 com.ibm.security.auth.module.Krb5LoginModule required
 useKeytab="file://D:/install/HbaseClientSample/conf/user.keytab"
 principal="hbaseuser1"
 credsType="both";
};
```

## 31.6.6 What Should I Do When the ZooKeeper Client Fails to Refresh a TGT?

### Question

The ZooKeeper client fails to refresh a TGT and therefore ZooKeeper cannot be accessed. The error message is as follows:

```
Login: Could not renew TGT due to problem running shell command: '*/kinit -R'; exception was:org.apache.zookeeper.Shell$ExitCodeException: kinit: Ticket expired while renewing credentials
```

### Answer

ZooKeeper uses the system command **kinit - R** to refresh a ticket. In the current version of MRS, the function of this command is canceled. If a long-term task needs to be executed, you are advised to implement the authentication function in keytab mode.

In the **jaas.conf** configuration file, set **useTicketCache** to **false**, **useKeyTab** to **true**, and specify the keytab path.

## 31.6.7 Why Is Message "Node does not exist" Displayed when A Large Number of Znodes Are Deleted Using the deleteall Command

### Question

When the client connects to a non-leader instance, run the **deleteall** command to delete a large number of znodes, the error message "Node does not exist" is displayed, but run the **stat** command, the node status can be obtained.

### Answer

The leader and follower data is not synchronized due to network problems or large data volume. To solve this problem, connect the client to the leader instance and delete the instance. To delete the leader node, view the IP address of the node where the leader resides by referring to [How Do I Check Which ZooKeeper Instance Is a Leader?](#), run the **zkCli.sh -server leader node IP address 2181** command to connect to the client, and then run the **deleteall** command to delete the leader node. For details, see [Using a ZooKeeper Client](#).

# 32 Appendix

---

## 32.1 Modifying Cluster Service Configuration Parameters

- You can modify service configuration parameters on the cluster management page of the MRS management console.
  - a. Log in to the MRS console. In the left navigation pane, choose **Clusters > Active Clusters**, and click a cluster name.
  - b. Choose **Components > Name of the desired service > Service Configuration**.

The **Basic Configuration** tab page is displayed by default. To modify more parameters, click the **All Configurations** tab. 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.
  - c. 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 system searches for the parameter in real time and displays the result.
  - d. Click **Save Configuration**. In the displayed dialog box, click **OK**.
  - e. Wait until the message **Operation successful** is displayed. Click **Finish**.

The configuration is modified.

Check whether there is any service whose configuration has expired in the cluster. If yes, restart the corresponding service or role instance for the configuration to take effect. You can also select **Restart the affected services or instances** when saving the configuration. .
- For MRS 3.x or earlier: You can log in to MRS Manager to modify service configuration parameters.
  - a. Log in to MRS Manager.
  - b. Click **Services**.

- c. Click the specified service name on the service management page.
- d. Click **Service Configuration**.

The **Basic Configuration** tab page is displayed by default. To modify more parameters, click the **All Configurations** tab. 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.

- e. 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 system searches for the parameter in real time and displays the result.

- f. Click **Save**. In the confirmation dialog box, click **OK**.
- g. Wait until the message **Operation successful** is displayed. Click **Finish**.

The configuration is modified.

Check whether there is any service whose configuration has expired in the cluster. If yes, restart the corresponding service or role instance for the configuration to take effect. You can also select **Restart the affected services or instances** when saving the configuration.

- For MRS 3.x or later: You can log in to FusionInsight Manager to modify service configuration parameters.

- a. You have logged in to FusionInsight Manager.
- b. Choose **Cluster > Service**.
- c. Click the specified service name on the service management page.
- d. Click **Configuration**.

The **Basic Configuration** tab page is displayed by default. To modify more parameters, click the **All Configurations** tab. 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.

- e. 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 system searches for the parameter in real time and displays the result.

- f. Click **Save**. In the confirmation dialog box, click **OK**.
- g. Wait until the message **Operation successful** is displayed. Click **Finish**.

The configuration is modified.

Check whether there is any service whose configuration has expired in the cluster. If yes, restart the corresponding service or role instance for the configuration to take effect.

## 32.2 Accessing Manager

## 32.2.1 Accessing MRS Manager (Versions Earlier Than MRS 3.x)

### Scenario

Clusters of versions earlier than MRS 3.x use MRS Manager to monitor, configure, and manage clusters. You can open the MRS Manager page on the MRS console.

### Accessing MRS manager

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

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

**Step 3** Click **Access Manager**. The **Access MRS Manager** page is displayed.

- If you have bound an EIP when creating a cluster
