

# MapReduce Service

## FAQs

|       |            |
|-------|------------|
| Issue | 01         |
| Date  | 2025-03-17 |



**Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.**

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

## **Trademarks and Permissions**



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

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

## **Notice**

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

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

---

# Contents

---

|                                                                                                                                                                           |           |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>1 MRS Basics.....</b>                                                                                                                                                  | <b>1</b>  |
| 1.1 What Is MRS Used For?.....                                                                                                                                            | 1         |
| 1.2 What Types of Distributed Storage Does MRS Support?.....                                                                                                              | 1         |
| 1.3 What Are Regions and AZs?.....                                                                                                                                        | 2         |
| 1.4 Can I Change the Network Segment of Nodes in an MRS Cluster?.....                                                                                                     | 3         |
| 1.5 Can I Downgrade the Specifications of an MRS Cluster Node?.....                                                                                                       | 4         |
| 1.6 Are Hive Components of Different Versions Compatible with Each Other?.....                                                                                            | 4         |
| 1.7 What Are the Differences Between OBS and HDFS in Data Storage?.....                                                                                                   | 4         |
| 1.8 What Are the Solutions for Processing 1 Billion Data Records?.....                                                                                                    | 5         |
| 1.9 What are the advantages of the compression ratio of zstd?.....                                                                                                        | 5         |
| <b>2 Billing.....</b>                                                                                                                                                     | <b>6</b>  |
| 2.1 Why Is the Price Not Displayed During MRS Cluster Creation?.....                                                                                                      | 6         |
| 2.2 How Is Auto Scaling Billed for an MRS Cluster?.....                                                                                                                   | 6         |
| 2.3 How Are Task Nodes in an MRS Cluster Billed?.....                                                                                                                     | 7         |
| 2.4 Why Does My Unsubscription from ECS Fail After I Unsubscribe from MRS?.....                                                                                           | 7         |
| <b>3 Cluster Creation.....</b>                                                                                                                                            | <b>8</b>  |
| 3.1 How Do I Create an MRS Cluster Using a Custom Security Group?.....                                                                                                    | 8         |
| 3.2 What Should I Do If HDFS, Yarn, and MapReduce Components Are Unavailable When I Buy an MRS Cluster?.....                                                              | 8         |
| 3.3 What Should I Do If the ZooKeeper Component Is Unavailable When I Buy an MRS Cluster?.....                                                                            | 9         |
| 3.4 What Should I Do If Invalid Authentication Is Reported When I Submit an Order for Purchasing an MRS Cluster? .....                                                    | 9         |
| <b>4 Web Page Access.....</b>                                                                                                                                             | <b>10</b> |
| 4.1 How Do I Change the Session Timeout Duration for an Open Source Component Web UI?.....                                                                                | 10        |
| 4.2 What Can I Do If the Dynamic Resource Plan Page in MRS Tenant Management Cannot Be Refreshed?.....                                                                    | 13        |
| 4.3 What Do I Do If the Kafka Topic Monitoring Tab Is Not Displayed on Manager?.....                                                                                      | 13        |
| 4.4 What CAN I DO IF an Error Is Reported or Some Functions Are Unavailable When I Access the Web UIs of Components such as HDFS, Hue, Yarn, Flink, and HetuEngine? ..... | 14        |
| 4.5 How Do I Switch the Methods to Access MRS Manager? .....                                                                                                              | 15        |
| 4.6 Why Cannot I Find the User Management Page on MRS Manager?.....                                                                                                       | 15        |
| 4.7 What Can I Do If the Excel File Downloaded by Hue Cannot Be Opened?.....                                                                                              | 16        |
| <b>5 Authentication and Permission.....</b>                                                                                                                               | <b>18</b> |

|                                                                                                                                                          |           |
|----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| 5.1 What Is the User for Logging in to FusionInsight Manager?.....                                                                                       | 18        |
| 5.2 How Do I Query and Change the Password Validity Period of a User In a Cluster?.....                                                                  | 18        |
| 5.3 Does an MRS Cluster Support Access Permission Control If Kerberos Authentication Is not Enabled?<br>.....                                            | 19        |
| 5.4 How Do I Add Tenant Management Permission to Users in a Cluster?.....                                                                                | 20        |
| 5.5 Does Hue Provide the Function of Configuring Account Permissions?.....                                                                               | 21        |
| 5.6 Why Can't I Submit Jobs on the Console After My IAM Account Is Assigned with MRS Permissions?<br>.....                                               | 21        |
| 5.7 How Do I View the Hive Table Created by Another User?.....                                                                                           | 21        |
| 5.8 How Do I Prevent Kerberos Authentication Expiration?.....                                                                                            | 23        |
| 5.9 How Do I Enable or Disable Kerberos Authentication for an Existing MRS Cluster?.....                                                                 | 24        |
| 5.10 What Are the Ports of the Kerberos Authentication Service?.....                                                                                     | 24        |
| <b>6 Client Usage.....</b>                                                                                                                               | <b>25</b> |
| 6.1 How Do I Disable SASL Authentication for ZooKeeper?.....                                                                                             | 25        |
| 6.2 What Can I Do If the Error Message "Permission denied" Is Displayed When kinit Is Executed on a<br>Client Outside the MRS Cluster?.....              | 25        |
| 6.3 What Should I Do If an Alarm Is Reported Indicating that the Memory Is Insufficient When I Execute a<br>SQL Statement on the ClickHouse Client?..... | 26        |
| 6.4 How Do I Connect to Spark Shell from MRS?.....                                                                                                       | 26        |
| 6.5 How Do I Connect to Spark Beeline from MRS?.....                                                                                                     | 27        |
| 6.6 What Should I Do If the Connection to the ClickHouse Server Fails and Error Code 516 Is Reported?<br>.....                                           | 28        |
| <b>7 Component Configurations.....</b>                                                                                                                   | <b>29</b> |
| 7.1 Does MRS Support Running Hive on Kudu?.....                                                                                                          | 29        |
| 7.2 Does an MRS Cluster Support Hive on Spark?.....                                                                                                      | 29        |
| 7.3 Can I Change the IP address of DBService?.....                                                                                                       | 29        |
| 7.4 What Access Protocols Does Kafka Support?.....                                                                                                       | 30        |
| 7.5 What Python Versions Are Supported by Spark Tasks in an MRS Cluster?.....                                                                            | 30        |
| 7.6 What Are the Restrictions on the Storm Log Size in an MRS 2.1.0 Cluster?.....                                                                        | 30        |
| 7.7 How Do I Modify the HDFS fs.defaultFS of an Existing Cluster?.....                                                                                   | 30        |
| 7.8 Can MRS Run Multiple Flume Tasks at a Time?.....                                                                                                     | 31        |
| 7.9 How Do I Change FlumeClient Logs to Standard Logs?.....                                                                                              | 31        |
| 7.10 Where Are the JAR Files and Environment Variables of Hadoop Stored?.....                                                                            | 31        |
| 7.11 How Do I View HBase Logs?.....                                                                                                                      | 32        |
| 7.12 How Do I Set the TTL for an HBase Table?.....                                                                                                       | 32        |
| 7.13 How Do I Change the Number of HDFS Replicas?.....                                                                                                   | 32        |
| 7.14 How Do I Modify the HDFS Active/Standby Switchover Class?.....                                                                                      | 32        |
| 7.15 What Data Type in Hive Tables Is Recommended for the Number Type of DynamoDB?.....                                                                  | 33        |
| 7.16 Can I Export the Query Result of Hive Data?.....                                                                                                    | 33        |
| 7.17 What Should I Do If an Error Occurs When Hive Runs the <b>beeline -e</b> Command to Execute Multiple<br>Statements?.....                            | 33        |
| 7.18 What Do I Do If "over max user connections" Is Displayed When Hue Connects to HiveServer?.....                                                      | 34        |
| 7.19 How Do I View MRS Hive Metadata?.....                                                                                                               | 34        |

|                                                                                                                                   |           |
|-----------------------------------------------------------------------------------------------------------------------------------|-----------|
| 7.20 How Do I Reset Kafka Data?.....                                                                                              | 35        |
| 7.21 What Should I Do If the Error Message "Not Authorized to access group XXX" Is Displayed When Kafka Topics Are Consumed?..... | 35        |
| 7.22 What Compression Algorithms Does Kudu Support?.....                                                                          | 36        |
| 7.23 How Do I View Kudu Logs?.....                                                                                                | 36        |
| 7.24 How Do I Handle the Kudu Service Exceptions Generated During Cluster Creation?.....                                          | 36        |
| 7.25 How Do I Configure Other Data Sources on Presto?.....                                                                        | 37        |
| 7.26 How Do I Update the Ranger Certificate in MRS 1.9.3?.....                                                                    | 39        |
| 7.27 How Do I Specify a Log Path When Submitting a Task in an MRS Storm Cluster?.....                                             | 41        |
| 7.28 How Do I Check the ResourceManager Configuration of Yarn?.....                                                               | 41        |
| 7.29 How Do I Modify the allow_drop_detached Parameter of ClickHouse?.....                                                        | 44        |
| 7.30 How Do I Add a Periodic Deletion Policy to Prevent Large ClickHouse System Table Logs?.....                                  | 45        |
| 7.31 How Do I Change the Time Zone of the ClickHouse Service?.....                                                                | 45        |
| <b>8 Cluster Management.....</b>                                                                                                  | <b>47</b> |
| 8.1 How Do I View All Clusters?.....                                                                                              | 47        |
| 8.2 How Do I View MRS Operation Logs?.....                                                                                        | 47        |
| 8.3 How Do I View MRS Cluster Configuration Information?.....                                                                     | 48        |
| 8.4 How Do I Add Components to an MRS Cluster?.....                                                                               | 48        |
| 8.5 How Do I Cancel Message Notification for Cluster Alarms?.....                                                                 | 48        |
| 8.6 Why Is the Resource Pool Memory Displayed in the MRS Cluster Smaller Than the Actual Cluster Memory?.....                     | 49        |
| 8.7 What Is the Python Version Installed for an MRS Cluster?.....                                                                 | 49        |
| 8.8 How Do I Upload a Local File to a Node Inside a Cluster?.....                                                                 | 49        |
| 8.9 What Can I Do If the Time Information of an MRS Cluster Node Is Incorrect?.....                                               | 50        |
| 8.10 What Are the Differences and Relationships Between the MRS Management Console and MRS Manager?.....                          | 50        |
| 8.11 How Do I Unbind an EIP from FusionInsight Manager of an MRS Cluster?.....                                                    | 52        |
| 8.12 How Do I Stop the Firewall Service?.....                                                                                     | 53        |
| 8.13 How Do I Switch the Login Mode of a Node in an MRS Cluster?.....                                                             | 53        |
| 8.14 How Do I Access an MRS Cluster from a Node Outside the Cluster?.....                                                         | 54        |
| 8.15 In an MRS Streaming Cluster, Can the Kafka Topic Monitoring Function Send Alarm Notifications?.....                          | 55        |
| 8.16 Where can I view the running resource queues when ALM-18022 Insufficient Yarn Queue Resources is generated?.....             | 56        |
| 8.17 How Do I Understand the Multi-Level Chart Statistics in the HBase Operation Requests Metric?.....                            | 56        |
| <b>9 Node Management.....</b>                                                                                                     | <b>58</b> |
| 9.1 What are the Operating Systems of Hosts in MRS Clusters of Different Versions?.....                                           | 58        |
| 9.2 Do I Need to Shut Down a Master Node Before Upgrading It?.....                                                                | 58        |
| 9.3 Can I Change MRS Cluster Nodes on the MRS Console?.....                                                                       | 59        |
| 9.4 How Do I Query the Startup Time of an MRS Node?.....                                                                          | 59        |
| 9.5 What Do I Do If Trust Relationships Between Nodes Are Abnormal?.....                                                          | 59        |
| 9.6 Can Master Node Specifications Be Adjusted in an MRS Cluster?.....                                                            | 60        |
| 9.7 Can Sudo Logs of Nodes in an MRS Cluster Be Cleared?.....                                                                     | 60        |

|                                                                                                                                                            |           |
|------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| 9.8 How Do I Partition Disks in an MRS Cluster?.....                                                                                                       | 61        |
| 9.9 Does an MRS Cluster Support System Reinstallation?.....                                                                                                | 63        |
| 9.10 Can I Change the OS of an MRS Cluster?.....                                                                                                           | 64        |
| <b>10 Component Management.....</b>                                                                                                                        | <b>65</b> |
| 10.1 Can I Delete Components Installed in an MRS Cluster?.....                                                                                             | 65        |
| 10.2 How Do I View the Configuration File Directory of Each Component?.....                                                                                | 65        |
| 10.3 Will Upper-Layer Services Be Affected If the Hive Service Status Is Partially Healthy?.....                                                           | 66        |
| 10.4 How Can I Obtain the IP Address and Port Number of a ZooKeeper Instance?.....                                                                         | 67        |
| <b>11 Job Management.....</b>                                                                                                                              | <b>69</b> |
| 11.1 What Types of Spark Jobs Can Be Submitted in a Cluster?.....                                                                                          | 69        |
| 11.2 What Should I Do If Error 408 Is Reported When an MRS Node Accesses OBS? .....                                                                        | 69        |
| 11.3 How Do I Enable Different Service Programs to Use Different Yarn Queues?.....                                                                         | 70        |
| 11.4 What Should I Do If a Job Fails to Be Submitted and the Error Is Related to OBS?.....                                                                 | 73        |
| 11.5 Can I Run Multiple Spark Tasks at the Same Time After the Minimum Tenant Resources of an MRS Cluster Is Changed to 0?.....                            | 74        |
| 11.6 What Should I Do If Job Parameters Separated By Spaces Cannot Be Identified? .....                                                                    | 74        |
| 11.7 What Are the Differences Between the Client Mode and Cluster Mode of Spark Jobs?.....                                                                 | 75        |
| 11.8 How Do I View MRS Job Logs?.....                                                                                                                      | 75        |
| 11.9 What Can I Do If the System Displays a Message Indicating that the Current User Does Not Exist on Manager When I Submit a Job?.....                   | 76        |
| 11.10 What Can I Do If LauncherJob Fails to Be Executed and the Error Message "jobPropertiesMap is null" Is Displayed?.....                                | 76        |
| 11.11 What Should I Do If the Flink Job Status on the MRS Console Is Inconsistent with That on Yarn?...77                                                  | 77        |
| 11.12 What Can I Do If a SparkStreaming Job Fails After Running for Dozens of Hours and Error 403 Is Reported for OBS Access?.....                         | 77        |
| 11.13 What Should I Do If Error Message "java.io.IOException: Connection reset by peer" Is Displayed During the Execution of a Spark Job?.....             | 77        |
| 11.14 What Should I Do If the Error Message "requestId=XXX" Is Displayed When a Spark Job Accesses OBS?.....                                               | 78        |
| 11.15 What Should I Do If the Error Message "UnknownScannerException" Is Displayed for Spark Jobs? .....                                                   | 78        |
| 11.16 What Can I Do If DataArts Studio Occasionally Fails to Schedule Spark Jobs?.....                                                                     | 78        |
| 11.17 What Should I Do If a Flink Job Fails to Execute and the Error Message "java.lang.NoSuchFieldError: SECURITY_SSL_ENCRYPT_ENABLED" Is Displayed?..... | 79        |
| 11.18 What Should I Do If Submitted Yarn Jobs Cannot Be Viewed on the Web UI?.....                                                                         | 79        |
| 11.19 What Can I Do If launcher-job Is Terminated by Yarn When a Flink Task Is Submitted?.....                                                             | 79        |
| 11.20 What Should I Do If the Error Message "slot request timeout" Is Displayed When I Submit a Flink Job?.....                                            | 80        |
| 11.21 FAQs About Importing and Exporting Data Using DistCP Jobs.....                                                                                       | 80        |
| 11.22 How Do I View SQL Statements of Hive Jobs on the Yarn Web UI?.....                                                                                   | 81        |
| 11.23 How Do I View Logs of a Specified Yarn Task?.....                                                                                                    | 82        |
| 11.24 What Should I Do If a HiveSQL/HiveScript Job Fails to be Submitted After Hive Is Added?.....                                                         | 82        |
| 11.25 Where Are the Execution Logs of Spark Jobs Stored?.....                                                                                              | 82        |

|                                                                                                                                            |            |
|--------------------------------------------------------------------------------------------------------------------------------------------|------------|
| 11.26 What Should I Do If an Alarm Indicating Insufficient Memory Is Reported During Spark Task Execution?.....                            | 83         |
| 11.27 What Can I Do If an Alarm is Generated Because the NameNode Is not Restarted on Time After the hdfs-site.xml File Is Modified? ..... | 83         |
| 11.28 What Should I Do If It Takes a Long Time for Spark SQL to Access Hive Partitioned Tables Before a Job Starts?.....                   | 83         |
| 11.29 How Does the System Select the Queue When an MRS Cluster User Is Associated to Multiple Queues?.....                                 | 84         |
| <b>12 Performance Tuning.....</b>                                                                                                          | <b>86</b>  |
| 12.1 How Do I Obtain the Hadoop Pressure Test Tool?.....                                                                                   | 86         |
| 12.2 How Do I Improve the Resource Utilization of Core Nodes in a Cluster?.....                                                            | 86         |
| 12.3 How Do I Configure the Knox Memory?.....                                                                                              | 87         |
| 12.4 How Do I Adjust the Memory Size of the manager-executor Process?.....                                                                 | 87         |
| 12.5 How Do I Configure Spark Jobs to Automatically Obtain More Resources During Execution?.....                                           | 88         |
| 12.6 What Should I Do If the spark.yarn.executor.memoryOverhead Setting Does Not Take Effect?.....                                         | 89         |
| 12.7 How Do I Improve Presto Resource Usage?.....                                                                                          | 89         |
| <b>13 Application Development.....</b>                                                                                                     | <b>91</b>  |
| 13.1 How Do I Get My Data into OBS or HDFS?.....                                                                                           | 91         |
| 13.2 Can MRS Write Data to HBase Through an HBase External Table of Hive?.....                                                             | 92         |
| 13.3 Where Can I Download the Dependency Package (com.huawei.gaussc10) in the Hive Sample Project?.....                                    | 92         |
| 13.4 Does MRS Support Python Code?.....                                                                                                    | 92         |
| 13.5 Does OpenTSDB Support Python APIs?.....                                                                                               | 92         |
| 13.6 How Do I Obtain a Spark JAR File?.....                                                                                                | 92         |
| 13.7 How Do I Configure the node_id Parameter When Calling the API for Adjusting Cluster Nodes?.....                                       | 93         |
| 13.8 How Do I Manage and Use Third-Party JAR Packages for MRS Cluster Components?.....                                                     | 93         |
| <b>14 Peripheral Service Interconnection.....</b>                                                                                          | <b>95</b>  |
| 14.1 Does MRS Support Read and Write Operations on DLI Service Tables?.....                                                                | 95         |
| 14.2 Does OBS Support the ListObjectsV2 Protocol?.....                                                                                     | 95         |
| 14.3 Can a Crawler Service Be Deployed on Nodes in an MRS Cluster?.....                                                                    | 95         |
| 14.4 Does MRS Support Secure Deletion?.....                                                                                                | 95         |
| 14.5 How Do I Use PySpark to Connect MRS Spark?.....                                                                                       | 96         |
| 14.6 Why Mapped Fields Do not Exist in the Database After HBase Synchronizes Data to CSS?.....                                             | 96         |
| 14.7 Can MRS Connect to an External KDC?.....                                                                                              | 96         |
| 14.8 What Can I Do If Jetty Is Incompatible with MRS when Open-Source Kylin 3.x Is Interconnected with MRS 1.9.3?.....                     | 96         |
| 14.9 What Should I Do If Data Failed to Be Exported from MRS to an OBS Encrypted Bucket?.....                                              | 97         |
| 14.10 How Do I Interconnect MRS with LTS? .....                                                                                            | 97         |
| 14.11 How Do I Install HSS on MRS Cluster Nodes?.....                                                                                      | 100        |
| 14.12 How Do I Connect to HBase of MRS Through HappyBase?.....                                                                             | 101        |
| 14.13 Can the Hive Driver Be Interconnected with DBCP2?.....                                                                               | 102        |
| <b>15 Upgrade and Patching.....</b>                                                                                                        | <b>103</b> |

15.1 How Do I Upgrade an MRS Cluster?..... 103

15.2 Does MRS Support Kernel Version Upgrade of Components in a Cluster?..... 103



# 1 MRS Basics

---

## 1.1 What Is MRS Used For?

### Question:

What Can MapReduce Service (MRS) Do?