  - a. Select the security group to which the security group rule to be added belongs. The security group is configured when the cluster is created.
  - b. Add a security group rule. By default, your public IP address used for accessing port 9022 is filled in the rule. To enable multiple IP address segments to access MRS Manager, see [Step 6](#) to [Step 9](#). If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

#### NOTE

- It is normal that the automatically generated public IP address is different from the local IP address and no action is required.
- If port 9022 is a Knox port, you need to enable the permission of port 9022 to access Knox for accessing MRS Manager.
- c. Select the checkbox stating that **I confirm that xx.xx.xx.xx is a trusted public IP address and MRS Manager can be accessed using this IP address**.
- If you have not bound an EIP when creating a cluster
  - a. Select an available EIP from the drop-down list or click **Manage EIP** to create one.
  - b. Select the security group to which the security group rule to be added belongs. The security group is configured when the cluster is created.
  - c. Add a security group rule. By default, your public IP address used for accessing port 9022 is filled in the rule. To enable multiple IP address segments to access MRS Manager, see [Step 6](#) to [Step 9](#). If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

 NOTE

- It is normal that the automatically generated public IP address is different from the local IP address and no action is required.
  - If port 9022 is a Knox port, you need to enable the permission of port 9022 to access Knox for accessing MRS Manager.
- d. Select the checkbox stating that **I confirm that xx.xx.xx.xx is a trusted public IP address and MRS Manager can be accessed using this IP address.**

**Step 4** Click **OK**. The MRS Manager login page is displayed.

**Step 5** Enter the default username **admin** and the password set during cluster creation, and click **Log In**. The MRS Manager page is displayed.

**Step 6** On the MRS console, click **Clusters** and choose **Active Clusters**. Click the target cluster name to access the cluster details page.

 NOTE

To assign MRS Manager access permissions to other users, follow instructions from [Step 6](#) to [Step 9](#) to add the users' public IP addresses to the trusted range.

**Step 7** Click **Add Security Group Rule** on the right of **EIP**.

**Step 8** On the **Add Security Group Rule** page, add the IP address segment for users to access the public network and select **I confirm that the authorized object is a trusted public IP address range. Do not use 0.0.0.0/0. Otherwise, security risks may arise.**

By default, the IP address used for accessing the public network is filled. You can change the IP address segment as required. To enable multiple IP address segments, repeat steps [Step 6](#) to [Step 9](#). If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

**Step 9** Click **OK**.

----End

## Granting the Permission to Access MRS Manager to Other Users

**Step 1** On the MRS console, click **Clusters** and choose **Active Clusters**. Click the target cluster name to access the cluster details page.

**Step 2** Click **Add Security Group Rule** on the right of **EIP**.

**Step 3** On the **Add Security Group Rule** page, add the IP address segment for users to access the public network and select **I confirm that the authorized object is a trusted public IP address range. Do not use 0.0.0.0/0. Otherwise, security risks may arise.**

By default, the IP address used for accessing the public network is filled. You can change the IP address segment as required. To enable multiple IP address segments, repeat steps [Step 1](#) to [Step 4](#). If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.



**Step 4** Click **OK**.

----End

## 32.2.2 Accessing FusionInsight Manager (MRS 3.x or Later)

### Scenario

In MRS 3.x or later, FusionInsight Manager is used to monitor, configure, and manage clusters. After the cluster is installed, you can use the account to log in to FusionInsight Manager.

#### NOTE

If you cannot log in to the WebUI of the component, access FusionInsight Manager by referring to [Accessing FusionInsight Manager from an ECS](#).

### Accessing FusionInsight Manager Using EIP

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

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

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

1. If no EIP is bound during MRS cluster creation, select an available EIP from the drop-down list on the right of **IEP**. If you have bound an EIP when you create a cluster, go to [Step 3.2](#).

#### NOTE

- If no EIPs are available, click **Manage EIP** to create one. Then, select the EIP from the drop-down list.
  - To unbind or release an EIP after using it, log in to the **EIPs** page, locate the row containing the target EIP, and click **Unbind** or choose **More > Release** in the **Operation** column.
  - If an EIP has been created but cannot be found during binding, the EIP may have been bound to another cluster. In this case, unbind the EIP on the **EIPs** page and then bind it to the current cluster.
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.

#### NOTE

- When creating a custom cluster, you can configure a security group created in advance or retain the default value **Auto create**. When you quickly create a cluster, the security group is automatically created by the cluster.
  - You can view the security group name in **Security Group** on the **Dashboard** tab page of the cluster.