### Answer:

MapReduce Service (MRS) is an enterprise-grade big data platform that allows you to quickly build and operate economical, secure, full-stack, cloud-native big data environments on the cloud. It provides engines such as ClickHouse, Spark, Flink, Kafka, and HBase, and supports convergence of data lake, data warehouse, business intelligence (BI), and artificial intelligence (AI). Fully compatible with open-source components, MRS helps you rapidly innovate and expand service growth.

## 1.2 What Types of Distributed Storage Does MRS Support?

### Question:

What Types of Distributed Storage Do MRS Clusters Support? What Versions Are Available?

### Answer:

The MRS clusters use the mainstream big data component Hadoop. Currently, Hadoop 3.x is supported and will be updated as the cluster evolves.

MRS allows you to store data in OBS and use an MRS cluster for data computing only. In this way, storage and compute are decoupled.

For details about the components supported by each MRS version, see [List of MRS Component Versions](#).


## 1.3 What Are Regions and AZs?

A region and an availability zone (AZ) identify the location of a data center. You can create resources in a specific AZ of a region.

- Regions are defined in terms of geographical locations and network latencies. Public services such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS) are shared within the same region. Regions are classified into common regions and dedicated regions. A common region provides common cloud services available to all tenants. A dedicated region provides services of a specific type or only for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. The computing, network, storage, and other resources in an AZ are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to allow you to build cross-AZ high-availability systems.

**Figure 1-1** shows the relationship between regions and AZs.

**Figure 1-1** Regions and AZs



HUAWEI CLOUD provides services worldwide. You can select a region and AZ as required. For more information, see [HUAWEI CLOUD Global Regions](#).

### What Are the Considerations for Selecting a Region?

When selecting a region, consider the following factors:

- Location  
You are advised to select a region closest to your target users. This reduces network latency and improves access rate. However, there is no distinct difference between the infrastructure, BGP network quality, and resource operations and configurations in the same region/country. Therefore, you do not need to consider the network latency when selecting a region in a country.

- If you or your target users are in the Asia Pacific area excluding the Chinese mainland, select the **CN-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
- If you or your target users are in Africa, select the **AF-Johannesburg** region.
- If you or your target users are in Europe, select the **EU-Paris** region.
- Resource price  
Resource prices may vary in different regions. For details, see [Product Pricing Details](#).

## What Are the Considerations for Selecting an AZ?

Consider the requirements for disaster recovery (DR) and network latency when selecting an AZ:

- Deploy resources in different AZs in the same region for DR purposes.
- Deploy resources in the same AZ for minimum latency.

## How Do I Obtain an Endpoint?

When you use resources by using an API, you must specify its regional endpoint. For details about Huawei Cloud regions and endpoints, see [Regions and Endpoints](#).

# 1.4 Can I Change the Network Segment of Nodes in an MRS Cluster?

### Question:

Can I Change the Network Segment of Nodes in an MRS Cluster?

### Answer:

Yes, you can. Perform the following steps to change the network segment of nodes in a cluster:

1. On the **Dashboard** page of the cluster, locate the **Network Information** pane, and click **Change Subnet** on the right of subnet name.
2. Select another subnet in the VPC to which the current cluster belongs to expand the available subnet IP addresses.

Selecting a new subnet will not change the IP addresses and subnets of existing nodes.

## 1.5 Can I Downgrade the Specifications of an MRS Cluster Node?

### Question:

Can I Downgrade the Specifications of an MRS Cluster Node?

### Answer:

Currently, the configuration specifications of nodes in an MRS cluster cannot be downgraded.

## 1.6 Are Hive Components of Different Versions Compatible with Each Other?

Hive 3.1 is incompatible with Hive 1.2 in the following items:

- Strings cannot be converted to int in Hive 3.1.
- The user-defined functions (UDFs) of the **Date** type are changed to Hive 3.1 built-in UDFs.
- Hive 3.1 does not provide the index function anymore.
- The JDBC drivers in Hive 3.1 and Hive 1.2 are incompatible.
- In Hive 3.1, column names in ORC files are case-sensitive and underscores-sensitive.
- Hive 3.1 does not allow columns named **time**.

## 1.7 What Are the Differences Between OBS and HDFS in Data Storage?

The data processed by MRS is from OBS or HDFS. OBS is an object-based storage service that provides secure, reliable, and cost-effective storage of huge amounts of data. MRS can directly process data in OBS. You can view, manage, and use data by using the OBS console or OBS client. In addition, you can use REST APIs independently or integrate APIs to service applications to manage and access data.

- Data stored in OBS: Data storage is decoupled from compute. The cluster storage cost is low, and storage capacity is not limited. Clusters can be deleted at any time. However, the computing performance depends on the OBS access performance and is lower than that of HDFS. OBS is recommended for applications that do not demand a lot of computation.
- Data stored in HDFS: Data storage is not decoupled from compute. The cluster storage cost is high, and storage capacity is limited. The computing performance is high. You must export data before you delete clusters. HDFS is recommended for computing-intensive scenarios.

## 1.8 What Are the Solutions for Processing 1 Billion Data Records?

### Question:

What Are the Solutions for processing 1 Billion Data Records?

### Answer:

- GaussDB (for MySQL) is recommended for scenarios, such as data updates, online transaction processing (OLTP), and complex analysis of 1 billion data records.
- Impala and Kudu in MRS also meet this requirement. Impala and Kudu can load all join tables to the memory in the join operation.

## 1.9 What are the advantages of the compression ratio of zstd?

### Question:

What are the advantages of the compression ratio of zstd?

### Answer:

The compression ratio of zstd is twice that of orc. zstd is open-source.

For details, see <https://github.com/L-Angel/compress-demo>.

CarbonData does not support lzo. MRS integrates zstd.

# 2 Billing

---

## 2.1 Why Is the Price Not Displayed During MRS Cluster Creation?

### Question:

Why Is the Price Not Displayed During MRS Cluster Creation?

### Answer:

If you select only the number of disks but not the number of cluster VM instances, the product quotation will not be charged.

The product quotation is displayed in the lower part of the browser only after the customer selects the number of instances.

## 2.2 How Is Auto Scaling Billed for an MRS Cluster?

### Question:

How Is Auto Scaling Billed for an MRS Cluster?

### Answer:

The price displayed at the bottom when you purchase an MRS cluster will not include the auto scaling fee if you specify only the auto scaling range for the Task nodes without configuring the instance count. If the number of instances is configured, the product quotation is displayed in the lower part of the browser.

If you add nodes through the auto scaling function, the added nodes will be billed by the actual usage duration per hour regardless of whether the cluster's billing mode is yearly/monthly or pay-per-use mode.

## 2.3 How Are Task Nodes in an MRS Cluster Billed?

### Question:

How Are Task Nodes in an MRS Cluster Billed?

### Answer:

Task nodes in a cluster are billed in pay-per-use mode, regardless of the cluster billing mode. That is, they are billed per hour based on the actual usage duration.

## 2.4 Why Does My Unsubscription from ECS Fail After I Unsubscribe from MRS?

### Question:

Why Does My Unsubscription from ECS Fail After I Unsubscribe from MRS?

### Answer:

1. Check that the ECS ID is not used by an MRS cluster in use.
2. On the ECS console, find the ECS to be unsubscribed from and click **Locked by MRS** to unlock it.
3. Click **Unsubscribe** again.
4. If the unsubscription still fails, collect the ECS ID and contact Huawei Cloud technical support.

# 3 Cluster Creation

---

## 3.1 How Do I Create an MRS Cluster Using a Custom Security Group?

### Question:

How Do I Create an MRS Cluster Using a Custom Security Group?

### Answer:

You can create an MRS cluster using a user-defined security group in either of the following ways:

- If you want to use a security group created by yourself, you need to enable port 9022 when you buy a cluster.
- When you buy a cluster, set **Security Group** to **Auto create**.

## 3.2 What Should I Do If HDFS, Yarn, and MapReduce Components Are Unavailable When I Buy an MRS Cluster?

### Question:

What Should I Do If HDFS, Yarn, and MapReduce Components Are Unavailable When I Buy an MRS Cluster?

### Answer:

The HDFS, YARN, and MapReduce components are integrated in Hadoop. If the three components are unavailable when an MRS cluster is bought, select Hadoop instead. After an MRS cluster is created, HDFS, YARN, and MapReduce are available in the **Components** page.



### 3.3 What Should I Do If the ZooKeeper Component Is Unavailable When I Buy an MRS Cluster?

#### Question:

Why Cannot I Find the ZooKeeper Component When I Buy an MRS Cluster?

#### Answer:

- If you create a cluster of a version earlier than MRS 3.x, ZooKeeper is installed by default and is not displayed on the GUI.  
After the cluster is created, the ZooKeeper component is available on the **Components** page.
- If you create a cluster of MRS 3.x or later, ZooKeeper is available on the GUI and is selected by default.

### 3.4 What Should I Do If Invalid Authentication Is Reported When I Submit an Order for Purchasing an MRS Cluster?

#### Question:

What Should I Do If Invalid Authentication Is Reported When I Submit an Order for Purchasing an MRS Cluster?

#### Answer:

1. Contact the administrator to check API request logs. The alarm information is as follows: IAM users have not been assigned the **mrs:cluster:create** permission in the fine-grained policy.
2. The IAM user belongs to multiple user groups, and different default MRS policies are assigned to these user groups. Low-permission policies are preferentially matched. This policy does not contain the **mrs:cluster:create** permission. As a result, the cluster creation operation cannot be submitted, and error 401 is reported.
3. After the user is removed from the user group corresponding to the low-permission policy, the cluster purchase order can be submitted.

# 4 Web Page Access

---

## 4.1 How Do I Change the Session Timeout Duration for an Open Source Component Web UI?

### Question:


How Do I Change the Session Timeout Interval on the Web Page of an Open-Source Component?

### Answer:

A proper web page timeout interval prevents information leakage caused by long-time web page exposure.

**Step 1** Check whether the session timeout duration can be adjusted for the cluster.

- For MRS cluster versions earlier than 3.x:
  - a. On the cluster details page, choose **Components > meta > Service Configuration**.
  - b. Switch **Basic** to **All**, and search for the **http.server.session.timeout.secs**. If **http.server.session.timeout.secs** does not exist, the cluster does not support change of the session timeout duration. If the parameter exists, perform the following steps to modify it.
- MRS 3.x and later: Log in to FusionInsight Manager and choose **Cluster > Services > meta**. On the displayed page, click **Configurations** and select **All Configurations**. Search for the **http.server.session.timeout.secs** configuration item. If this configuration item exists, perform the following steps to modify it. If the configuration item does not exist, the version does not support dynamic adjustment of the session duration.



You are advised to set all session timeout durations to the same value. Otherwise, the settings of some parameters may not take effect due to value conflict.

## Step 2 Modify the timeout duration on Manager and the authentication center page.

- **For clusters of versions earlier than MRS 3.x**
  - a. Log in to each master node in the cluster and perform [Step 2.b](#) to [Step 2.d](#).
  - b. Change the value of `<session-timeout>20</session-timeout>` in the `/opt/Bigdata/apache-tomcat-7.0.78/webapps/cas/WEB-INF/web.xml` file. `<session-timeout>20</session-timeout>` indicates the session timeout duration, in minutes. Change it based on service requirements. The maximum value is 480 minutes.
  - c. Change the value of `<session-timeout>20</session-timeout>` in the `/opt/Bigdata/apache-tomcat-7.0.78/webapps/web/WEB-INF/web.xml` file. `<session-timeout>20</session-timeout>` indicates the session timeout duration, in minutes. Change it based on service requirements. The maximum value is 480 minutes.
  - d. Change the values of `p:maxTimeToLiveInSeconds="$ {tgt.maxTimeToLiveInSeconds:1200}"` and `p:timeToKillInSeconds="$ {tgt.timeToKillInSeconds:1200}"` in the `/opt/Bigdata/apache-tomcat-7.0.78/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml` file. The maximum value is 28,800 seconds.
  - e. Restart the Tomcat node on the active master node.
    - i. On the active master node, run the `netstat -anp |grep 28443 |grep LISTEN | awk '{print $7}'` command as user `omm` to query the Tomcat process ID.
    - ii. Run the `kill -9 {pid}` command, in which `{pid}` indicates the Tomcat process ID obtained in [Step 2.e.i](#).
    - iii. Wait until the process automatically restarts. You can run the `netstat -anp |grep 28443 |grep LISTEN` command to check whether the process is successfully restarted. If the process can be queried, the process is successfully restarted. If the process cannot be queried, query the process again later.
- **For clusters of MRS 3.x or later**
  - a. Log in to each master node in the cluster and perform [Step 2.b](#) to [Step 2.c](#) on each master node.
  - b. Change the value of `<session-timeout>20</session-timeout>` in the `/opt/Bigdata/om-server_xxx/apache-tomcat-xxx/webapps/web/WEB-INF/web.xml` file. `<session-timeout>20</session-timeout>` indicates the session timeout duration, in minutes. Change it based on service requirements. The maximum value is 480 minutes.

- c. Add **ticket.tgt.timeToKillInSeconds=28800** to the **/opt/Bigdata/om-server\_XXX/apache-tomcat-8.5.63/webapps/cas/WEB-INF/classes/config/application.properties** file. **ticket.tgt.timeToKillInSeconds** indicates the validity period of the authentication center, in seconds. Change it based on service requirements. The maximum value is 28,800 seconds.
- d. Restart the Tomcat node on the active master node.
  - i. On the active master node, run the **netstat -anp |grep 28443 |grep LISTEN | awk '{print \$7}'** command as user **omm** to query the Tomcat process ID.
  - ii. Run the **kill -9 {pid}** command, in which **{pid}** indicates the Tomcat process ID obtained in [Step 2.d.i](#).
  - iii. Wait until the process automatically restarts.

You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is successfully restarted. If the process is displayed, the process is successfully restarted. If the process is not displayed, query the process again later.

**Step 3** Modify the timeout duration for an open-source component web UI.

1. Access the **All Configurations** page.
  - For MRS cluster versions earlier than MRS 3.x:

On the cluster details page, choose **Components > Meta > Service Configuration**.
  - For MRS cluster version 3.x or later:

Log in to FusionInsight Manager and choose **Cluster > Services > meta**. On the displayed page, click **Configurations** and select **All Configurations**.
2. Change the value of **http.server.session.timeout.secs** under **meta** as required. The unit is second.
3. Save the settings, deselect **Restart the affected services or instances**, and click **OK**.

You are advised to perform the restart during off-peak hours.
4. (Optional) If you need to use the Spark web UI, search for **spark.session.maxAge** on the **All Configurations** page of Spark and change the value (unit: second).

Save the settings, deselect **Restart the affected services or instances**, and click **OK**.
5. Restart the meta service and components on web UI, or restart the cluster during off-peak hours.

To prevent service interruption, restart the service during off-peak hours or perform a rolling restart. For details, see [Performing Rolling Restart](#).

----End

 **NOTE**

The timeout duration of the Flink component web UI cannot be changed.

## 4.2 What Can I Do If the Dynamic Resource Plan Page in MRS Tenant Management Cannot Be Refreshed?

### Question:

What Can I Do If the Dynamic Resource Plan Page in MRS Tenant Management Cannot Be Refreshed?

### Answer:

- Step 1** Log in to the Master1 and Master2 nodes as user **root**.
  - Step 2** Run the **ps -ef |grep aos** command to check the AOS process ID.
  - Step 3** Run the **kill -9 AOS process ID** command to end the AOS process.
  - Step 4** Wait until the AOS process is automatically restarted. You can run the **ps -ef |grep aos** command to check whether the AOS process exists. If the process exists, the restart is successful. If the process does not exist, retry later.
- End

## 4.3 What Do I Do If the Kafka Topic Monitoring Tab Is Not Displayed on Manager?

### Question:

What Do I Do If the Kafka Topic Monitoring Tab Is Not Displayed on Manager?

### Answer:

- Step 1** Log in to the Master nodes of the cluster and switch to user **omm**.
  - Step 2** Go to the **/opt/Bigdata/apache-tomcat-7.0.78/webapps/web/WEB-INF/lib/components/Kafka/** directory.  

```
cd /opt/Bigdata/apache-tomcat-7.0.78/webapps/web/WEB-INF/lib/components/Kafka/
```
  - Step 3** Copy the ZooKeeper package to the directory.  

```
cp /opt/share/zookeeper-3.5.1-mrs-2.0/zookeeper-3.5.1-mrs-2.0.jar ./
```
  - Step 4** Restart Tomcat.  

```
sh /opt/Bigdata/apache-tomcat-7.0.78/bin/shutdown.sh  
sh /opt/Bigdata/apache-tomcat-7.0.78/bin/startup.sh
```
- End

## 4.4 What CAN I DO IF an Error Is Reported or Some Functions Are Unavailable When I Access the Web UIs of Components such as HDFS, Hue, Yarn, Flink, and HetuEngine?

### Question:


What CAN I DO IF an Error Is Reported or Some Functions Are Unavailable When I Access the Web UIs of Components such as HDFS, Hue, Yarn, Flink, and HetuEngine?

### Answer:


Users who access the WebUI of components such as HDFS, Hue, Yarn, Flink, and HetuEngine do not have the management permission of the corresponding components. As a result, an error is reported or some functions are unavailable.

Examples:

- After you log in to the web UI of Flink as the current user, some content cannot be displayed, and you do not have the permission to create applications, cluster connections, or data connections.



- An error is reported when the user accesses the HDFS web UI. Failed to retrieve data from /jmx?qry=java.lang:type=Memory, cause: Forbidden



- After you log in to the web UI of YARN as the current user, you cannot view job information.



**SUBMITTED Applications**

Cluster Metrics

| Apps Submitted | Apps Pending | Apps Running | Apps Completed | Containers Running | Used Resources         | Total Resources           | Reserved Resources     |
|----------------|--------------|--------------|----------------|--------------------|------------------------|---------------------------|------------------------|
| 0              | 0            | 0            | 0              | 0                  | <memory>0 B, vCores<0> | <memory>96 GB, vCores<96> | <memory>0 B, vCores<0> |

Cluster Node Metrics

| Active Nodes | Decommissioning Nodes | Decommissioned Nodes | Lost Nodes | Unhealthy Nodes | Rebooted Nodes | Shutdown Nodes |
|--------------|-----------------------|----------------------|------------|-----------------|----------------|----------------|
| 0            | 0                     | 0                    | 0          | 0               | 0              | 0              |

User Metrics for admin

| Apps Submitted | Apps Pending | Apps Running | Apps Completed | Containers Running | Containers Pending | Containers Reserved | Memory Used | Memory Pending | Memory Reserved | VCores Used | VCores Pending | VCores Reserved |
|----------------|--------------|--------------|----------------|--------------------|--------------------|---------------------|-------------|----------------|-----------------|-------------|----------------|-----------------|
| 0              | 0            | 0            | 0              | 0                  | 0                  | 0                   | 0 B         | 0 B            | 0 B             | 0           | 0              | 0               |

Scheduler Metrics


| Scheduler Type     | Scheduling Resource Type                   | Minimum Allocation     | Maximum Allocation                                            | Maximum Cluster Application Priority |
|--------------------|--------------------------------------------|------------------------|---------------------------------------------------------------|--------------------------------------|
| SupernodeScheduler | hadoop.io/gpu, memory-mb (unit-MB), vcores | <memory>512, vCores<1> | <memory>65536, vCores<32, hadoop.io/gpu: 9223372036854775807> | 5                                    |

Show 20 entries

| ID                         | User | Name | Application Type | Queue | Application Priority | StartTime | FinishTime | State | FinalStatus | Running Containers | Allocated CPU | Allocated Memory | Allocated GPU | Reserved CPU | Reserved Memory | Reserved GPU | % of Queue | % of Cluster | Progress | Tracking UI | Blacklisted Nodes |
|----------------------------|------|------|------------------|-------|----------------------|-----------|------------|-------|-------------|--------------------|---------------|------------------|---------------|--------------|-----------------|--------------|------------|--------------|----------|-------------|-------------------|
| No data available in table |      |      |                  |       |                      |           |            |       |             |                    |               |                  |               |              |                 |              |            |              |          |             |                   |

Showing 0 to 0 of 0 entries

- After you log in to the web UI of Hue as the current user, click  in the navigation pane on the left, and select **Workflow**, an error message is displayed.



You are advised to log in to the web UIs of the components as a user with corresponding management permissions. For example, you can create a service user who has the management permissions on HDFS and you can log in to the web UI of HDFS as the created user. For details, see [Creating a User](#).

## 4.5 How Do I Switch the Methods to Access MRS Manager?

### Question:

How Do I Access MRS Manager Through an EIP After Accessing It Through a Direct Connect Connection in an MRS 3.x Cluster?

### Answer:

On the **Dashboard** tab page of the cluster, click  next to **Access Manager** to switch the access mode.

## 4.6 Why Cannot I Find the User Management Page on MRS Manager?

### Question:

Why Cannot I Find the User Management Page on MRS Manager?

**Answer:**

The current login user does not have the **Manager\_administrator** role permission. Log in to MRS Manager again as user **admin** or other users who have the Manager management permission.

## 4.7 What Can I Do If the Excel File Downloaded by Hue Cannot Be Opened?

**NOTE**

This section applies only to versions earlier than MRS 3.x.

1. Log in to a Master node as user **root** and switch to user **omm**.  
**su - omm**

2. Check whether the current node is the active OMS node.

```
sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh
```

If **active** is displayed in the command output, the node is the active node. Otherwise, log in to the other Master node.

Figure 4-1 Active OMS node

| nodeName       | HostName       | HAVersion      | StartTime           | HAActive       | HAAllResOK | HARunPhase  |
|----------------|----------------|----------------|---------------------|----------------|------------|-------------|
| 192-168-67-136 |                |                | 2022-09-27 20:23:41 | active         | normal     | Activated   |
| 192-168-67-142 |                |                | 2022-09-27 20:24:39 | standby        | normal     | Deactivated |
| nodeName       | ResName        | ResStatus      | ResHStatus          | ResType        |            |             |
| 192-168-67-136 | acs            | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | aos            | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | cep            | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | controller     | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | disaster       | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | floatip        | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | fms            | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | gaussDB        | Active_normal  | Normal              | Active_standby |            |             |
| 192-168-67-136 | heartbeatcheck | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | httpd          | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | iam            | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | ntp            | Active_normal  | Normal              | Active_standby |            |             |
| 192-168-67-136 | okerberos      | Active_normal  | Normal              | Active_standby |            |             |
| 192-168-67-136 | oldap          | Active_normal  | Normal              | Active_standby |            |             |
| 192-168-67-136 | oms            | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-136 | tomcat         | Normal         | Normal              | Single_active  |            |             |
| 192-168-67-142 | acs            | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | aos            | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | cep            | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | controller     | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | disaster       | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | floatip        | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | fms            | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | gaussDB        | Standby_normal | Normal              | Active_standby |            |             |
| 192-168-67-142 | heartbeatcheck | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | httpd          | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | iam            | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | ntp            | Standby_normal | Normal              | Active_standby |            |             |
| 192-168-67-142 | okerberos      | Standby_normal | Normal              | Active_standby |            |             |
| 192-168-67-142 | oldap          | Standby_normal | Normal              | Active_standby |            |             |
| 192-168-67-142 | oms            | Stopped        | Normal              | Single_active  |            |             |
| 192-168-67-142 | tomcat         | Stopped        | Normal              | Single_active  |            |             |

3. Go to the **{BIGDATA\_HOME}/Apache-httpd-\*/conf** directory.  
**cd \${BIGDATA\_HOME}/Apache-httpd-\*/conf**
4. Open the **httpd.conf** file.  
**vim httpd.conf**
5. Search for **21201** in the file and delete the following content from the file. (The values of *proxy\_ip* and *proxy\_port* are the same as those in the actual environment.)

**ProxyHTMLEnable On**

**SetEnv** PROXY\_PREFIX=https://[proxy\_ip]:[proxy\_port]

**ProxyHTMLURLMap** (https?://\[^\]\*:[0-9]\*.\*) \${PROXY\_PREFIX}/proxyRedirect=\$1 RV



Figure 4-2 Content to be deleted

```
494 <VirtualHost *:80>
495     ServerName https://192.168.0.175:8080
496     SSLProxyEngine On
497     ProxyRequests Off
498     TraceEnable Off
499     ProxyTimeout 1200
500     RewriteEngine On
501     ProxyHTMLEnable On
502     # LogLevel alert rewrite:trace6
503     RewriteMap proxylist dbm:/opt/bigdata/apache-httpd-2.4.26/conf/proxylist.dbm
504
505     SetEnv PROXY_PREFIX=https://192.168.0.175:20026
506     ProxyHTMLURLMap (https://\/*:[^:]*:[0-9]*.*) ${PROXY_PREFIX}/proxyRedirect=$1 RV
507
508     RewriteRule ^(\/*)$ ${proxylist:Hue}$1 [E=TARGET_PATH:$1,L,P]
509
510     Header edit Location "(?!(https://192.168.0.175:20009|https://192.168.0.175:8080|https://192.168.0.175:8080)https://\/*:[^:]*:[0-9]*.*)$ https://192.168.0.175:8080$1
511
512     ProxyPassReverseCookiePath / / interpolate
513
```

6. Save the modification and exit.
7. Open the **httpd.conf** file again, search for **proxy\_hue\_port**, and delete the following content:  
ProxyHTMLEnable On  
SetEnv PROXY\_PREFIX=https://[proxy\_ip]:[proxy\_port]  
ProxyHTMLURLMap (https://\/\*:[^:]\*:[0-9]\*.\*) \${PROXY\_PREFIX}/proxyRedirect=\$1 RV

Figure 4-3 Content to be deleted

```
494 <VirtualHost *:80>
495     ServerName https://[proxy_ip]:[proxy_hue_port]
496     SSLProxyEngine On
497     ProxyRequests Off
498     TraceEnable Off
499     ProxyTimeout 1200
500     RewriteEngine On
501     ProxyHTMLEnable On
502     # LogLevel alert rewrite:trace6
503     RewriteMap proxylist dbm:/httpd_home/conf/proxylist.dbm
504
505     SetEnv PROXY_PREFIX=https://[proxy_ip]:[proxy_port]
506     ProxyHTMLURLMap (https://\/*:[^:]*:[0-9]*.*) ${PROXY_PREFIX}/proxyRedirect=$1 RV
507
508     RewriteRule ^(\/*)$ ${proxylist:Hue}$1 [E=TARGET_PATH:$1,L,P]
509
510     Header edit Location "(?!(https://[cas_ip]:[cas_port]|https://[proxy_ip]:[proxy_hue_port]|https://\/*:[^:]*:[0-9]*.*)$ https://[proxy_ip]:[proxy_hue_port]$1
511
512     ProxyPassReverseCookiePath / / interpolate
513
```

8. Save the modification and exit.
9. Run the following command to restart the **httpd** process:  
**sh \${BIGDATA\_HOME}/Apache-httpd-\*/setup/restarthttpd.sh**
10. Check whether the **httpd.conf** file on the standby Master node is modified. If the file is modified, no further action is required. If the file is not modified, modify the **httpd.conf** file on the standby Master node in the same way. You do not need to restart the **httpd** process.
11. Download the Excel file again. You can open the file successfully.

# 5 Authentication and Permission

---

## 5.1 What Is the User for Logging in to FusionInsight Manager?

### Question:

What Is the User for Logging in to FusionInsight Manager?

### Answer:

The default account for logging in to Manager is **admin**, and the password is the one you set when you created the cluster.

## 5.2 How Do I Query and Change the Password Validity Period of a User In a Cluster?

### Querying the Password Validity Period

Querying the password validity period of a component running user (human-machine user or machine-machine user):

- Step 1** Log in to the node where the client is installed as the client installation user.
- Step 2** Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:  
**cd /opt/Bigdata/client**
- Step 3** Run the following command to configure environment variables:  
**source bigdata\_env**
- Step 4** Run the following command and enter the password of user **kadmin/admin** to log in to the kadmin console:  
**kadmin -p kadmin/admin**

 NOTE

The default password of user **kadmin/admin** is **Admin@123**. Change the password upon your first login or as prompted and keep the new password secure.

**Step 5** Run the following command to view the user information:

**getprinc** *Internal system username*

Example: **getprinc** user1

```
kadmin: getprinc user1
.....
Expiration date: [never]
Last password change: Sun Oct 09 15:29:54 CST 2022
Password expiration date: [never]
.....
```

----End

**Querying the password validity period of an OS user:**

**Step 1** Log in to any master node in the cluster as user **root**.

**Step 2** Run the following command to view the password validity period (value of **Password expires**):

**chage -l** *Username*

For example, to view the password validity period of user **root**, run the **chage -l root** command. The command output is as follows:

```
[root@xxx ~]# chage -l root
Last password change           : Sep 12, 2021
Password expires                : never
Password inactive              : never
Account expires                : never
Minimum number of days between password change : 0
Maximum number of days between password change  : 99999
Number of days of warning before password expires : 7
```

----End

## Changing the Password Validity Period

- The password of a machine-machine user is randomly generated and never expires by default.
- The password validity period of a human-machine user can be changed by modifying the password policy on Manager.

The operations are as follows:

Change the password validity period of a human-machine user by referring to [Modifying a Password Policy](#).

## 5.3 Does an MRS Cluster Support Access Permission Control If Kerberos Authentication Is not Enabled?

MRS clusters with versions earlier than 1.8.0 do not support access permission control if Kerberos authentication is not enabled. MRS clusters with Kerberos

authentication enabled support role based permission control. For MRS 1.8.0 or later, all clusters support role based permission control.

For MRS cluster 2.1.0 or earlier, choose **System > Configuration > Permission** on MRS Manager.

For MRS cluster 3.x or later, choose **System > Permission** on FusionInsight Manager.

For details about how to use clusters with Kerberos authentication enabled, see [Clusters with Kerberos Authentication Enabled](#).

## 5.4 How Do I Add Tenant Management Permission to Users in a Cluster?

You can assign tenant management permission only in analysis or hybrid clusters, but not in streaming clusters.

The operations vary depending on the MRS cluster version:

### Procedure for versions earlier than MRS cluster 3.x:

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

**Step 2** Choose **System > Manage User**. Select the new user, and click **Modify** in the **Operation** column.

**Step 3** In **Assign Rights by Role**, click **Select and Add Role**.

- If you bind the **Manager\_tenant** role to the account, the account will have permission to view tenant management information.
- If you bind the **Manager\_administrator** role to the account, the account will have permission to view and perform tenant management.

**Step 4** Click **OK**.

----End

### Procedure for MRS cluster 3.x and later versions:

**Step 1** Log in to FusionInsight Manager and choose **System > Permission > User**.

**Step 2** Locate the user and click **Modify**.

Modify the parameters based on service requirements.

If you bind the **Manager\_tenant** role to the account, the account will have permission to view tenant management information. If you bind the **Manager\_administrator** role to the account, the account will have permission to view and perform tenant management.

### NOTE

It takes about three minutes for the settings to take effect after user group or role permission are modified.

**Step 3** Click **OK**.

----End

## 5.5 Does Hue Provide the Function of Configuring Account Permissions?

### Question:

Does Hue Provide the Function of Configuring Account Permissions?

### Answer:

The Hue service does not provide the function of configuring account permissions.

You can configure account permissions related to Hue by configuring user roles and user groups in **System Settings** on FusionInsight Manager.

## 5.6 Why Can't I Submit Jobs on the Console After My IAM Account Is Assigned with MRS Permissions?

### Question:

Why can't I submit jobs on the MRS console after my IAM account is assigned with the **MRS ReadOnlyAccess** and **MRS FullAccess** permissions?

### Answer:

The account has the **MRS ReadOnlyAccess** and **MRS FullAccess** permissions. Due to the permission priority, the account cannot add jobs on the MRS console.

The group to which the IAM account belongs has the **MRS FullAccess**, **MRS ReadOnlyAccess**, and **MRS Administrator** permissions. **MRS FullAccess** and **MRS ReadOnlyAccess** are fine-grained permissions, and **MRS Administrator** is a role-based access control (RBAC) policy. Fine-grained permissions have higher priorities over RBAC policies. If fine-grained permissions and RBAC policies are configured, fine-grained permissions take effect first. Fine-grained permissions follow the deny priority principle. Therefore, the **MRS ReadOnlyAccess** permission takes effect finally. So, a message is displayed, indicating that the account does not have the permission.

Delete the **MRS ReadOnlyAccess** permission, log out of the console, and log in to the console again.

## 5.7 How Do I View the Hive Table Created by Another User?

Versions earlier than MRS 3.x:

1. Log in to MRS Manager and choose **System > Permission > Manage Role**.
2. Click **Create Role**, and set **Role Name** and **Description**.

3. In the **Permission** table, choose **Hive > Hive Read Write Privileges**.
4. In the database list, click the name of the database where the table created by user B is stored. The table is displayed.
5. In the **Permission** column of the table created by user B, select **SELECT**.
6. Click **OK**, and return to the **Role** page.
7. Choose **System > Manage User**. Locate the row containing user A, click **Modify** to bind the new role to user A, and click **OK**. After about 5 minutes, user A can access the table created by user B.

MRS 3.x or later:

1. Log in to FusionInsight Manager and choose **Cluster > Services**. On the page that is displayed, choose **Hive**. On the displayed page, choose **More**, and check whether **Enable Ranger** is grayed out.
  - If yes, go to [9](#).
  - If no, perform [2](#) to [8](#).
2. Log in to FusionInsight Manager and choose **System > Permission > Role**.
3. Click **Create Role**, and set **Role Name** and **Description**.
4. In the **Configure Resource Permission** table, choose *Name of the desired cluster* > **Hive > Hive Read Write Privileges**.
5. In the database list, click the name of the database where the table created by user B is stored. The table is displayed.
6. In the **Permission** column of the table created by user B, select **Select**.
7. Click **OK**, and return to the **Role** page.
8. Choose **Permission > User**. On the **Local User** page that is displayed, locate the row containing user A, click **Modify** in the **Operation** column to bind the new role to user A, and click **OK**. After about 5 minutes, user A can access the table created by user B.
9. Perform the following steps to add the Ranger access permission policy of Hive:
  - a. Log in to FusionInsight Manager as a Hive administrator and choose **Cluster > Services**. On the page that is displayed, choose **Ranger**. On the displayed page, click the URL next to **Ranger WebUI** to go to the Ranger management page.
  - b. On the home page, click the component plug-in name in the **HADOOP SQL** area, for example, **Hive**.
  - c. On the **Access** tab page, click **Add New Policy** to add a Hive permission control policy.
  - d. In the **Create Policy** dialog box that is displayed, set the following parameters:
    - **Policy Name**: Enter a policy name, for example, **table\_test\_hive**.
    - **database**: Enter or select the database where the table created by user B is stored, for example, **default**.
    - **table**: Enter or select the table created by user B, for example, **test**.

- **column:** Enter and select a column, for example, `*`.
  - In the **Allow Conditions** area, click **Select User**, select user A, click **Add Permissions**, and select **select**.
  - Click **Add**.
10. Perform the following steps to add the Ranger access permission policy of HDFS:
- a. Log in to FusionInsight Manager as user **rangeradmin** and choose **Cluster > Services**. On the page that is displayed, choose **Ranger**. On the displayed page, click the URL next to **Ranger WebUI** to go to the Ranger management page.
  - b. On the home page, click the component plug-in name in the **HDFS** area, for example, **hacluster**.
  - c. Click **Add New Policy** to add an HDFS permission control policy.
  - d. In the **Create Policy** dialog box that is displayed, set the following parameters:
    - **Policy Name:** Enter a policy name, for example, **tablehdfs\_test**.
    - **Resource Path:** Set this parameter to the HDFS path where the table created by user B is stored, for example, **/user/hive/warehouse/Database name/ Table name**.
    - In the **Allow Conditions** area, select user A for **Select User**, click **Add Permissions** in the **Permissions** column, and select **Read** and **Execute**.
    - Click **Add**.
11. View basic information about the policy in the policy list. After the policy takes effect, user A can view the table created by user B.