3. Add a security group rule. By default, the filled-in rule is used to access the EIP. To enable multiple IP address segments to access Manager, see steps [Step 6](#) to [Step 9](#). If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

4. Select the information to be confirmed and click **OK**.

**Step 4** Click **OK**. The Manager login page is displayed.

**Step 5** Enter the default username **admin** and the password set during cluster creation, and click **Log In**. The Manager page is displayed.

**Step 6** On the MRS management console, choose **Clusters > Active Clusters**. Click the target cluster name to access the cluster details page.

 **NOTE**

To grant other users the permission to access Manager, perform **Step 6** to **Step 9** to add the users' public IP addresses to the trusted IP address range.

**Step 7** Click **Add Security Group Rule** next to **EIP**.

**Step 8** 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.**

By default, the IP address used for accessing the public network is filled. You can change the IP address segment as required. To enable multiple IP address segments, repeat steps **Step 6** to **Step 9**. If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

**Step 9** Click **OK**.

----End

## Accessing FusionInsight Manager from an ECS

**Step 1** On the MRS management console, click **Clusters**.

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

Record the **AZ, VPC, MRS ManagerSecurity 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 **Elastic Cloud Server > User Guide > Getting Started > Creating and Logging In to a Windows ECS**.

 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 **Elastic Cloud Server > User Guide > Security Group > 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 **Virtual Private Cloud > User Guide > Security > Security Group > Adding a Security Group Rule**.

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

For details, see **Virtual Private Cloud > User Guide > Elastic IP > Assigning an EIP and Binding It to an ECS**.

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


The Windows system account, password, EIP, and the security group rules are required for logging in to the ECS.

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

The address for accessing Manager is the address of the **MRS Manager** page. Enter the name and password of the cluster user, for example, user **admin**.

 NOTE

- 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

## 32.3 Using an MRS Client

### 32.3.1 Installing a Client (MRS 3.x or Later)

#### Scenario

This section describes how to install clients of all services (excluding Flume) in an MRS cluster. For details about how to install the Flume client, see [Installing the Flume Client](#).

A client can be installed on a node inside or outside the cluster. This section uses the installation directory **//opt/client** as an example. Replace it with the actual one.

## Prerequisites

- A Linux ECS has been prepared. For details about the supported OS of the ECS, see [Table 32-1](#).

**Table 32-1** Reference list

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

In addition, sufficient disk space is allocated for the ECS, for example, 40 GB.

- The ECS and the MRS cluster are in the same VPC.
- The security group of the ECS must be the same as that of the master node in the MRS cluster.
- 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.
- A user can log in to the Linux ECS using the password (in SSH mode).
- All ports in the inbound direction of the MRS cluster security group are open to the client node.

## Installing a Client on a Node Inside a Cluster

1. Obtain the software package.

Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Click the name of the cluster to be operated in the **Cluster** drop-down list.

Choose **More > Download Client**. The **Download Cluster Client** dialog box is displayed.

### NOTE

In the scenario where only one client is to be installed, choose **Cluster > Service > Service name > More > Download Client**. The **Download Client** dialog box is displayed.

2. Set the client type to **Complete Client**.

**Configuration Files Only** is to download client configuration files in the following scenario: After a complete client is downloaded and installed and

administrators modify server configurations on Manager, developers need to update the configuration files during application development.

The platform type can be set to **x86\_64** or **aarch64**.

- **x86\_64**: indicates the client software package that can be deployed on the x86 servers.
- **aarch64**: indicates the client software package that can be deployed on the TaiShan servers.

 **NOTE**

The cluster supports two types of clients: **x86\_64** and **aarch64**. The client type must match the architecture of the node for installing the client. Otherwise, client installation will fail.

3. Select **Save to Path** and click **OK** to generate the client file.

The generated file is stored in the **/tmp/FusionInsight-Client** directory on the active management node by default. You can also store the client file in a directory on which user **omm** has the read, write, and execute permissions. Copy the software package to the file directory on the server where the client is to be installed as user **omm** or **root**.

The name of the client software package is in the follow format:

**FusionInsight\_Cluster\_<Cluster ID>\_Services\_Client.tar**. In this section, the cluster ID **1** is used as an example. Replace it with the actual cluster ID.

The following steps and sections use

**FusionInsight\_Cluster\_1\_Services\_Client.tar** as an example.

- Copy the client installation package to another directory on the current node, for example, **opt/Bigdata/client**.