## 5.8 How Do I Prevent Kerberos Authentication Expiration?

- Java applications:

Before connecting to HBase, HDFS, or other big data components, call `loginUserFromKeytab()` to create a UGI. Then, start a scheduled thread to periodically check whether the Kerberos Authentication expires. Log in to the system again before the Kerberos Authentication expires.

```
private static void startCheckKeytabTgtAndReloginJob() {  
    //The credential is checked every 10 minutes, and updated before the expiration time.  
    ThreadPool.updateConfigThread.scheduleWithFixedDelay(() -> {  
        try {  
            UserGroupInformation.getLoginUser().checkTGTAndReloginFromKeytab();  
            logger.warn("get tgt:{", UserGroupInformation.getLoginUser().getTGT());  
            logger.warn("Check Kerberos Tgt And Relogin From Keytab Finish.");  
        } catch (IOException e) {  
            logger.error("Check Kerberos Tgt And Relogin From Keytab Error", e);  
        }  
    }, 0, 10, TimeUnit.MINUTES);  
    logger.warn("Start Check Keytab TGT And Relogin Job Success.");  
}
```

- Tasks executed in shell client:
  - a. Run the **kinit** command to authenticate the user.
  - b. Create a scheduled task of the operating system or any other scheduled task to run the **kinit** command to authenticate the user periodically.
  - c. Submit jobs to execute big data tasks.
- Spark jobs:

If you submit jobs using **spark-shell**, **spark-submit**, or **spark-sql**, you can specify **Keytab** and **Principal** in the command to perform authentication and periodically update the login credential and authorization tokens to prevent authentication expiration.

The following is an example:

```
spark-shell --principal spark2x/hadoop.<System domain name>@<System domain name> --keytab ${BIGDATA_HOME}/FusionInsight_Spark2x_XXX/install/FusionInsight-Spark2x-2.4.5/keytab/spark2x/SparkResource/spark2x.keytab --master yarn
```

## 5.9 How Do I Enable or Disable Kerberos Authentication for an Existing MRS Cluster?

### Question

How Do I Enable or Disable Kerberos Authentication for an Existing MRS Cluster?

### Answer

You cannot enable or disable the Kerberos service after the MRS cluster is created. You need to create an MRS cluster, enable or disable the Kerberos service, and migrate data from the original one.

## 5.10 What Are the Ports of the Kerberos Authentication Service?

### Question

What Are the Ports of the Kerberos Authentication Service?

### Answer

The Kerberos authentication service uses ports 21730 (TCP), 21731 (TCP/UDP), and 21732 (TCP/UDP).



# 6 Client Usage

---

## 6.1 How Do I Disable SASL Authentication for ZooKeeper?

### Question:

How Do I Disable ZooKeeper SASL Authentication?

### Answer:

1. Log in to FusionInsight Manager.
2. Choose Cluster > Service > ZooKeeper > Configuration > All Configurations.
3. In the navigation tree on the left, choose quorumpeer > Customization and add zookeeper.sasl.disable = false.  
Save the configuration and restart the ZooKeeper service.

## 6.2 What Can I Do If the Error Message "Permission denied" Is Displayed When kinit Is Executed on a Client Outside the MRS Cluster?

### Symptom

After the client is installed on a node outside an MRS cluster and the **kinit** command is executed, the following error information is displayed:

```
-bash: kinit: Permission denied
```

The following error information is displayed when the **java** command is executed:

```
-bash: /xxx/java: Permission denied
```

After running the **ll /Java installation path/JDK/jdk/bin/java** command, it is found that the file execution permission is correct.

## Fault Locating

Run the **mount | column -t** command to check the status of the mounted partition. It is found that the partition status of the mount point where the Java execution file is located is **noexec**. In the current environment, the data disk where the MRS client is installed is set to **noexec**, that is, binary file execution is prohibited. As a result, Java commands cannot be executed.

## Solution

1. Log in to the node where the MRS client is located as user **root**.
2. Remove the configuration item **noexec** of the data disk where the MRS client is located from the **/etc/fstab** file.
3. Run the **umount** command to detach the data disk, and then run the **mount -a** command to remount the data disk.

## 6.3 What Should I Do If an Alarm Is Reported Indicating that the Memory Is Insufficient When I Execute a SQL Statement on the ClickHouse Client?

### Symptom

The ClickHouse client restricts the memory used by GROUP BY statements. When a SQL statement is executed on the ClickHouse client, the following error information is displayed:

```
Progress: 1.83 billion rows, 85.31 GB (68.80 million rows/s., 3.21 GB/s.)      6%Received exception from server:
Code: 241. DB::Exception: Received from localhost:9000, 127.0.0.1.
DB::Exception: Memory limit (for query) exceeded: would use 9.31 GiB (attempt to allocate chunk of 1048576 bytes), maximum: 9.31 GiB:
(while reading column hits):
```

### Solution

- Run the following command before executing an SQL statement on condition that the cluster has sufficient memory:  
`SET max_memory_usage = 128000000000; #128G`
- If no sufficient memory is available, ClickHouse enables you to overflow data to disk to free up the memory: You are advised to set the value of **max\_memory\_usage** to twice the size of **max\_bytes\_before\_external\_group\_by**.  
`set max_bytes_before_external_group_by=20000000000; #20G`  
`set max_memory_usage=40000000000; #40G`
- If the customer has a large amount of data and the entire table is queried, you are advised to query the data by partition or upgrade the specifications of the core node in the cluster.

## 6.4 How Do I Connect to Spark Shell from MRS?

1. Log in to the Master node in the cluster as user **root**.

2. Run the following command to configure environment variables:  
**source** *Client installation directory*/**bigdata\_env**
3. If Kerberos authentication is enabled for the cluster, authenticate the user. If Kerberos authentication is disabled, skip this step.  
Run the **kinit** *MRS cluster user* command.  
Examples:
  - For a machine-machine user, run the **kinit -kt** *user.keytab sparkuser* command.
  - For a human-machine user, run the **kinit** *sparkuser* command.
4. Run the following command to connect to Spark shell:  
**spark-shell**

## 6.5 How Do I Connect to Spark Beeline from MRS?

1. Log in to the master node in the cluster as user **root**.
2. Run the following command to configure environment variables:  
**source** *Client installation directory*/**bigdata\_env**
3. If Kerberos authentication is enabled for the cluster, authenticate the user. If Kerberos authentication is disabled, skip this step.  
Run the **kinit** *MRS cluster user* command.  
Examples:
  - For a machine-machine user, run the **kinit -kt** *user.keytab sparkuser* command.
  - For a human-machine user, run the **kinit** *sparkuser* command.
4. Run the following command to connect to Spark Beeline:  
**spark-beeline**
5. Run commands on Spark Beeline. For example, create the table **test** in the **obs://mrs-word001/table/** directory.  
**create table test(id int) location 'obs://mrs-word001/table/';**
6. Query all tables.  
**show tables;**  
If the table **test** is displayed in the command output, OBS is successfully accessed.

Figure 6-1 Returned table name

```
0: jdbc:hive2://ha-cluster/default> create table test(id int) location 'obs://mrs-word001/table/';
+-----+
| Result |
+-----+
+-----+
No rows selected (2.515 seconds)
0: jdbc:hive2://ha-cluster/default> show tables;
+-----+
| database | tableName | isTemporary |
+-----+
| default  | test      | false       |
| default  | test_obs  | false       |
+-----+
2 rows selected (0.127 seconds)
```

7. Press **Ctrl+C** to exit the Spark Beeline.

## 6.6 What Should I Do If the Connection to the ClickHouse Server Fails and Error Code 516 Is Reported?

### Symptom

When the **clickhouse client** command is executed to connect to the ClickHouse server, the following error is reported:

```
ClickHouse exception, code: 516, host: 192.168.0.198, port: 8443; Code: 516, e.displayText() = DB::Exception: clickDeveloper: Authentication failed: password is incorrect or there is no user with such name
```

### Fault Locating

The username or password used for connecting the ClickHouse server is incorrect.

### Procedure

Enter the correct username or password . For details, see [Using ClickHouse from Scratch](#).

# 7 Component Configurations

---

## 7.1 Does MRS Support Running Hive on Kudu?

### Question:

Does MRS Support Running Hive on Kudu?

### Answer:

MRS does not support Hive on Kudu.

Currently, MRS supports only the following two methods to access Kudu:

- Access Kudu through Impala tables.
- Access and operate Kudu tables using the client application.

## 7.2 Does an MRS Cluster Support Hive on Spark?

### Question:

Does an MRS Cluster Support Hive on Spark?

### Answer:

- Clusters of MRS 1.9.x support Hive on Spark.
- Clusters of MRS 3.x or later support Hive on Spark.
- You can use Hive on Tez for the clusters of other versions.

## 7.3 Can I Change the IP address of DBService?

### Question:

Can I Change the IP address of DBService?

**Answer:**

The IP address of DBService cannot be changed in the MRS cluster.

## 7.4 What Access Protocols Does Kafka Support?

**Question:**

What Access Protocols Does Kafka Support?

**Answer:**

Kafka supports the following access protocols: PLAINTEXT, SSL, SASL\_PLAINTEXT, and SASL\_SSL.

## 7.5 What Python Versions Are Supported by Spark Tasks in an MRS Cluster?

**Question:**

Which Python Versions Are Supported by Spark Tasks in MRS 3.1.0 Clusters?

**Answer:**

For MRS 3.1.0 clusters, Python 2.7 or 3.x is recommended for Spark tasks.

## 7.6 What Are the Restrictions on the Storm Log Size in an MRS 2.1.0 Cluster?

**Question:**

What Are the Restrictions on the Storm Log Size in an MRS 2.1.0 Cluster?

**Answer:**

In MRS 2.1.0, the size of Storm logs cannot exceed 20 GB. If the size of Storm logs exceeds 20 GB, the logs will be deleted cyclically.

Logs are stored on the system disk, and the space is limited. If logs need to be stored for a long time, mount the logs to a data disk.

## 7.7 How Do I Modify the HDFS fs.defaultFS of an Existing Cluster?

**Question:** How do I modify the HDFS NameSpace (fs.defaultFS) of an existing cluster?

**Answer:** You can modify or add the HDFS NameSpace (fs.defaultFS) of the cluster by modifying the parameters of the **core-site.xml** and **hdfs-site.xml** files on the client. However, you are not advised to perform this operation on the server.

## 7.8 Can MRS Run Multiple Flume Tasks at a Time?

The Flume client supports multiple independent data flows. You can configure and link multiple sources, channels, and sinks in the **properties.properties** configuration file. These components can be linked to form multiple flows.

The following is an example of configuring two data flows in a configuration file:

```
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
```

## 7.9 How Do I Change FlumeClient Logs to Standard Logs?

1. Log in to the node where the Flume client is installed.
2. Go to the Flume client installation directory, for example, **/opt/FlumeClient**. Run the following command:  
**cd /opt/FlumeClient/fusioninsight-flume-1.9.0/bin**
3. Run the **./flume-manage.sh stop force** command to stop FlumeClient.
4. Run the **vi ../conf/log4j.properties** command to open the **log4j.properties** file and change the **flume.root.logger** value to **\${flume.log.level},console**.
5. Run the **./flume-manage.sh start force** command to restart FlumeClient.
6. After the modification, check whether the Docker configuration is correct.

## 7.10 Where Are the JAR Files and Environment Variables of Hadoop Stored?

- **hadoopstreaming.jar**: **/opt/share/hadoop-streaming-\*** (\* indicates the Hadoop version.)
- JDK environment variables: **/opt/client/JDK/component\_env**
- Hadoop environment variables: **/opt/client/HDFS/component\_env**
- Hadoop client: **/opt/client/HDFS/hadoop**

## 7.11 How Do I View HBase Logs?

1. Log in to the Master node in the cluster as user **root**.
2. Run the **su - omm** command to switch to user **omm**.
3. Run the **cd /var/log/Bigdata/hbase/** command to go to the **/var/log/Bigdata/hbase/** directory and view HBase logs.

## 7.12 How Do I Set the TTL for an HBase Table?

- Set the time to live (TTL) when creating a table:  
Create the **t\_task\_log** table, set the column family to **f**, and set the TTL to **86400** seconds.

```
create 't_task_log',{NAME=>'f', TTL=>'86400'}
```

- Set the TTL for an existing table:  

```
alter "t_task_log",NAME=>'data',TTL=>'86400' # Set the TTL value for the column family data.
```

## 7.13 How Do I Change the Number of HDFS Replicas?

1. Go to the HDFS service configuration page.
  - For MRS 1.8.10 or earlier clusters  
Log in to MRS Manager (see [Accessing MRS Manager](#)), choose **Services > HDFS > Service Configuration**, and select **All** from the **Basic** drop-down list.
  - For MRS 1.8.10 or later and MRS 2.x, click the cluster name. On the cluster details page that is displayed, choose **Components > HDFS > 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 cluster version 3.x or later:  
Log in to FusionInsight Manager, and choose **Cluster > Name of the desired cluster > Services > HDFS > Configurations > All Configurations**.
2. Search for **dfs.replication**, change the value (value range: 1 to 16), and restart the HDFS instance.

## 7.14 How Do I Modify the HDFS Active/Standby Switchover Class?

If the **org.apache.hadoop.hdfs.server.namenode.ha.AdaptiveFailoverProxyProvider** class is unavailable when a cluster of MRS 3.x connects to NameNodes using



HDFS, the cause is that the HDFS active/standby switchover class of the cluster is configured improperly. To solve the problem, perform the following operations:

- Method 1: Add the **hadoop-plugins-xxx.jar** package to the **classpath** or **lib** directory of your program.

The **hadoop-plugins-xxx.jar** package is stored in the HDFS client directory, for example, **\$HADOOP\_HOME/share/hadoop/common/lib/hadoop-plugins-8.0.2-302023.jar**.

- Method 2: Change the configuration item of HDFS to the corresponding open source class, as shown in the follows:

`dfs.client.failover.proxy.provider.hacluster=org.apache.hadoop.hdfs.server.name  
node.ha.ConfiguredFailoverProxyProvider`

## 7.15 What Data Type in Hive Tables Is Recommended for the Number Type of DynamoDB?

**Question:** What data type in Hive tables is recommended for the number type of DynamoDB?

**Answer:** **smallint** is recommended.

## 7.16 Can I Export the Query Result of Hive Data?

Run the following statement to export the query result of Hive data:

```
insert overwrite local directory "/tmp/out/" row format delimited fields terminated by "\t" select * from  
table;
```

## 7.17 What Should I Do If an Error Occurs When Hive Runs the beeline -e Command to Execute Multiple Statements?

When Hive uses Beeline to run the **beeline -e "use default;show tables;"** command on an MRS 3.x cluster, the following error is reported:

```
Error while compiling statement: FAILED: ParseException line 1:11 missing EOF at ',' near 'default'  
(state=42000,code=40000)
```

Solutions:

- Method 1: Replace the **beeline -e " use default;show tables;"** command with **beeline --entirelineascommand=false -e "use default;show tables;"**.
- Method 2:
  - a. In the **/opt/Bigdata/client/Hive** directory on the Hive client, change **export CLIENT\_HIVE\_ENTIRELINEASCOMMAND=true** in the **component\_env** file to **export CLIENT\_HIVE\_ENTIRELINEASCOMMAND=false**.

**Figure 7-1** Changing the `component_env` file

```
PATH_NEW="echo $PATH | sed "s|/opt/Bigdata/client/Hive/Beeline/bin:||g" | sed "s|/opt/Bigdata/client/Hive/Beeline/bin:||g"
PATH_NEW="echo $PATH_NEW | sed "s|/opt/Bigdata/client/Hive/HCatalog/bin:||g" | sed "s|/opt/Bigdata/client/Hive/HCatalog/bin:||g"
export PATH=/opt/Bigdata/client/Hive/Beeline/bin:/opt/Bigdata/client/Hive/HCatalog/bin:$PATH_NEW
export CLIENT_HIVE_URI=jdbc:hive2://192.168.0.88:2181,192.168.0.9:2181,192.168.0.258:2181/;serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2
export HIVE_HOME=/opt/Bigdata/client/Hive/Beeline
export HIVE_LIB=/opt/Bigdata/client/Hive/Beeline/lib
export HCAT_CONF_DIR=/opt/Bigdata/client/Hive/HCatalog/conf/
export CLIENT_HIVE_ENTIRELINEASCOMMAND=false
```

- b. Run the following command to verify the configuration:
- ```
source /opt/Bigdata/client/bigdata_env
beeline -e "use default;show tables;"
```

## 7.18 What Do I Do If "over max user connections" Is Displayed When Hue Connects to HiveServer?

Applicable versions: MRS 3.1.0 and earlier

1. Modify the following configuration files on all Hue nodes:  
`/opt/Bigdata/FusionInsight_Porter_8.*/install/FusionInsight-Hue-*/hue/apps/ beeswax/src/beeswax/models.py`
2. Change the configurations in lines 396 and 404.

```
q = self.filter(owner=user, application=application).exclude(guid="").exclude(secret=")
```

Modify the content as follows:

```
q = self.filter(owner=user, application=application).exclude(guid=None).exclude(secret=None)
```

**Figure 7-2** Modifying the Hue configuration file

```
apps/beeswax/src/beeswax/models.py -2-2
@@ -393,7 +393,7 @@ class SessionManager(models.Manager):
394 def get_session(self, user, application='beeswax', filter_open=True):
395     try:
396         q = self.filter(owner=user, application=application).exclude(guid="").exclude(secret=")
397         if filter_open:
398             q = q.filter(status_code=0)
399         return q.latest("last_used")
400     except:
401         return None
402
403 @@ -401,7 +401,7 @@ class SessionManager(models.Manager):
404 def get_session(self, user, application='beeswax', filter_open=True):
405     try:
406         q = self.filter(owner=user, application=application).exclude(guid="").exclude(secret=")
407         if filter_open:
408             q = q.filter(status_code=0)
409         return q.latest("last_used")
410     except:
411         return None
```

## 7.19 How Do I View MRS Hive Metadata?

- Hive metadata is stored in the GaussDB database of the MRS cluster. You can perform the following steps to view the metadata:
  - a. Log in to the DBServer active node as user **root**.
  - b. Log in to the Hive database:

```
su - omm
source $DBSERVER_HOME/.dbservice_profile
gsql -p 20051 -U USER -W PASSWD -d hivemeta
```
  - c. View all metadata tables in the Hive metadata database:

```
\d+
```
  - d. Run **q** to exit the page for viewing the metadata table.
  - e. View data in the metadata table:

```
select * from Table name;
```

 **NOTE**

To query the IP address of the active DBServer node, log in to FusionInsight Manager, choose **Cluster > Services > DBService**, and click the **Instances** tab.

You can view the default username (USER) and initial password (PASSWD) of the DBService database in [User Accounts](#).

- If Hive metadata is stored in an external relational database, perform the following steps:
  - a. On the cluster **Dashboard** page, click **Manage** on the right of **Data Connection**.
  - b. On the displayed page, obtain the value of **Data Connection ID**.
  - c. On the MRS console, click **Data Connections**.
  - d. In the data connection list, locate the data connection based on the data connection ID.
  - e. Click **Edit** in the **Operation** column of the data connection.

The **RDS Instance** and **Database** indicate the relational database in which the Hive metadata is stored.

## 7.20 How Do I Reset Kafka Data?

You can reset Kafka data by deleting Kafka topics.

- Delete a topic: **kafka-topics.sh --delete --zookeeper ZooKeeper Cluster service IP address:2181/kafka --topic *topicname***
- Query all topics: **kafka-topics.sh --zookeeper ZooKeeper cluster service IP address:2181/kafka --list**

After the deletion command is executed, empty topics will be deleted immediately. If a topic has data, the topic will be marked for deletion and will be deleted by Kafka later.

## 7.21 What Should I Do If the Error Message "Not Authorized to access group XXX" Is Displayed When Kafka Topics Are Consumed?

This issue is caused by the conflict between the Ranger authentication and ACL authentication of a cluster. If a Kafka cluster uses ACL for permission access control and Ranger authentication is enabled for the Kafka component, all authentications of the component are managed by Ranger. The permissions set by the original authentication plug-in are invalid. As a result, ACL authorization does not take effect. You can disable Ranger authentication of Kafka and restart the Kafka service to rectify the fault. The procedure is as follows:

1. Log in to FusionInsight Manager and choose **Cluster > Services > Kafka**.
2. In the upper right corner of the **Dashboard** page, click **More** and choose **Disable Ranger**. In the displayed dialog box, enter the password and click **OK**. After the operation is successful, click **Finish**.

3. In the upper right corner of the **Dashboard** page, click **More** and choose **Restart Service** to restart the Kafka service.

## 7.22 What Compression Algorithms Does Kudu Support?

**Question:** What compression algorithms does Kudu support?

**Answer:** Kudu supports Snappy, LZ4, and zlib. LZ4 is used by default.

## 7.23 How Do I View Kudu Logs?

1. Log in to the Master node in the cluster.
2. Run the **su - omm** command to switch to user **omm**.
3. Run the **cd /var/log/Bigdata/kudu/** command to go to the **/var/log/Bigdata/kudu/** directory and view Kudu logs.

## 7.24 How Do I Handle the Kudu Service Exceptions Generated During Cluster Creation?

### Viewing the Kudu Service Exception Logs

1. Log in to the MRS console.
2. Click the name of the cluster.
3. On the page displayed, choose **Components > Kudu > Instances** and locate the IP address of the abnormal instance.

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.)

4. Log in to the node where the abnormal instance resides, and view the Kudu log.

```
cd /var/log/Bigdata/Kudu
[root@node-master1AERu kudu]# ls
healthchecklog  runninglog  startlog
```

You can find the Kudu health check logs in the **healthchecklog** directory, the startup logs in the **startlog** directory, and the Kudu process run logs in the **runninglog** directory.

```
[root@node-master1AERu logs]# pwd
/var/log/Bigdata/kudu/runninglog/master/logs
[root@node-master1AERu logs]# ls -al
kudu-master.ERROR  kudu-master.INFO  kudu-master.WARNING
```

Run logs are classified into three types: ERROR, INFO, and WARNING. Each type of run logs is recorded in the corresponding file. You can run the **cat** command to view run logs of each type.

### Handling Kudu Service Exceptions

The **/var/log/Bigdata/kudu/runninglog/master/logs/kudu-master.INFO** file contains the following error information:

"Unable to init master catalog manager: not found: Unable to initialize catalog manager: Failed to initialize sys tables async: Unable to load consensus metadata for tablet 000000000000000000000000: xxx"

If this exception occurs when the Kudu service is installed for the first time, the KuduMaster service is not started. Data inconsistency causes the startup failure. To solve the problem, perform the following steps to clear the data directories and restart the Kudu service. If the Kudu service is not installed for the first time, clearing the data directories will cause data loss. In this case, migrate data and clear the data directory.

1. Search for the data directories **fs\_data\_dir**, **fs\_wal\_dir**, and **fs\_meta\_dir**.  
**find /opt -name master.gflagfile**  
**cat /opt/Bigdata/FusionInsight\_Kudu\_\*/\*\_KuduMaster/etc/master.gflagfile**  
**| grep fs\_**
2. On the cluster details page, choose **Components > Kudu** and click **Stop Service**.
3. Clear the Kudu data directories on all KuduMaster and KuduTserver nodes. The following command uses two data disks as an example.  
**rm -Rvf /srv/Bigdata/data1/kudu, rm -Rvf /srv/Bigdata/data2/kudu**
4. On the cluster details page, choose **Components > Kudu** and choose **More > Restart Service**.
5. Check the Kudu service status and logs.


## 7.25 How Do I Configure Other Data Sources on Presto?

This topic uses the MySQL data source as an example.

- For MRS 1.x and 3.x clusters, do the following:
  - a. Log in to the MRS management console.
  - b. Click the name of the cluster to go to its details page.
  - c. Click the **Components** tab and then **Presto** in the component list. On the page that is displayed, click the **Configurations** tab then the **All Configurations** sub-tab.
  - d. On the Presto configuration page that is displayed, find **connector-customize**.
  - e. Set **Name** and **Value** as follows:

**Name: mysql.connector.name**

**Value: mysql**



- f. Click the plus sign (+) to add three more fields and set **Name** and **Value** according to the table below. Then click **Save**.




Table 7-1 Parameter description

| Name                      | Value                             | Description              |
|---------------------------|-----------------------------------|--------------------------|
| mysql.connection-url      | jdbc:mysql://xxx.xxx.xxx.xxx:3306 | Database connection pool |
| mysql.connection-user     | xxxx                              | Database username        |
| mysql.connection-password | xxxx                              | Database password        |

- g. Restart the Presto service.
- h. Run the following command to connect to the Presto Server of the cluster:  
**presto\_cli.sh --krb5-config-path {krb5.conf path} --krb5-principal {User principal} --krb5-keytab-path {user.keytab path} --user {presto username}**
- i. Log in to Presto and run the **show catalogs** command to check whether the data source list of Presto can be queried.

```
[root@node-master2uoHG bin]# ./presto_cli.sh
--server http://15xxxxxx20
show catalogs;

Catalog
-----
hive
jmx
mysql
system
tpcds
tpch
(6 rows)

Query 20220422_121338_00002_ra2vb, FINISHED, 3 nodes
Splits: 53 total, 53 done (100.00%)
0:00 [0 rows, 0B] [0 rows/s, 0B/s]
```

Run the **show schemas from mysql** command to query the MySQL database.

- For MRS 2.x clusters, do the following:
  - a. Create the **mysql.properties** configuration file containing the following content:  
connector.name=mysql  
connection-url=jdbc:mysql://IP address of MySQL:3306  
connection-user=Username  
connection-password=Password

 NOTE

- The IP Address of MySQL is that of the MySQL instance, which must be able to communicate with the MRS cluster.
  - The username and password are those used to log in to the MySQL database.
  - Configuration files containing authentication passwords pose security risks. Delete such files after configuration or store them securely.
- b. Upload the configuration file to the **/opt/Bigdata/MRS\_Current/1\_14\_Coordinator/etc/catalog/** directory on the Master node (where the Coordinator instance resides) and the **/opt/Bigdata/MRS\_Current/1\_14\_Worker/etc/catalog/** directory on the Core node (depending on the actual directory in the cluster), and change the file owner group to **omm:wheel**.
- c. Restart the Presto service.

## 7.26 How Do I Update the Ranger Certificate in MRS 1.9.3?

MRS 1.9.3 is used as an example. Replace it with the actual cluster version. After the certificate is updated, manually clear the alarm indicating that the certificate file is invalid or about to expire.

 NOTE

- The validity period of the Ranger certificate is 10 years.
- After the Ranger certificate expires, the Ranger web UI can still be accessed and functions properly. However, there will be a message indicating that the certificate is untrusted upon your access.
- If Ranger is not installed in the cluster, log in to each master node and run the following command to rename the certificate file:  
**mv /opt/Bigdata/MRS\_1.9.3/install/MRS-Ranger-1.0.1/ranger/ranger-1.0.1-admin/ranger-admin-keystore.jks /opt/Bigdata/MRS\_1.9.3/install/MRS-Ranger-1.0.1/ranger/ranger-1.0.1-admin/ranger-admin-keystore.jks\_bak**
- If Ranger has been installed in the cluster, update the certificate as follows:
  - a. Download **MRS\_1.9\_Patch\_UpdateRangerJks\_All\_20210203.tar.gz** from the obs-patch bucket and upload it to the **/tmp** directory on the node where the active RangerAdmin instance of the cluster runs.  
On MRS Manager, choose **Service > Ranger > Instance** and obtain the IP address of the node where the active RangerAdmin instance runs.
    - CN-Hong Kong: [https://mrs-patch-ap-southeast-1.obs.ap-southeast-1.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_1.9\\_Patch\\_UpdateRangerJks\\_All\\_20210203.tar.gz](https://mrs-patch-ap-southeast-1.obs.ap-southeast-1.myhuaweicloud.com/MRS_Common_Script/MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz)
    - AP-Bangkok: [https://mrs-patch-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_1.9\\_Patch\\_UpdateRangerJks\\_All\\_20210203.tar.gz](https://mrs-patch-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/MRS_Common_Script/MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz)
    - AP-Singapore: [https://mrs-patch-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_1.9\\_Patch\\_UpdateRangerJks\\_All\\_20210203.tar.gz](https://mrs-patch-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/MRS_Common_Script/MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz)

- LA-Sao Paulo: [https://mrs-container1-patch-sa-brazil-1.obs.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_1.9\\_Patch\\_UpdateRangerJks\\_All\\_20210203.tar.gz](https://mrs-container1-patch-sa-brazil-1.obs.myhuaweicloud.com/MRS_Common_Script/MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz)
  - LA-Mexico City: [https://mrs-container1-patch-na-mexico-1.obs.myhuaweicloud.com/MRS\\_Common\\_Script/MRS\\_1.9\\_Patch\\_UpdateRangerJks\\_All\\_20210203.tar.gz](https://mrs-container1-patch-na-mexico-1.obs.myhuaweicloud.com/MRS_Common_Script/MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz)
- b. Log in to the node where the active RangerAdmin instance is located and run the following commands:
- ```
cd /tmp
chmod 700 MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz
chown omm:wheel
MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz
su - omm
cd /tmp
tar -zxvf MRS_1.9_Patch_UpdateRangerJks_All_20210203.tar.gz
```
- c. Replace the certificate files.
- ```
cd updateRangerJks
sh updateRangerJks.sh ${IP address of the active Master node} ${IP
address of the active RangerAdmin node} ${Certificate password}
```
-  **NOTE**
- This script will restart the controller process. During the restart process, the MRS Manager page may not be viewed.
  - Obtain the IP address of the active Master node from **Hosts** on MRS Manager.
  - To obtain the IP address of the active RangerAdmin node, choose **Services > Ranger > Instances** on MRS Manager.
  - *\${Certificate password}* is a user-defined password. Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.
- d. Log in to the MRS console.
- e. Choose **Active Clusters** and click a cluster name to go to the cluster details page.
- f. Choose **Components > Ranger > Service Configuration** and modify the RangerAdmin configuration.
- i. Search for the **polycmgr\_https\_keystore\_password** and change its value to the certificate password entered in **c**, that is, *\${Certificate password}*.  
You are advised to copy and paste the password. If the passwords are different, Ranger will fail to restart.
  - ii. Save the configuration and perform a rolling restart of RangerAdmin.
- g. Verify that you can log in to the RangerAdmin web UI.
- i. Choose **Components > Ranger > Service Status**. In **Ranger Summary**, click **RangerAdmin** corresponding to **Ranger Web UI**.
  - ii. On the Ranger web UI login page, the default username for MRS cluster 1.9.2 is **admin** and the password is **admin@12345**. The



default username for MRS cluster 1.9.3 or later 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.

- h. Log in to the node where the RangerAdmin instance is located and delete the temporary files.

```
rm -rf /tmp/updateRangerJks
```

```
rm -rf /tmp/updateRangerJks.tar.gz
```

For a cluster with a custom topology, if the active master and RangerAdmin instances are not on the same node, log in to the active master node and delete temporary files.

## 7.27 How Do I Specify a Log Path When Submitting a Task in an MRS Storm Cluster?

You can modify the `/opt/Bigdata/MRS_XXX/1_XX_Supervisor/etc/worker.xml` file on the streaming Core node of MRS, set the value of **filename** to the path, and restart the corresponding instance on Manager.


You are advised not to modify the default log configuration of MRS. Otherwise, the log system may become abnormal.

## 7.28 How Do I Check the ResourceManager Configuration of Yarn?

This section applies only to versions earlier than MRS 3.x.

**Step 1** Log in to MRS Manager and choose **Services > Yarn > Instance**.

**Step 2** Click the names of the two ResourceManager nodes, choose **More > Synchronize Configuration**, and deselect **Restart services or instances whose configurations have expired**.

**Figure 7-3** Synchronizing the configuration

**Step 3** Click **Yes** to synchronize the configuration.

**Step 4** Log in to the Master nodes as user **root**.

**Step 5** Run the **cd /opt/Bigdata/MRS\_Current/\*\_ResourceManager/etc\_UPDATED/** command to go to the **etc\_UPDATED** directory.

**Step 6** Run the **grep '\.queues' capacity-scheduler.xml -A2** command to display all configured queues and check whether the queues are consistent with those displayed on Manager.

**root-default** is hidden on the Manager page.

```
[omm@node-master1ilza etc]$  
[omm@node-master1ilza etc]$ grep '\.queues' capacity-scheduler.xml -A2  
<name>yarn.scheduler.capacity.root.queues</name>  
<value>default,root-default,launcher-job,test1,test2,test3,test4</value>  
</property>  
[omm@node-master1ilza etc]$  
[omm@node-master1ilza etc]$
```

**Step 7** Run the **grep '\.capacity</name>' capacity-scheduler.xml -A2** command to display the value of each queue and check whether the value of each queue is the same as that displayed on Manager. Check whether the sum of the values configured for all queues is **100**.

- If the sum is **100**, the configuration is correct.
- If the sum is not **100**, the configuration is incorrect. Perform the following steps to rectify the fault.

```
[omm@node-master111ZA etc]$  
[omm@node-master111ZA etc]$ grep '\.capacity</name>' capacity-scheduler.xml -A2  
<name>yarn.scheduler.capacity.root.root-default.accessible-node-labels.zhaolu.capacity</name>  
<value>0.0</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.launcher-job.capacity</name>  
<value>10</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.accessible-node-labels.zhaolu.capacity</name>  
<value>100</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.test1.capacity</name>  
<value>10</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.test2.capacity</name>  
<value>10</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.test3.capacity</name>  
<value>10</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.capacity</name>  
<value>100</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.root-default.capacity</name>  
<value>40.0</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.test4.accessible-node-labels.zhaolu.capacity</name>  
<value>100</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.test4.capacity</name>  
<value>0</value>  
</property>  
--  
<name>yarn.scheduler.capacity.root.default.capacity</name>  
<value>20</value>  
</property>  
[omm@node-master111ZA etc]$
```

**Step 8** Log in to MRS Manager, and select **Hosts**.

**Step 9** Determine the active Master node. The host name of the active Master node starts with a solid pentagon.

**Step 10** Log in to the active Master node as user **root**.

**Step 11** Run the **su - omm** command to switch to user **omm**.

**Step 12** Run the **sh /opt/Bigdata/om-0.0.1/sbin/restart-controller.sh** command to restart the controller when no operation is being performed on Manager.

Restarting the controller will not affect the big data component services.

**Step 13** Repeat **Step 1** to **Step 7** to synchronize ResourceManager configurations and check whether the configurations are correct.

If the latest configuration has not been loaded after the configuration synchronization is complete, a message will be displayed on the Manager page indicating that the configuration has expired. However, this will not affect services. The latest configuration will be automatically loaded when the component restarts.

----End

## 7.29 How Do I Modify the allow\_drop\_detached Parameter of ClickHouse?

- Step 1** Log in to the node where the ClickHouse client is located as user **root**.
- Step 2** Run the following commands to go to the client installation directory and set the environment variables:

```
cd /opt/Client installation directory
```

```
source bigdata_env
```

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

```
kinit MRS cluster user
```

### NOTE

The user must have the ClickHouse administrator permissions.

- Step 4** Run the **clickhouse client --host 192.168.42.90 --secure -m** command, in which **192.168.42.90** indicates the IP address of the ClickHouseServer instance node. The command output is as follows:

```
[root@server-2110082001-0017 hadoopclient]# clickhouse client --host 192.168.42.90 --secure -m
ClickHouse client version 21.3.4.25.
Connecting to 192.168.42.90:21427.
Connected to ClickHouse server version 21.3.4 revision 54447.
```

- Step 5** Run the following command to set the value of the **allow\_drop\_detached** parameter, for example, **1**:

```
set allow_drop_detached=1;
```

- Step 6** Run the following command to query the value of the **allow\_drop\_detached** parameter:

```
SELECT * FROM system.settings WHERE name = 'allow_drop_detached';
```

```
server-2110081635-0001 :) SELECT * FROM system.settings WHERE name = 'allow_drop_detached';
SELECT *
FROM system.settings
WHERE name = 'allow_drop_detached'
Query id: 8211d1ff-5717-49af-929f-8e4170c6e1d1

+-----+-----+-----+-----+-----+-----+-----+-----+
| name          | value | changed | description                                     | min  | max  | readonly | type  |
+-----+-----+-----+-----+-----+-----+-----+-----+
| allow_drop_detached | 1     | 1       | Allow ALTER TABLE ... DROP DETACHED PART[ITION] ... queries | NULL | NULL | 0        | Bool |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 rows in set. Elapsed: 0.004 sec.
```

- Step 7** Run the **q;** command to exit the ClickHouse client.

----End

## 7.30 How Do I Add a Periodic Deletion Policy to Prevent Large ClickHouse System Table Logs?

### Symptom

ClickHouse logs generated in the system table is too large. Deleting logs at a time takes a long time.

### Handling Procedure

- For MRS 3.3.0 and later versions, on FusionInsight Manager, choose **Cluster > Services > ClickHouse > Configurations > All Configurations**. You can set **query\_log.ttl** and **trace\_log.ttl** to configure the data clearing period, the default value is 90 days.
- In versions earlier than MRS 3.3.0, run the following statement on the ClickHouse client to modify the TTL of a table in ClickHouse:  
**alter table *system.Table name* modify TTL event\_date + INTERVAL *Number of days for storing data* day;**

---

**CAUTION**

- This statement is used only on a SQL node to configure the TTL of the system table. If the TTL needs to be configured on all nodes, run this statement on each node. Do not run the statement with **on cluster** on each node. Otherwise the ClickHouse will keep running.
  - Run the preceding statement during off-peak hours. This operation may take a long time if there is a large amount of data.
- 

## 7.31 How Do I Change the Time Zone of the ClickHouse Service?

### Symptom

When data is inserted to ClickHouse, the time zone is incorrect. The time difference is eight hours.