```
cp -p /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_Services_Client.tar /opt/Bigdata/client
```

- Copy the client installation package to a directory on another node in the cluster, for example, **opt/Bigdata/client**.

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

 **NOTE**

If you cannot obtain the permissions of user **root**, use user **omm**.

4. Log in to the server where the client software package is located as user **user\_client**.
5. Decompress the software package.

Go to the directory where the installation package is stored, for example, **/opt/Bigdata/client**. Run the following command to decompress the installation package to a local directory:

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

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.tar.sha256
```

FusionInsight\_Cluster\_1\_Services\_ClientConfig.tar: OK

7. Decompress the obtained installation file.

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

8. Go to the directory where the installation package is stored, and run the following command to install the client to a specified directory (an absolute path), for example, `/opt/client`:

```
cd/opt/Bigdata/client/FusionInsight_Cluster_1_Services_ClientConfig
```

Run the `./install.sh /opt/client` command to install the client. The client is successfully installed if information similar to the following is displayed:

The component client is installed successfully

#### NOTE

- If the clients of all or some services use the `/opt/client` directory, other directories must be used when you install other service clients.
- You must delete the client installation directory when uninstalling a client.
- To ensure that an installed client can only be used by the installation user (for example, `user_client`), add parameter `-o` during the installation. That is, run the `./install.sh /opt/client -o` command to install the client.
- If an HBase client is installed, it is recommended that the client installation directory contain only uppercase and lowercase letters, digits, and characters (`_-?.@+=`) due to the limitation of the Ruby syntax used by HBase.

## Using a Client

1. On the node where the client is installed, run the `sudo su - omm` command to switch the user. Run the following command to go to the client directory:

```
cd /opt/client
```

2. Run the following command to configure environment variables:

```
source bigdata_env
```

3. If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit MRS cluster user
```

Example: `kinit admin`

#### NOTE

User `admin` is created by default for MRS clusters with Kerberos authentication enabled and is used for administrators to maintain the clusters.

4. Run the client command of a component directly.

For example, run the `hdfs dfs -ls /` command to view files in the HDFS root directory.

## Installing a Client on a Node Outside a Cluster

1. Create an ECS that meets the requirements in [Prerequisites](#).
2. Perform NTP time synchronization to synchronize the time of nodes outside the cluster with that of the MRS cluster.

- a. Run the `vi /etc/ntp.conf` command to edit the NTP client configuration file, add the IP addresses of the master node in the MRS cluster, and comment out the IP address of other servers.

```
server master1_ip prefer
server master2_ip
```

Figure 32-1 Adding the master node IP addresses

```
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 10.9.2.38 prefer
#server 10.9.2.39
#broadcast 192.168.1.255 autokey # broadcast server
#broadcastclient # broadcast client
#broadcast autokey # multicast server
#multicastclient # multicast client
#manycastserver # manycast server
#manycastclient autokey # manycast client

Enable public key cryptography.
#crypto
```

- b. Run the `service ntpd stop` command to stop the NTP service.
- c. Run the following command to manually synchronize the time:  
`/usr/sbin/ntpdate 192.168.10.8`

 NOTE

192.168.10.8 indicates the IP address of the active Master node.

- d. Run the `service ntpd start` or `systemctl restart ntpd` command to start the NTP service.
  - e. Run the `ntpstat` command to check the time synchronization result.
3. Perform the following steps to download the cluster client software package from FusionInsight Manager, copy the package to the ECS node, and install the client:
    - a. Log in to FusionInsight Manager and download the cluster client to the specified directory on the active management node by referring to [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#) and [Installing a Client on a Node Inside a Cluster](#).

- b. Log in to the active management node as user **root** and run the following command to copy the client installation package to the target node:  