### Fault Locating

The Clickhouse time zone is Europe/Moscow by default. You need to change it to the local time zone.

### Procedure

- Step 1** Log in to FusionInsight Manager, choose **Cluster > Services > ClickHouse**, and click **Configurations** and then **All Configurations**.

**Step 2** Search for **clickhouse-config-customize** in the search box on the left and add the parameter name **timezone** and value **Asia/Shanghai**.

 **NOTE**

The value of **timezone** must be the same as the local time zone. For example, the value for UTC+8 is **Asia/Shanghai**, and the value for UTC+0 is **Europe/London**.

**Figure 7-4** Configuring the ClickHouse time zone

| Parameter                                                                                                     | Value                                                                                                                                                                                       |      |       |                                       |                                            |
|---------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|-------|---------------------------------------|--------------------------------------------|
| ClickHouse->ClickHouseServer                                                                                  |                                                                                                                                                                                             |      |       |                                       |                                            |
|  clickhouse-config-customize | <table><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td><input type="text" value="timezone"/></td><td><input type="text" value="Asia/Shanghai"/></td></tr></tbody></table> | Name | Value | <input type="text" value="timezone"/> | <input type="text" value="Asia/Shanghai"/> |
| Name                                                                                                          | Value                                                                                                                                                                                       |      |       |                                       |                                            |
| <input type="text" value="timezone"/>                                                                         | <input type="text" value="Asia/Shanghai"/>                                                                                                                                                  |      |       |                                       |                                            |

----End

# 8 Cluster Management

---

## 8.1 How Do I View All Clusters?

All MRS clusters are displayed on the **Clusters** page of the MRS management console. You can view clusters in different status.

- **Active Clusters:** all clusters except clusters in **Failed** and **Terminated** states.
- **Cluster History:** clusters in the **Deleted** state. Only the clusters deleted within the last six months are displayed. If you want to view clusters deleted more than six months ago, contact technical support engineers.
- **Failed Tasks:** tasks in **Failed** state. The failed tasks include the following:
  - Tasks failed to create clusters
  - Tasks failed to terminate clusters
  - Tasks failed to scale out clusters
  - Tasks failed to scale in clusters

## 8.2 How Do I View MRS Operation Logs?

You can view operation logs of clusters and jobs on the **Operation Logs** page. The MRS operation logs record the following operations:

- Cluster operations
  - Create, terminate, and scale out or in clusters
  - Create directories and delete directories or files
- Job operations: Create, stop, and delete jobs
- Data operations: IAM user tasks, add users, and add user groups

**Figure 8-1** shows the operation logs.

**Figure 8-1** Log information

| Operation Type | Operator IP Address | Operation Description                                   | Time                |
|----------------|---------------------|---------------------------------------------------------|---------------------|
| Cluster        |                     | Create id is : aae6 and name as: bigdata_xq318 cluster  | 2016-03-18 17:17:46 |
| Cluster        |                     | Delete the id for 7631 name for bigdata_DVwu cluster    | 2016-03-10 16:45:24 |
| Job            |                     | createJob.jobId: 93a2.jobName:distcp_clusterid: 7631    | 2016-03-10 10:26:28 |
| Job            |                     | createJob.jobId: c981.jobName:job_spark_clusterid: 7631 | 2016-03-07 11:02:28 |

## 8.3 How Do I View MRS Cluster Configuration Information?

- After a cluster is created, click the cluster name on the MRS console. On the page displayed, you can view basic configuration information about the cluster, including the name, ID, billing mode, AZ, creation time, Hadoop component version, as well as the instance specifications and capacity of node. The instance specifications and node capacity determine the data analysis and processing capability. Higher instance specifications and larger capacity enable faster data processing at a higher cost.
- On the basic information page, click **Access Manager** to access the MRS cluster management page. On MRS Manager, you can view and handle alarms, and modify cluster configuration.

## 8.4 How Do I Add Components to an MRS Cluster?

**Question:** How do I add services to an existing MRS cluster?

**Answer:**

You cannot install new components for a created cluster whose version is MRS 3.1.0 or earlier. If you want to use a component that is not installed in your cluster, you need to create a cluster and select the required component.

You can add components to MRS 3.1.2-LTS.3 or later custom clusters. For details, see [Managing Services](#).

## 8.5 How Do I Cancel Message Notification for Cluster Alarms?

- Log in to the MRS console.
- Click the name of the cluster.
- On the page displayed, choose **Alarms > Notification Rules**.
- Locate the row that contains the rule you want to modify, click **Edit** in the **Operation** column, and deselect the alarm or event severity levels.
- Click **OK**.



## 8.6 Why Is the Resource Pool Memory Displayed in the MRS Cluster Smaller Than the Actual Cluster Memory?

**Question:** Why is the resource pool memory displayed in the MRS cluster smaller than the actual cluster memory?

**Answer:** In an MRS cluster, MRS allocates 50% of the cluster memory to Yarn by default. You manage Yarn nodes logically by resource pool. Therefore, the total memory of the resource pool displayed in the cluster is only 50% of the total memory of the cluster.

## 8.7 What Is the Python Version Installed for an MRS Cluster?

Log in to a Master node as user **root** and run the **Python3** command to query the Python version.

**Table 8-1** Python versions supported by MRS

| MRS Cluster Version | Python Version |
|---------------------|----------------|
| MRS 3.1.0           | Python 3.8.0   |
| MRS 3.0.5           | Python 3.7.0   |
| MRS 3.0.2           | Python 3.7.0   |
| MRS 2.1.1           | Python 3.6.8   |
| MRS 2.1.0           | Python 3.6.8   |
| MRS 1.9.3           | Python 3.6.8   |

## 8.8 How Do I Upload a Local File to a Node Inside a Cluster?

- Step 1** Log in to the MRS console.
- Step 2** In the navigation pane on the left, choose **Active Clusters** and click the name of the target cluster to go to its details page.
- Step 3** On the **Nodes** page, click a node name to log in to the ECS console.
- Step 4** Bind an EIP to a node in the cluster. For details, see [Assigning an EIP](#).
- Step 5** Upload the local file to the cluster node. For details, see [How Can I Upload a File to an ECS?](#)
- End

## 8.9 What Can I Do If the Time Information of an MRS Cluster Node Is Incorrect?

- If the time on a node inside the cluster is incorrect, log in to the node and rectify the fault from 2.
  - If the time on a node inside the cluster is different from that on a node outside the cluster, log in to the node and rectify the fault from 1.
1. 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 8-2 Adding the master node IP addresses



```
# 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
```

2. Run the **service ntpd stop** command to stop the NTP service.
3. Run the **/usr/sbin/ntpdate IP address of the active master node** command to manually synchronize time.
4. Run the **service ntpd start** or **systemctl restart ntpd** command to start the NTP service.
5. Run the **ntpstat** command to check the time synchronization result:

## 8.10 What Are the Differences and Relationships Between the MRS Management Console and MRS Manager?

### Question:

What Are the Differences and Relationships Between the MRS Management Console and MRS Manager?

### Answer:

You can access Manager from the MRS management console.

Manager is classified as MRS Manager and FusionInsight Manager.

- MRS Manager is the manager page of MRS 2.x or earlier clusters.
- FusionInsight Manager is the manager page of MRS 3.x or later clusters.

The following table lists the differences and relationships between the management console and FusionInsight Manager.

| Common Operation                                                                                                           | MRS Console   | FusionInsight Manager |
|----------------------------------------------------------------------------------------------------------------------------|---------------|-----------------------|
| Changing subnets, adding security group rules, controlling OBS permissions, managing agencies, and synchronizing IAM users | Supported     | Not supported         |
| Adding node groups, scaling out, scaling in, and upgrading specifications                                                  | Supported     | Not supported         |
| Isolating hosts, starting all roles, and stopping all roles                                                                | Supported     | Supported             |
| Downloading the client, starting services, stopping services, and perform rolling restart of services                      | Supported     | Supported             |
| Viewing the instance status of services, configuring parameters, and synchronizing configurations                          | Supported     | Supported             |
| Viewing cleared alarms and events                                                                                          | Supported     | Supported             |
| Viewing the alarm help                                                                                                     | Supported     | Supported             |
| Setting thresholds for threshold alarms                                                                                    | Not supported | Supported             |
| Adding message subscription specifications                                                                                 | Supported     | Not supported         |
| Managing files                                                                                                             | Supported     | Not supported         |
| Managing jobs                                                                                                              | Supported     | Not supported         |
| Managing tenants                                                                                                           | Supported     | Supported             |
| Managing tags                                                                                                              | Supported     | Not supported         |

| Common Operation                                                        | MRS Console   | FusionInsight Manager |
|-------------------------------------------------------------------------|---------------|-----------------------|
| Setting permissions (adding and deleting users, user groups, and roles) | Not supported | Supported             |
| Backing up component data and restoring data                            | Not supported | Supported             |
| Logging metadata operation audit                                        | Not supported | Supported             |
| Monitoring resources                                                    | Supported     | Supported             |

## 8.11 How Do I Unbind an EIP from FusionInsight Manager of an MRS Cluster?

### Question:


How Do I Unbind an EIP from FusionInsight Manager of an MRS Cluster?

### Answer:

After an EIP is bound to FusionInsight Manager, do not unbind the EIP. Otherwise, other users may fail to access FusionInsight Manager.


#### NOTE

After an EIP is bound to FusionInsight Manager on the MRS console, the EIP cannot be unbound from FusionInsight Manager on the **EIPs** page.



To unbind an EIP, call the related API of the EIP service.

**Step 1** Log in to the EIP management console. In the EIP list, view and record the ID of the desired EIP.



**Step 2** Unbind the EIP from API Explorer by referring to [Updating an EIP](#).

For example, log in to API Explorer as the current user, select a region based on site requirements, and set **project\_id** to the project ID of the current region and **publicip\_id** to the ID of the EIP to be unbound queried in [Step 1](#).

**Step 3** After the API is successfully called, log in to the EIP management console. In the EIP list, you can see that the EIP is in the **Unbound** state.

**Step 4** Log in to the MRS management console, go to the **Dashboard** tab page of the cluster, and bind another EIP to the cluster.

----End

## 8.12 How Do I Stop the Firewall Service?

### Question:

How Do I Stop the Firewall Service?

### Answer:

**Step 1** Log in to each node of a cluster as user **root**.

**Step 2** Check whether the firewall service is started.

For example, to check the firewall status on EulerOS, run the **systemctl status firewalld.service** command.

**Step 3** Stop the firewall service.

For example, to stop the firewall service on EulerOS, run the **systemctl stop firewalld.service** command.

----End

## 8.13 How Do I Switch the Login Mode of a Node in an MRS Cluster?

**Question:** How do I switch the login mode (password or key pair) of a node in an MRS cluster?

**Answer:** Login mode cannot be changed. You can select the login mode when creating the cluster, but cannot change the login mode after you create the cluster.

## 8.14 How Do I Access an MRS Cluster from a Node Outside the Cluster?

## Creating a Linux ECS Outside the Cluster to Access the MRS Cluster

**Step 1** Create an ECS outside the cluster. For details, see [Purchasing an ECS](#).


Set **AZ**, **VPC**, and **Security Group** of the ECS to the same values as those of the cluster to be accessed.

**Step 2** On the VPC management console, apply for an EIP and bind it to the ECS.


For details, see [Assigning an EIP](#).

### Step 3 Configure security group rules for the cluster.

1. On the **Dashboard** tab page, click **Add Security Group Rule**. In the **Add Security Group Rule** dialog box that is displayed, click **Manage Security Group Rule**.



- Click the **Inbound Rules** tab, and click **Add Rule**. In the **Add Inbound Rule** dialog box, configure the IP address of the ECS and enable all ports.



- After the security group rule is added, you can download and install the client on the ECS. For details, see [Installing a Client](#).
- Use the client.

Log in to the client node as the client installation user and run the following command to switch to the client directory:

```
cd /opt/hadoopclient
```

Run the following command to load environment variables:

```
source bigdata_env
```

If Kerberos authentication is enabled for the cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, authentication is not required.

```
kinit MRS cluster user
```

Example:

```
kinit admin
```

Run the client command of a component.

Example:

Run the following command to view files in the HDFS root directory:

```
hdfs dfs -ls /
```

```
Found 15 items
drwxrwx--x - hive      hive      0 2021-10-26 16:30 /apps
drwxr-xr-x - hdfs      hadoop   0 2021-10-18 20:54 /datasets
drwxr-xr-x - hdfs      hadoop   0 2021-10-18 20:54 /datastore
drwxrwx---+ - flink     hadoop   0 2021-10-18 21:10 /flink
drwxr-x--- - flume      hadoop   0 2021-10-18 20:54 /flume
drwxrwx--x - hbase     hadoop   0 2021-10-30 07:31 /hbase
...
```

----End

## Creating a Windows ECS Outside the Cluster to Access MRS Manager

You can create a Windows ECS on the ECS console to access MRS Manager.

## 8.15 In an MRS Streaming Cluster, Can the Kafka Topic Monitoring Function Send Alarm Notifications?

### Question:

In an MRS Streaming Cluster, Can the Kafka Topic Monitoring Function Send Alarm Notifications?

### Answer:

The Kafka topic monitoring cannot send notifications by email or SMS message. Currently, you can view alarm information on Manager.

## 8.16 Where can I view the running resource queues when ALM-18022 Insufficient Yarn Queue Resources is generated?

### Question:

Where can I view the running resource queues when ALM-18022 Insufficient Yarn Queue Resources is generated?

### Answer:

Log in to FusionInsight Manager and choose **Cluster > Services > Yarn**. In the navigation pane on the left, choose **ResourceManager(Active)** and log in to the native Yarn page.

For details, see the online help.

## 8.17 How Do I Understand the Multi-Level Chart Statistics in the HBase Operation Requests Metric?


### Question:

How Do I Understand the Multi-Level Chart Statistics in the HBase Operation Requests Metric?

### Answer:

The following uses the **Operation Requests on RegionServers** monitoring item as an example:


1. Log in to FusionInsight Manager and choose **Cluster > Services > HBase > Resource**. On the displayed page, you can view the **Operation Requests on RegionServers** chart. If you click **all**, the top 10 RegionServers ranked by the total number of operation requests in the current cluster are displayed, the statistics interval is 5 minutes.






- Click a point in the chart. A level-2 chart is displayed, showing the number of operation requests of all RegionServers in the past 5 minutes.


## Operation Requests on RegionServers



- Click an operation statistics bar chart. A level-3 chart is displayed, showing the distribution of operations in each region within the period.



- Click a region name. The distribution chart of operations performed every 5 minutes in the last 12 hours is displayed. You can view the number of operations performed in the period.



# 9 Node Management

## 9.1 What are the Operating Systems of Hosts in MRS Clusters of Different Versions?

### Question:

What are the Operating Systems of Hosts in MRS Clusters of Different Versions?

### Answer:

The host operating system varies depending on the cluster version. [Table 1](#) shows you the mapping relationship.

**Table 9-1** Mappings between cluster versions and host OSs

| Cluster Version | x86          | Kunpeng (Arm) |
|-----------------|--------------|---------------|
| MRS 3.2.0-LTS.1 | EulerOS 2.10 | EulerOS 2.10  |
| MRS 3.1.5       | EulerOS 2.9  | EulerOS 2.9   |
| MRS 3.1.2-LTS.3 | EulerOS 2.9  | EulerOS 2.9   |
| MRS 3.1.0       | EulerOS 2.5  | EulerOS 2.8   |
| MRS 1.9.2       | EulerOS 2.2  | EulerOS 2.8   |

## 9.2 Do I Need to Shut Down a Master Node Before Upgrading It?

**Question:** Do I need to shut down a Master node before upgrading it?

**Answer:** The Master node automatically shuts down during specifications upgrade, and automatically starts after the upgrade. No manual intervention is required.

## 9.3 Can I Change MRS Cluster Nodes on the MRS Console?

You cannot change MRS cluster nodes on the MRS console. You are also advised not to change MRS cluster nodes on the ECS console. Manually stopping or deleting an ECS, modifying or reinstalling the ECS OS, or modifying ECS specifications for a cluster node on the ECS console will affect the cluster stability.

If an ECS is deleted, the ECS OS is modified or reinstalled, or the ECS specifications are modified on the ECS console, MRS will automatically identify and delete the node. You can log in to the MRS console and restore the deleted node through scale-out. Do not perform operations on the nodes that are being scaled out.