```
scp -p /tmp/FusionInsight-Client/
FusionInsight_Cluster_1_Services_Client.tar IP address of the node
where the client is to be installed:/tmp
```
- c. Log in to the node on which the client is to be installed as the client user. Run the following commands to install the client. If the user does not have operation permissions on the client software package and client installation directory, grant the permissions using the **root** user.  

```
cd /tmp
tar -xvf FusionInsight_Cluster_1_Services_Client.tar
tar -xvf FusionInsight_Cluster_1_Services_ClientConfig.tar
cd FusionInsight_Cluster_1_Services_ClientConfig
./install.sh /opt/client
```
- d. Run the following commands to switch to the client directory and configure environment variables:  

```
cd /opt/client
source bigdata_env
```
- e. If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step.  

```
kinit MRS cluster user
```

Example: **kinit admin**
- f. Run the client command of a component directly.  
For example, run the **hdfs dfs -ls /** command to view files in the HDFS root directory.

## 32.3.2 Installing a Client (Versions Earlier Than 3.x)

### Scenario

An MRS client is required. The MRS cluster client can be installed on the Master or Core node in the cluster or on a node outside the cluster.

After a cluster of versions earlier than MRS 3.x is created, a client is installed on the active Master node by default. You can directly use the client. The installation directory is **/opt/client**.

For details about how to install a client of MRS 3.x or later, see [Installing a Client \(MRS 3.x or Later\)](#).

#### NOTE

If a client has been installed on the node outside the MRS cluster and the client only needs to be updated, update the client using the user who installed the client, for example, user **root**.



## Prerequisites

- An ECS has been prepared. For details about the OS and its version of the ECS, see [Table 32-2](#).

**Table 32-2** Reference list

OS	Supported Version
EulerOS	<ul style="list-style-type: none"><li>• Available: EulerOS 2.2</li><li>• Available: EulerOS 2.3</li><li>• Available: EulerOS 2.5</li></ul>

For example, a user can select an ECS running the EulerOS.

In addition, sufficient disk space is allocated for the ECS, for example, 40 GB.

- The ECS and the MRS cluster are in the same VPC.
- The security group of the ECS is the same as that of the Master node of the MRS cluster.

If this requirement is not met, modify the ECS security group or configure the inbound and outbound rules of the ECS security group to allow the ECS security group to be accessed by all security groups of MRS cluster nodes.

- To enable users to log in to a Linux ECS using a password (SSH), see *Instances > Logging In to a Linux ECS > Login Using an SSH Password in the Elastic Cloud Server User Guide*.
- All ports in the inbound direction of the MRS cluster security group are open to the client node.

## Installing a Client on the Core Node

1. Log in to MRS Manager and choose **Services > Download Client** to download the client installation package to the active management node.

### NOTE

If only the client configuration file needs to be updated, see method 2 in [Updating a Client \(Versions Earlier Than 3.x\)](#).

2. Use the IP address to search for the active management node, and log in to the active management node using VNC.
3. Log in to the active management node, and run the following command to switch the user:

```
sudo su - omm
```

4. On the MRS management console, view the IP address on the **Nodes** tab page of the specified cluster.

Record the IP address of the Core node where the client is to be used.

5. On the active management node, run the following command to copy the client installation package to the Core node:

```
scp -p /tmp/MRS-client/MRS_Services_Client.tar IP address of the Core node:/opt/client
```

6. Log in to the Core node as user **root**.  
Master nodes support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation.
7. Run the following commands to install the client:  

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

  
For example, run the following command:  

```
./install.sh /opt/client
```
8. For details about how to use the client, see [Using an MRS Client](#).


## Using an MRS Client

1. On the node where the client is installed, run the **sudo su - omm** command to switch the user. Run the following command to go to the client directory:  

```
cd /opt/client
```
2. Run the following command to configure environment variables:  

```
source bigdata_env
```
3. If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step.  

```
kinit MRS cluster user
```

  
Example: **kinit admin**  
  
 **NOTE**  
  
User **admin** is created by default for MRS clusters with Kerberos authentication enabled and is used for administrators to maintain the clusters.
4. Run the client command of a component directly.  
For example, run the **hdfs dfs -ls /** command to view files in the HDFS root directory.

## Installing a Client on a Node Outside the Cluster

- Step 1** Create an ECS that meets the requirements in the prerequisites.
- Step 2** Log in to MRS Manager. For details, see [Accessing MRS Manager \(Versions Earlier Than MRS 3.x\)](#). Then, choose **Services**.
- Step 3** Click **Download Client**.
- Step 4** In **Client Type**, select **All client files**.
- Step 5** In **Download To**, select **Remote host**.
- Step 6** Set **Host IP Address** to the IP address of the ECS, **Host Port** to **22**, and **Save Path** to **/tmp**.

- If the default port **22** for logging in to an ECS using SSH has been changed, set **Host Port** to the new port.
- **Save Path** contains a maximum of 256 characters.

**Step 7** Set **Login User** to **root**.

If other users are used, ensure that the users have read, write, and execute permission on the save path.

**Step 8** 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 9** Click **OK** to generate a client file.

If the following information is displayed, the client package is saved. Click **Close**. Obtain the client file from the save path on the remote host that is set when the client is downloaded.

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.

 **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 10** Log in to the ECS using VNC. For details, see **Instance > Logging In to a Linux > Logging In to a Linux** in the *Elastic Cloud Server User Guide*

All images support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation. It is recommended that you change the password upon the first login.

**Step 11** Perform NTP time synchronization to synchronize the time of nodes outside the cluster with the time of the MRS cluster.

1. Check whether the NTP service is installed. If it is not installed, run the **yum install ntp -y** command to install it.
2. Run the **vim /etc/ntp.conf** command to edit the NTP client configuration file, add the IP address of the Master node in the MRS cluster, and comment out the IP addresses of other servers.

```
server master1_ip prefer
server master2_ip
```

Figure 32-2 Adding the Master node IP addresses

```
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 # manycast server
#manycastclient [redacted] autokey # manycast client
#
Enable public key cryptography.
#crypto
```

3. Run the **service ntpd stop** command to stop the NTP service.
4. Run the following command to manually synchronize the time:  
**/usr/sbin/ntpdate 192.168.10.8**

 NOTE

192.168.10.8 indicates the IP address of the active Master node.

5. Run the **service ntpd start** or **systemctl restart ntpd** command to start the NTP service.
6. Run the **ntpstat** command to check the time synchronization result:

**Step 12** On the ECS, switch to user **root** and copy the installation package in **Save Path** in **Step 6** to the **/opt** directory. For example, if **Save Path** is set to **/tmp**, run the following commands:

```
sudo su - root
cp /tmp/MRS_Services_Client.tar /opt
```

**Step 13** Run the following command in the **/opt** directory to decompress the package and obtain the verification file and the configuration package of the client:

```
tar -xvf MRS_Services_Client.tar
```

**Step 14** Run the following command to verify the configuration file package of the client:

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

The command output is as follows:

```
MRS_Services_ClientConfig.tar: OK
```

**Step 15** Run the following command to decompress **MRS\_Services\_ClientConfig.tar**:

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

**Step 16** Run the following command to install the client to a new directory, for example, **/opt/Bigdata/client**. A directory is automatically generated during the client installation.

```
sh /opt/MRS_Services_ClientConfig/install.sh /opt/Bigdata/client
```

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

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

**Step 17** Check whether the IP address of the ECS node is connected to the IP address of the cluster Master node.

For example, run the following command: **ping** *Master node IP address*.

- If yes, go to **Step 18**.
- If no, check whether the VPC and security group are correct and whether the ECS and the MRS cluster are in the same VPC and security group, and go to **Step 18**.

**Step 18** Run the following command to configure environment variables:

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

**Step 19** If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit MRS cluster user
```

Example: **kinit admin**

**Step 20** Run the client command of a component.

For example, run the following command to query the HDFS directory:

```
hdfs dfs -ls /
```

```
----End
```

### 32.3.3 Updating a Client (Version 3.x or Later)

A 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.

#### Updating the Client Configuration

**Method 1:**

**Step 1** Log in to FusionInsight Manager. For details, see [Accessing FusionInsight Manager \(MRS 3.x or Later\)](#). Click the name of the cluster to be operated in the **Cluster** drop-down list.

**Step 2** Choose **More > Download Client > Configuration Files Only**.

The generated compressed file contains the configuration files of all services.

**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 FusionInsight Manager.

- 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

## 32.3.4 Updating a Client (Versions Earlier Than 3.x)

### NOTE

This section applies to clusters of versions earlier than MRS 3.x. For MRS 3.x or later, see [Updating a Client \(Version 3.x or Later\)](#).

## Updating a Client Configuration File

### Scenario

An MRS cluster provides a client for you to connect to a server, view task results, or manage data. Before using an MRS client, you need to download and update the client configuration file if service configuration parameters are modified and a service is restarted or the service is merely restarted on MRS Manager.

During cluster creation, the original client is stored in the **/opt/client** directory on all nodes in the cluster by default. 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.

### Procedure

#### Method 1:

- Step 1** Log in to MRS Manager. For details, see [Accessing MRS Manager \(Versions Earlier Than MRS 3.x\)](#). Then, 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.

- 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** Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

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

**Step 6** Run the following command to update client configurations:

```
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-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** Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

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

**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 successfully.

```
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

## Fully Updating the Original Client of the Active Master Node

### Scenario

During cluster creation, the original client is stored in the **/opt/client** directory on all nodes in the cluster by default. The following uses **/opt/Bigdata/client** as an example.

- For a normal MRS cluster, you will use the pre-installed client on a Master node to submit a job on the management console page.
- You can also use the pre-installed client on the Master node to connect to a server, view task results, and manage data.

After installing the patch on the cluster, you need to update the client on the Master node to ensure that the functions of the built-in client are available.

### Procedure

**Step 1** Log in to MRS Manager. For details, see [Accessing MRS Manager \(Versions Earlier Than MRS 3.x\)](#). Then, choose **Services**.

**Step 2** Click **Download Client**.

Set **Client Type** to **All client 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.

**Step 3** Query and log in to the active Master node.

**Step 4** On the ECS, switch to user **root** and copy the installation package to the **/opt** directory.

```
sudo su - root
```

```
cp /tmp/MRS-client/MRS_Services_Client.tar /opt
```

**Step 5** Run the following command in the **/opt** directory to decompress the package and obtain the verification file and the configuration package of the client:

```
tar -xvf MRS_Services_Client.tar
```

**Step 6** Run the following command to verify the configuration file package of the client:

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

The command output is as follows:

```
MRS_Services_ClientConfig.tar: OK
```

**Step 7** Run the following command to decompress **MRS\_Services\_ClientConfig.tar**:

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

**Step 8** Run the following command to move the original client to the **/opt/Bigdata/client\_bak** directory:

```
mv /opt/Bigdata/client /opt/Bigdata/client_bak
```

**Step 9** Run the following command to install the client in a new directory. The client path must be **/opt/Bigdata/client**.

```
sh /opt/MRS_Services_ClientConfig/install.sh /opt/Bigdata/client
```

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

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

**Step 10** Run the following command to modify the user and user group of the **/opt/Bigdata/client** directory:

```
chown omm:wheel /opt/Bigdata/client -R
```

**Step 11** Run the following command to configure environment variables:

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

**Step 12** If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step.

```
kinit MRS cluster user
```

Example: **kinit admin**

**Step 13** Run the client command of a component.

For example, run the following command to query the HDFS directory:

```
hdfs dfs -ls /
```

```
----End
```

## Fully Updating the Original Client of the Standby Master Node

**Step 1** Repeat **Step 1** to **Step 3** to log in to the standby Master node, and run the following command to switch to user **omm**:

```
sudo su - omm
```

**Step 2** Run the following command on the standby master node to copy the downloaded client package from the active master node:

```
scp omm@master1 nodeIP address:/tmp/MRS-client/
MRS_Services_Client.tar /tmp/MRS-client/
```

 NOTE

- In this command, **master1** node is the active master node.
- **/tmp/MRS-client/** is an example target directory of the standby master node.

**Step 3** Repeat **Step 4** to **Step 13** to update the client of the standby Master node.

----End