## 9.4 How Do I Query the Startup Time of an MRS Node?

Log in to the target node and run the following command to query the startup time:

```
date -d "$(awk -F. '{print $1}' /proc/uptime) second ago" +"%Y-%m-%d %H:%M:%S"
```

```
[root@server-2110082001-0018 ~]#date -d "$(awk -F. '{print $1}' /proc/uptime) second ago" +"%Y-%m-%d %H:%M:%S"  
2021-12-13 15:56:23
```

## 9.5 What Do I Do If Trust Relationships Between Nodes Are Abnormal?

If "ALM-12066 Inter-Node Mutual Trust Fails" is reported on Manager or there is no SSH trust relationship between nodes, rectify the fault by performing the following operations:

1. Run the **ssh-add -l** command on both nodes of the trusted cluster to check whether there are identities.

```
[omm@node-group-2eU40 ~]$  
[omm@node-group-2eU40 ~]$  
[omm@node-group-2eU40 ~]$  
[omm@node-group-2eU40 ~]$ ll .ssh/  
total 32  
srw----- 1 omm wheel 0 Dec 29 14:17 agent.pid  
-rw----- 1 omm wheel 12901 Mar 9 14:48 authorized_keys  
-rw----- 1 omm wheel 54 Sep 24 11:42 config  
-rw----- 1 omm wheel 1766 Sep 24 11:43 id_rsa  
-rw----- 1 omm wheel 402 Sep 24 11:42 id_rsa.pub  
-rw----- 1 omm wheel 88 Jun 8 2020 id_rsa.sha256  
[omm@node-group-2eU40 ~]$  
[omm@node-group-2eU40 ~]$ ssh-add -l  
The agent has no identities.  
[omm@node-group-2eU40 ~]$  
[omm@node-group-2eU40 ~]$  
[omm@node-group-2eU40 ~]$  
[omm@node-group-2eU40 ~]$ vim /var/log/Bigdata/nodeagent/  
agentlog/ alarmlog/ monitorlog/ scriptlog/  
[omm@node-group-2eU40 ~]$ vim /var/log/Bigdata/nodeagent/scriptlog/  
agent_alarm_py.log install.log  
agent_alarm_py.log.1 installntp.log
```

2. If no identities are displayed, run the **ps -ef|grep ssh-agent** command to find the ssh-agent process, kill the process, and wait for the process to automatically restart.

```
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$ ssh-add -l  
The agent has no identities.  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$ ps -ef|grep ssh-agent  
omm 18729 1 0 14:53 ? 00:00:00 ssh-agent -a /home/omm/.ssh/agent.pid  
omm 25098 1 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor-startup.sh  
omm 25206 25098 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor.sh  
omm 27201 4913 0 14:54 pts/0 00:00:00 grep --color=auto ssh-agent  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$ ssh-add -l
```

3. Run the **ssh-add -l** command to check whether the identities have been added. If yes, manually run the **ssh** command to check whether the trust relationship is normal.

```
omm 22276 4913 0 14:53 pts/0 00:00:00 grep --color=auto ssh-agent  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$ ssh-add -l  
The agent has no identities.  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$ ps -ef|grep ssh-agent  
omm 18729 1 0 14:53 ? 00:00:00 ssh-agent -a /home/omm/.ssh/agent.pid  
omm 25098 1 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor-startup.sh  
omm 25206 25098 0 14:54 ? 00:00:00 bash /opt/Bigdata/om-agent/nodeagent/bin/ssh-agent-monitor.sh  
omm 27201 4913 0 14:54 pts/0 00:00:00 grep --color=auto ssh-agent  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$ ssh-add -l  
2048 SHA256:uchnRbhh1HYxpt0Z1bS0zym1kXMIaFyvn0IMpiZjg /home/omm/.ssh/id_rsa (RSA)  
omm@node-group-2eu40 ~]$  
omm@node-group-2eu40 ~]$ ssh 10.33.109.226  
Warning: Permanently added '10.33.109.226' (ECDSA) to the list of known hosts.
```

4. If identities exist, check whether the **authorized\_keys** file in the **/home/omm/.ssh** directory contains the information in the **id\_rsa.pub** file in the **/home/omm/.ssh** of the peer node. If no, manually add the information about the peer node.
5. Check whether the permissions on the files in **/home/omm/.ssh** directory are correct.
6. Check the **/var/log/Bigdata/nodeagent/scriptlog/ssh-agent-monitor.log** file.
7. If the **ommuser** directory has been deleted, contact MRS support personnel.

## 9.6 Can Master Node Specifications Be Adjusted in an MRS Cluster?

Yes.

You can scale up the specifications of a master node in an existing MRS cluster containing at least two master nodes. For details, see [Scaling Up Master Node Specifications](#).

## 9.7 Can Sudo Logs of Nodes in an MRS Cluster Be Cleared?

### Question:

Can Sudo Logs of Nodes in an MRS Cluster Be Cleared?

### Answer:

The sudo log file on the nodes in the MRS cluster records the operations performed by user **omm** to facilitate fault locating. You can clear the file as you need.

You can delete the logs in earlier days to release storage space.

1. If the log file is large, add the log file directory to **/etc/logrotate.d/syslog** to enable periodical deletion.  
Run **sed -i '3 a/var/log/sudo/sudo.log' /etc/logrotate.d/syslog**.
2. Set the maximum number and size of logs in **/etc/logrotate.d/syslog**. If the threshold is exceeded, the logs will be automatically deleted. By default, logs are aged based on the size and number of archived logs. You can use **size** and **rotate** to limit the size and number of archived logs, respectively. If required, you can also add **daily/weekly/monthly** to specify how often the logs are cleared.

## 9.8 How Do I Partition Disks in an MRS Cluster?

MRS clusters can be used immediately after being delivered. You do not need to plan disk partitions. [Table 9-2](#) describes the OS disk partitions of a created node.

**Table 9-2** OS disk partitions of an MRS cluster node

| Partition Type | Partition Directory | Capacity                                                                          | Usage                                                                                                                                                             |
|----------------|---------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OS partition   | /                   | 220 GB                                                                            | OS root partition directory and program storage directory, including all directories (except the directories specified below)                                     |
|                | /tmp                | 10 GB                                                                             | Directory for storing temporary files                                                                                                                             |
|                | /var                | 10 GB                                                                             | OS run directory                                                                                                                                                  |
|                | /var/log            | The remaining space of the OS disk is allocated to the <b>/var/log</b> partition. | Directory for storing logs                                                                                                                                        |
|                | /srv/BigData        | 60 GB                                                                             | Data directory of FusionInsight Manager, which stores data such as ldapData, Manager, and metric_agent, and provides mount points for component data directories. |

After an MRS node is created, the non-OS disks of the node are mounted to the **/srv/BigData/dataN** directory. For example, if the node has four data disks, the disk mounting directories are **/srv/BigData/data1**, **/srv/BigData/data2**, **/srv/BigData/data3**, and **/srv/BigData/data4**.

The metadata and data directories of each component deployed on the node are allocated to different disk partition directories based on certain mapping rules.

**Table 9-3** lists the data directory description of each component.

**Table 9-3** Non-OS disk partitions of an MRS cluster node

| Partition Type     | Drive Partition Mount Directory | Data Catalog      | Usage                                                                                                                                                                                                     |
|--------------------|---------------------------------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Metadata partition | /srv/BigData/data1              | dbdata_om         | Directory for storing OMS database data. If Manager is to be installed on two nodes, both two nodes contain this partition directory.                                                                     |
|                    |                                 | LocalBackup       | If the cluster data is backed up in <b>LocalDir</b> , the backup data is stored in the directory by default. If Manager is to be installed on two nodes, both two nodes contain this partition directory. |
|                    |                                 | containers        | Node where the WebContainer role is, which stores the WAR files and configuration sets of container BLU applications                                                                                      |
|                    |                                 | doris/fe          | Directory for storing Doris metadata                                                                                                                                                                      |
|                    | /srv/BigData/data2              | journalnode       | Node where the JournalNode role of HDFS is, which stores the JournalNode metadata of HDFS.                                                                                                                |
|                    |                                 | dbdata_service    | Node where the DBServer role of DBService is, and the database directory of DBService                                                                                                                     |
|                    |                                 | iotdb/iotdbserver | Directory for storing IoTDB metadata                                                                                                                                                                      |
|                    |                                 | iotdb/confignode  | Directory for storing metadata of the IoTDB ConfigNode role                                                                                                                                               |
|                    | /srv/BigData/data3              | namenode          | Node where the NameNode role of HDFS is, which stores NameNode data.                                                                                                                                      |
|                    |                                 | iotdb/iotdbserver | Directory for storing IoTDBServer log data                                                                                                                                                                |
|                    | /srv/BigData/data4              | zookeeper         | Node where the quorumpeer role of ZooKeeper is, which stores ZooKeeper data.                                                                                                                              |
|                    |                                 | hetuengine/qas    | Node where the QAS role of HetuEngine is, which stores QAS data.                                                                                                                                          |

| Partition Type         | Drive Partition Mount Directory | Data Catalog                                                      | Usage                                                                                                                                                         |
|------------------------|---------------------------------|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Service data partition | /srv/BigData/dataN              | <ul style="list-style-type: none"><li>• dn</li><li>• nm</li></ul> | Directory for storing DataNode data and intermediate MapReduce data                                                                                           |
|                        |                                 | kafka-logs                                                        | Directory for storing Kafka broker data.                                                                                                                      |
|                        |                                 | clickhouse<br>clickhouse_path                                     | Directory for storing ClickHouse database data<br>The <b>clickhouse_path</b> directory storing ClickHouse metadata exists only in the <b>data1</b> directory. |
|                        |                                 | iotdb/<br>iotdbserver                                             | Directory for storing IoTDB service data                                                                                                                      |
|                        |                                 | doris/be                                                          | Directory for storing Doris database data                                                                                                                     |

 NOTE

- The metadata partition directory uses a maximum of four disks (data1 to data4). The metadata directories are mapped to directories from **/srv/BigData/data1** to **/srv/BigData/data4** in sequence according to [Table 9-3](#). If only three data disks are mounted to the current node, the directories under data4 and data2 are combined. If only two data disks are mounted, the directories under data3 and data1 are combined, and those under data4 and data2 are combined.

For example, if the ZooKeeper node has four data disks, the ZooKeeper data directory is **/srv/BigData/data4/zookeeper**. If the node has only three data disks, the ZooKeeper data directory is **/srv/BigData/data2/zookeeper**.

- The mapping rules of service data directories are as follows:

For HDFS, Kafka, ClickHouse, and IoTDB, mount points that comply with the **/srv/BigData/dataN** directory are automatically identified as data directories based on the number of mounted disks on the current node.

For example, if disks are mounted to the **/srv/BigData/data1** to **/srv/BigData/data3** directories, the DataNode data directories are **/srv/BigData/data1/dn**, **/srv/BigData/data2/dn**, and **/srv/BigData/data3/dn**, and the Kafka data directories are **/srv/BigData/data1/kafka-logs**, **/srv/BigData/data2/kafka-logs**, and **/srv/BigData/data3/kafka-logs**.

## 9.9 Does an MRS Cluster Support System Reinstallation?

### Question:

Does an MRS Cluster Support System Reinstallation?

**Answer:**

An MRS cluster does not support system reinstallation.

## **9.10 Can I Change the OS of an MRS Cluster?**

**Question:**

Can I Change the OS of an MRS Cluster?

**Answer:**

The OS of an MRS cluster cannot be changed.



# 10 Component Management

## 10.1 Can I Delete Components Installed in an MRS Cluster?

You cannot delete any component from a created MRS cluster of MRS 3.1.0. If a component is not required, log in to MRS Manager and stop the component on the **Services** page.

You can add components to MRS 3.1.2-LTS.3 or later custom clusters. For details, see [Managing Services](#).

## 10.2 How Do I View the Configuration File Directory of Each Component?

The configuration file paths of commonly used components are as follows:

| Component  | Configuration File Directory                                      |
|------------|-------------------------------------------------------------------|
| ClickHouse | <i>Client installation directory/ClickHouse/clickhouse/config</i> |
| Flink      | <i>Client installation directory/Flink/flink/conf</i>             |
| Flume      | <i>Client installation directory/fusioninsight-flume-xxx/conf</i> |
| HBase      | <i>Client installation directory/HBase/hbase/conf</i>             |
| HDFS       | <i>Client installation directory/HDFS/hadoop/etc/hadoop</i>       |
| Hive       | <i>Client installation directory/Hive/config</i>                  |
| Hudi       | <i>Client installation directory/Hudi/hudi/conf</i>               |
| Kafka      | <i>Client installation directory/Kafka/kafka/config</i>           |


| Component | Configuration File Directory                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Loader    | <ul style="list-style-type: none"><li>• <i>Client installation directory/Loader/loader-tools-xxx/loader-tool/conf</i></li><li>• <i>Client installation directory/Loader/loader-tools-xxx/schedule-tool/conf</i></li><li>• <i>Client installation directory/Loader/loader-tools-xxx/shell-client/conf</i></li><li>• <i>Client installation directory/Loader/loader-tools-xxx/sqoop-shell/conf</i></li></ul> |
| Oozie     | <i>Client installation directory/Oozie/oozie-client-xxx/conf</i>                                                                                                                                                                                                                                                                                                                                           |
| Spark2x   | <i>Client installation directory/Spark2x/spark/conf</i>                                                                                                                                                                                                                                                                                                                                                    |
| Yarn      | <i>Client installation directory/Yarn/config</i>                                                                                                                                                                                                                                                                                                                                                           |
| ZooKeeper | <i>Client installation directory/Zookeeper/zookeeper/conf</i>                                                                                                                                                                                                                                                                                                                                              |

## 10.3 Will Upper-Layer Services Be Affected If the Hive Service Status Is Partially Healthy?

### Question

Will Upper-Layer Services Be Affected If the Hive Service Status Is Partially Healthy?

**Figure 10-1** Hive service status



### Answer:

If the cluster service status is **Subhealthy**, some enhancements of the service cannot work properly.

Log in to FusionInsight Manager to check the status of other services in the cluster. In this example, the HBase service is stopped. If HBase is installed but is abnormal, Hive, Spark, and Loader are in the **Subhealthy** state.

In this case, upper-layer services depending on the Hive service are not affected. After you restart HBase, the Hive service recovers.


## 10.4 How Can I Obtain the IP Address and Port Number of a ZooKeeper Instance?

You can obtain the IP address and port number of a ZooKeeper instance through the MRS console or FusionInsight Manager.

Method 1: Obtaining the IP address and port number of a ZooKeeper through the MRS console

1. On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.
2. Click the **Components** tab and choose **ZooKeeper**. On the displayed page, click **Instances** to view the business IP address of a ZooKeeper instance.

Figure 10-2 ZooKeeper addresses



The screenshot shows the MRS console interface. At the top, there is a navigation bar with tabs: Dashboard, Nodes, Components (selected), Alarms, Files, Jobs, Tenants, Bootstrap Actions, and Tags. Below the navigation bar, there is a message: "You can manage all services installed in the cluster. Learn more". Underneath, there is a section for "Service ZooKeeper / Instances". There are three tabs: Service Status, Instances (selected), and Service Configuration. Below the tabs, there is a "More" button. The main content area displays a table with the following columns: Role, Host Name, OM IP Address, and Business IP Address. The table contains three rows of data, all with the role "quorumpeer".


| <input type="checkbox"/> Role       | Host Name    | OM IP Address | Business IP Address |
|-------------------------------------|--------------|---------------|---------------------|
| <input type="checkbox"/> quorumpeer | node-master1 | 192.          | 192.                |
| <input type="checkbox"/> quorumpeer | node-master2 | 192.          | 192.                |
| <input type="checkbox"/> quorumpeer | node-master3 | 192.          | 192.                |

3. Click the **Service Configuration** tab. On the displayed page, search for the **clientPort** parameter to view the port number of the ZooKeeper instance.

Method 2: Obtaining the IP address and port number of a ZooKeeper through FusionInsight Manager

1. Log in to FusionInsight Manager.
2. Perform the following operations to obtain the IP address and port number of a ZooKeeper instance.
  - For clusters of MRS 3.x or earlier
    - i. Choose **Services > ZooKeeper**. On the displayed page, click the **Instance** tab to view the business IP address of a ZooKeeper instance.

Figure 10-3 ZooKeeper addresses



The screenshot shows the FusionInsight Manager interface. It displays a table with the following columns: Role, Host Name, OM IP Address, and Business IP Address. The table contains three rows of data, all with the role "quorumpeer".

| <input type="checkbox"/> Role       | Host Name    | OM IP Address | Business IP Address |
|-------------------------------------|--------------|---------------|---------------------|
| <input type="checkbox"/> quorumpeer | node-master1 | 192.          | 192.                |
| <input type="checkbox"/> quorumpeer | node-master2 | 192.          | 192.                |
| <input type="checkbox"/> quorumpeer | node-master3 | 192.          | 192.                |

- ii. Click the **Service Configuration** tab. On the displayed page, search for the **clientPort** parameter to view the port number of the ZooKeeper instance.
- For clusters of MRS 3.x or later
  - i. Choose **Cluster > Services > ZooKeeper**. On the displayed page, click the **Instance** tab to view the business IP address of a ZooKeeper instance.
  - ii. Click the **Configurations** tab. On the displayed page, search for the **clientPort** parameter to view the port number of the ZooKeeper instance.

# 11 Job Management

---

## 11.1 What Types of Spark Jobs Can Be Submitted in a Cluster?

### Question:

What Types of Spark Jobs Can Be Submitted in a Cluster?

### Answer:

MRS clusters support Spark jobs submitted in Spark, Spark Script, or Spark SQL mode.

## 11.2 What Should I Do If Error 408 Is Reported When an MRS Node Accesses OBS?

### Question:

What Should I Do If Error 408 Is Reported When an MRS Node Accesses OBS?

### Answer:

If error 408 is reported when an MRS node accesses OBS, you need to change the OBS domain name.

Change the OBS domain name to the domain name suffixed with myhuaweicloud.com.

## 11.3 How Do I Enable Different Service Programs to Use Different Yarn Queues?

### Question:

How Do I Enable Different Service Programs to Use Different Yarn Queues?

### Answer:

Create a tenant on FusionInsight Manager and bind different cluster service users to different tenants.

### Procedure

- Step 1** Log in to FusionInsight Manager and choose **Tenant Resources**.
- Step 2** In the tenant list on the left, select a parent tenant and click . On the page for adding a sub-tenant, set properties for the sub-tenant according to [Table 11-1](#).

**Table 11-1** Sub-tenant parameters

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster       | Cluster to which the parent tenant belongs.                                                                                                                                                                                                                                                                                                                                                         |
| Parent Tenant | Name of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                          |
| Name          | <ul style="list-style-type: none"><li>Name of the current tenant. The value consists of 3 to 50 characters, including digits, letters, and underscores (_).</li><li>Plan a sub-tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or Yarn queue that exists in the current cluster.</li></ul>                                                |
| Tenant Type   | <p>Whether the tenant is a leaf tenant.</p> <ul style="list-style-type: none"><li>When <b>Leaf Tenant</b> is selected, the current tenant is a leaf tenant and no sub-tenant can be added.</li><li>When <b>Non-leaf Tenant</b> is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant. However, the tenant depth cannot exceed 5 levels.</li></ul> |

| Parameter                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Computing Resource                     | <p>Dynamic compute resources for the current tenant.</p> <ul style="list-style-type: none"><li>When <b>Yarn</b> is selected, the system automatically creates a queue in YARN and the queue is named the same as the sub-tenant name.<ul style="list-style-type: none"><li>A leaf tenant can directly submit jobs to the queue.</li><li>A non-leaf tenant cannot directly submit jobs to the queue. However, YARN adds an extra queue (hidden) named <b>default</b> for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue.</li></ul></li><li>If <b>Yarn</b> is not selected, the system does not automatically create a queue.</li></ul>                           |
| Default Resource Pool Capacity (%)     | Percentage of computing resources used by the current tenant. The base value is the total resources of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Default Resource Pool Max Capacity (%) | Maximum percentage of computing resources used by the current tenant. The base value is the total resources of the parent tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Storage Resource                       | <p>Storage resources for the current tenant.</p> <ul style="list-style-type: none"><li>When <b>HDFS</b> is selected, the system automatically creates a folder named after the sub-tenant in the HDFS parent tenant directory.</li><li>When <b>HDFS</b> is not selected, the system does not automatically allocate storage resources.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                   |
| Quota                                  | Quota for files and directories.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Space Quota                            | <p>Quota for the HDFS storage space used by the current tenant.</p> <ul style="list-style-type: none"><li>If the unit is set to <b>MB</b>, the value ranges from <b>1</b> to <b>8796093022208</b>. If the unit is set to <b>GB</b>, the value ranges from <b>1</b> to <b>8589934592</b>.</li><li>This parameter indicates the maximum HDFS storage space that can be used by the tenant, but not the actual space used.</li><li>If its value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk.</li><li>If this quota is greater than the quota of the parent tenant, the actual storage space does not exceed the quota of the parent tenant.</li></ul> |

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage Path | <p>HDFS storage directory for the tenant.</p> <ul style="list-style-type: none"><li>• The system automatically creates a folder named after the sub-tenant name in the directory of the parent tenant by default. For example, if the sub-tenant is <b>ta1s</b> and the parent directory is <b>/tenant/ta1</b>, the storage path for the sub-tenant is then <b>/tenant/ta1/ta1s</b>.</li><li>• The storage path is customizable in the parent directory.</li></ul> |
| Description  | Description of the current tenant.                                                                                                                                                                                                                                                                                                                                                                                                                                 |

 **NOTE**

Roles, computing resources, and storage resources are automatically created when tenants are created.

- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System > Permission > Role**. The role name is in the format of *Tenant name\_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and bind the user to the role of the tenant.
- The sub-tenant can further allocate the resources of its parent tenant. The sum of the resource percentages of direct sub-tenants under a parent tenant at each level cannot exceed 100%. The sum of the computing resource percentages of all level-1 tenants cannot exceed 100%.

**Step 3** Check whether the current tenant needs to be associated with resources of other services.

- If yes, go to [Step 4](#).
- If no, go to [Step 5](#).

**Step 4** Click **Associate Service** to configure other service resources used by the current tenant.

1. Set **Services** to **HBase**.
2. Set **Association Type** as follows:
  - **Exclusive** indicates that the service resources are used by the tenant exclusively and cannot be associated with other tenants.
  - **Shared** indicates that the service resources can be shared with other tenants.



 **NOTE**

- Only HBase can be associated with a new tenant. However, HDFS, HBase, and Yarn can be associated with existing tenants.
- To associate an existing tenant with service resources, click the target tenant in the tenant list, switch to the **Service Associations** page, and click **Associate Service** to configure resources to be associated with the tenant.
- To disassociate an existing tenant from service resources, click the target tenant in the tenant list, switch to the **Service Associations** page, and click **Delete** in the **Operation** column. In the displayed dialog box, select **I have read the information and understand the impact** and click **OK**.

3. Click **OK**.

**Step 5** Click **OK**. Wait until the system displays a message indicating that the tenant is successfully created.

----End

## 11.4 What Should I Do If a Job Fails to Be Submitted and the Error Is Related to OBS?

### Question:

Why Does a Job Fail to Be Submitted and the Error Is Related to OBS?

### Answer:

OBS errors are generally related to read and write. You can perform the following steps to locate the fault:

**Step 1** View the failed job details on the **Jobs** page and check the actual job ID. If the job ID is empty, go to [Step 4](#).

#### View Details

|                 |                                      |
|-----------------|--------------------------------------|
| Job Type        | SparkSql                             |
| Job ID          | 9b6cdf6a-e381-4a12-98a4-3e47eccf8af1 |
| Launcher Job ID | application_1719654297461_0006       |
| Actual Job ID   | --                                   |
| Submit Time     | Jul 01, 2024 18:07:53 GMT+08:00      |
| Started         | Jul 01, 2024 18:07:54 GMT+08:00      |
| Ended           | Jul 01, 2024 18:11:00 GMT+08:00      |

**Step 2** Log in to FusionInsight Manager, choose **Cluster** > **Services** > **Yarn** > **Overview** > **ResourceManager Web UI**, and search for **logs** by the job ID.

| Show 20 ▾ entries | Attempt ID                          | Started                       | Node                         | Logs | Nodes blacklisted by the app | Nodes blacklisted by the system |
|-------------------|-------------------------------------|-------------------------------|------------------------------|------|------------------------------|---------------------------------|
|                   | appattempt_171497451991_0003_000002 | Mon May 6 22:41:13 +0800 2024 | http://node-master3yZr26006  | Logs | 0                            | 0                               |
|                   | appattempt_171497451991_0003_000001 | Mon May 6 22:38:04 +0800 2024 | http://node-master2KxRP26006 | Logs | 0                            | 0                               |

**Step 3** Check whether the following error information is displayed in the job log: If the following error message is displayed, bind an OBS agency to the cluster and submit the job again. For details, see [Configuring a Storage-Compute Decoupled Cluster \(Agency\)](#).

```
2024-05-06 16:49:58,842 INFO [block-manager-storage-async-thread-pool-0] | Removing RDD 3 | org.apache.spark.storage.BlockManager$LogInfo(Logging.scala:61)
2024-05-06 16:49:58,846 INFO [Driver] | Running: SuperSQL : [create external table src_wordcount11(line string) row format delimited fields terminated by "\n" stored as textfile location 'obs://mrs-lmq-c
2024-05-06 16:49:58,860 INFO [Driver] | It took 50 ms to build plan. | org.apache.spark.sql.execution.QueryExecution$LogInfo(Logging.scala:31)
2024-05-06 16:49:59,248 INFO [Driver] | Initialising http4i with default configuration | com.huawei.mrs.lasshttpclient.<init>()lassHttpClient.java:49)
2024-05-06 16:50:00,223 INFO [Driver] | The user info: lmqcloud (auth:SIMPLE) Agency mapping properties: {} | com.huawei.mrs.AgencyMappingLoader.<init>()AgencyMappingLoader.java:68)
2024-05-06 16:50:00,407 WARN [Driver] | Get ecs meta data error: org.apache.http.client.HttpResponseException: status code: 500, reason phrase: Exception when fetch security key from meta server | com.hu
2024-05-06 16:50:00,408 ERROR [Driver] | Failed to get security from mrs meta, will not get security key from ECS directly | com.huawei.mrs.MrsObsCredentialsProvider.getSecurityKey(MrsObsCredentialsProvide
2024-05-06 16:50:00,408 ERROR [Driver] | Initialising OBSFileSystem fail | org.apache.hadoop.fs.obs.OBSFileSystem.initialize(OBSFileSystem.java:545)
com.obs.services.exception.ObsException: Exception when get temporary security key
at com.huawei.mrs.MrsObsCredentialsProvider.cacheLimitedTimeSecurityKey(MrsObsCredentialsProvider.java:203) ~[mrs-obs-provider-8.3.1-331.r10.jar:?]
at com.huawei.mrs.MrsObsCredentialsProvider.getSecurityKey(MrsObsCredentialsProvider.java:195) ~[mrs-obs-provider-8.3.1-331.r10.jar:?]
at com.obs.services.ObsClient.<init>(ObsClient.java:161) ~[hadoop-huaweicloud-3.1.1-hw-53.7.jar:?]
at org.apache.hadoop.fs.obs.DefaultObsClientFactory.createObsClient(DefaultObsClientFactory.java:190) ~[hadoop-huaweicloud-3.1.1-hw-53.7.jar:?]
at org.apache.hadoop.fs.obs.OBSFileSystem.initialize(OBSFileSystem.java:431) ~[hadoop-huaweicloud-3.1.1-hw-53.7.jar:?]
at org.apache.hadoop.fs.FileSystem.createFileSystem(FileSystem.java:3564) ~[hadoop-common-3.3.1-h0.cbu.mrs.331.r10.jar:?]
at org.apache.hadoop.fs.FileSystem.access$300(FileSystem.java:174) ~[hadoop-common-3.3.1-h0.cbu.mrs.331.r10.jar:?]
at org.apache.hadoop.fs.FileSystem$Cache.getInternal(FileSystem.java:3672) ~[hadoop-common-3.3.1-h0.cbu.mrs.331.r10.jar:?]
at org.apache.hadoop.fs.FileSystem$Cache.get(FileSystem.java:3623) ~[hadoop-common-3.3.1-h0.cbu.mrs.331.r10.jar:?]
at org.apache.hadoop.fs.FileSystem.get(FileSystem.java:347) ~[hadoop-common-3.3.1-h0.cbu.mrs.331.r10.jar:?]
```

**Step 4** View the logs of to the Launcher job. On the **Jobs** page, view the logs to check whether the **stdout** or **stderr** log of the Launcher job contains error information. Locate the fault based on the error logs.

----End

## 11.5 Can I Run Multiple Spark Tasks at the Same Time After the Minimum Tenant Resources of an MRS Cluster Is Changed to 0?

### Question:

Can I Run Multiple Spark Tasks at the Same Time After the Minimum Tenant Resources of an MRS Cluster Is Changed to 0?

### Answer:

You can run only one Spark task at a time after the minimum tenant resources of an MRS cluster is changed to 0.

## 11.6 What Should I Do If Job Parameters Separated By Spaces Cannot Be Identified?

### Question:

Will Parameters Be Identified If Job Submission Parameters Are Separated by Spaces?

### Answer:

Use spaces to separate parameters. To prevent parameters from being saved as plaintext, add an at sign (@) before parameters, for example, **@password=XXXXXX**.

## 11.7 What Are the Differences Between the Client Mode and Cluster Mode of Spark Jobs?

You need to understand the concept ApplicationMaster before understanding the essential differences between Yarn-client and Yarn-cluster.

In Yarn, each application instance has an ApplicationMaster process, which is the first container started by the application. It interacts with ResourceManager and requests resources. After obtaining resources, it instructs NodeManager to start containers. The essential difference between the Yarn-cluster and Yarn-client modes lies in the ApplicationMaster process.


In Yarn-cluster mode, Driver runs in ApplicationMaster, which requests resources from Yarn and monitors the running status of a job. After a user submits a job, the client can be stopped and the job continues running on Yarn. Therefore, the Yarn-cluster mode is not suitable for running interactive jobs.

In Yarn-client mode, ApplicationMaster requests only Executor from Yarn. The client communicates with the requested containers to schedule tasks. Therefore, the client cannot be stopped.

## 11.8 How Do I View MRS Job Logs?


**Step 1** On the **Jobs** page of the MRS console, you can view logs of each job, including launcherJob and realJob logs.

- Generally, error logs are printed in **stderr** and **stdout** for launcherJob jobs, as shown in the following figure:



The screenshot shows a log viewer interface with tabs for different log types: container-localizer-syslog, directory.info, launch\_container.sh, prelaunch.err, prelaunch.out, stderr, stdout, and syslog. The 'stderr' tab is selected. The log content shows a sequence of Java stack traces. The first line of the log is highlighted in red and contains the error message: 'Output directory hdfs://hacluster/user/mr-0610-100 already exists'. The subsequent lines show stack traces from org.apache.hadoop.mapred.FileAlreadyExistsException through various Hadoop and Java classes.

- You can view realJob logs on the ResourceManager web UI provided by the Yarn service on MRS Manager.



**Step 2** Log in to the Master node of the cluster to obtain the job log files in **Step 1**. The HDFS path is `/tmp/logs/{submit_user}/logs/{application_id}`.

**Step 3** After the job is submitted, if the job application ID cannot be found on the Yarn web UI, the job fails to be submitted. You can log in to the active Master node of the cluster and view the job submission process log `/var/log/executor/logs/exe.log`.

----End

## 11.9 What Can I Do If the System Displays a Message Indicating that the Current User Does Not Exist on Manager When I Submit a Job?

**Analysis:** If IAM synchronization is not performed when a job is submitted to a cluster in security mode, the error message "The current user does not exist on MRS Manager. Grant the user sufficient permissions on IAM and then perform IAM user synchronization on the Dashboard tab page." is displayed.

**Solution:** Before you submit the job, go to the **Dashboard** page, and click **Synchronize** on the right side of **IAM User Sync**.

## 11.10 What Can I Do If LauncherJob Fails to Be Executed and the Error Message "jobPropertiesMap is null" Is Displayed?

**Question:** What should I do if the Launcher job fails because the user who submits the job does not have the write permission on the `hdfs /mrs/job-properties` directory?

**Answer:** This problem has been fixed in the 2.1.0.6 patch. Additionally, you can also grant the write permission on the `/mrs/job-properties` directory to the synchronized user who submits the job on MRS Manager.

## 11.11 What Should I Do If the Flink Job Status on the MRS Console Is Inconsistent with That on Yarn?

**Question:** What should I do if the Flink job status on the MRS console is inconsistent with that on Yarn?

**Answer:** To save storage space, the Yarn configuration item **yarn.resourcemanager.max-completed-applications** is modified to reduce the number of historical job records stored on Yarn. Flink jobs are long-term jobs. The real job is still running on Yarn, but the launcher job has been deleted. As a result, the launcher job cannot be found on Yarn, and the job status fails to be updated. This problem is fixed in the 2.1.0.6 patch.

**Workaround:** Stop the jobs where launcher job cannot be found. The statuses of jobs submitted later will be updated.

## 11.12 What Can I Do If a SparkStreaming Job Fails After Running for Dozens of Hours and Error 403 Is Reported for OBS Access?

**Question:** What can I do if a SparkStreaming job fails after running for dozens of hours and error 403 is reported for OBS access?

**Answer:** When a user submits a job that needs to read and write OBS, the job submission program adds the temporary access key (AK) and secret key (SK) for accessing OBS by default. The temporary AK and SK have expiration time.

If you want to run long-term jobs such as Flink and SparkStreaming jobs, you can enter the AK and SK in **Service Parameter** to prevent jobs execution failure due to key expiration.

## 11.13 What Should I Do If Error Message "java.io.IOException: Connection reset by peer" Is Displayed During the Execution of a Spark Job?

### Symptom

The execution of a Spark job cannot finish and error message "java.io.IOException: Connection reset by peer" is displayed.

### Solution

Add the **executor.memory Overhead** parameter to the parameters for submitting a job.

## 11.14 What Should I Do If the Error Message "requestId=XXX" Is Displayed When a Spark Job Accesses OBS?

### Symptom

Error message "requestId=4971883851071737250" is displayed when a Spark job accesses OBS.

### Solution

Log in to the node where the Spark client is located, go to the **conf** directory, and change the value of the **fs.obs.metrics.switch** parameter in the **core-site.xml** configuration file to **false**.

## 11.15 What Should I Do If the Error Message "UnknownScannerException" Is Displayed for Spark Jobs?

### Symptom

Spark jobs run slowly. Warning information is printed in run logs, and the error cause is **UnknownScannerException**.

### Solution

Before running a Spark job, adjust the value of **hbase.client.scanner.timeout.period** (for example, from 60 seconds to 120 seconds).

Log in to FusionInsight Manager and choose **Cluster > Services > HBase**. Click **Configurations** then **All Configurations**, search for **hbase.client.scanner.timeout.period**, and change its value to **120000** (unit: ms).

## 11.16 What Can I Do If DataArts Studio Occasionally Fails to Schedule Spark Jobs?

### Symptom

DataArtsStudio occasionally fails to schedule Spark jobs and the rescheduling also fails. The following error information is displayed:

Caused by: org.apache.spark.SparkException: Application application\_1619511926396\_2586346 finished with failed status

## Solution

Log in to the node where the Spark client is located as user **root** and increase the value of the **spark.driver.memory** parameter in the **spark-defaults.conf** file.

## 11.17 What Should I Do If a Flink Job Fails to Execute and the Error Message "java.lang.NoSuchFieldError: SECURITY\_SSL\_ENCRYPT\_ENABLED" Is Displayed?

### Symptom

A Flink job fails to be executed and the following error message is displayed:

Caused by: java.lang.NoSuchFieldError: SECURITY\_SSL\_ENCRYPT\_ENABLED

### Solution

The third-party dependency package in the customer code conflicts with the cluster package. As a result, the job fails to be submitted to the MRS cluster. You need to modify the dependency package, set the scope of the open source Hadoop package and Flink package in the POM file to **provide**, and pack and execute the job again.

## 11.18 What Should I Do If Submitted Yarn Jobs Cannot Be Viewed on the Web UI?

**Question:** Why a created Yarn job cannot be viewed when I log in to Manager as the **admin** user?

**Answer:**

- The **admin** user is a user on the cluster management page. Check whether the user has the **supergroup** permission. Generally, only the user with the **supergroup** permission can view jobs.
- Log in to Yarn as the user who submits the job. Do not use the **admin** user to view the job.

## 11.19 What Can I Do If launcher-job Is Terminated by Yarn When a Flink Task Is Submitted?

### Symptom

Launcher-job is terminated by Yarn due to insufficient heap size when a Flink task is submitted on the management plane.

### Solution

Increase the heap size of the launcher-job queue.

1. Log in to the active OMS node as user **omm**.
2. Change the value of **job.launcher.resource.memory.mb** in **/opt/executor/webapps/executor/WEB-INF/classes/servicebroker.xml** to **2048**.
3. Run the **sh /opt/executor/bin/restart-executor.sh** command to restart the executor process.

## 11.20 What Should I Do If the Error Message "slot request timeout" Is Displayed When I Submit a Flink Job?

### Symptom

When a Flink job is submitted, JobManager is started successfully. However, TaskManager remains in the starting state until timeout. The following error information is displayed:

```
org.apache.flink.runtime.jobmanager.scheduler.NoResourceAvailableException: Could not allocate the required slot within slot request timeout. Please make sure that the cluster has enough resources
```

### Possible Causes

1. The resources in the YARN queue are insufficient. As a result, TaskManager fails to start.
2. Your JAR files conflict with those in the environment. You can execute the WordCount program to determine whether the issue occurs.
3. If the cluster is in security mode, the SSL certificate of Flink may be incorrectly configured or has expired.

### Solution

1. Add resources to the YARN queue.
2. Exclude the Flink and Hadoop dependencies in your JAR files so that Flink and Hadoop can depend only on the JAR files in the environment.
3. Reconfigure the SSL certificate of Flink. For details, see [Using Flink from Scratch](#).

## 11.21 FAQs About Importing and Exporting Data Using DistCP Jobs

- **Question:** Does a DistCP job check data consistency during data import and export?  
**Answer:** No. DistCP jobs only copy data but do not modify it.
- **Question:** When data is exported from a DistCP job, if some files already exist in OBS, how will the job process the files?  
**Answer:** DistCP jobs will overwrite the files in OBS.




## 11.22 How Do I View SQL Statements of Hive Jobs on the Yarn Web UI?

**Step 1** Log in to FusionInsight Manager as a service user.

**Step 2** Choose **Cluster > Services > Yarn**. Click **ResourceManager(xxx,Active)** next to **ResourceManager WebUI**.


**Step 3** On the YARN web UI, click the ID of the job to be viewed.



The screenshot shows the Hadoop YARN web UI. On the left, there is a navigation menu with options like Cluster, Nodes, Labels, Applications, and Tools. The main area displays a table of applications. The application 'application\_0004' is highlighted in red. The table columns include ID, User, QueueUser, Name, Application Type, Application Tags, Queue, Application Priority, StartTime, LaunchTime, and FinishTime.


| ID               | User | QueueUser | Name                                                   | Application Type | Application Tags | Queue   | Application Priority | StartTime                      | LaunchTime                     | FinishTime                     |
|------------------|------|-----------|--------------------------------------------------------|------------------|------------------|---------|----------------------|--------------------------------|--------------------------------|--------------------------------|
| application_0005 |      |           | Flink session cluster                                  |                  |                  | default | 0                    | Sat Oct 8 16:01:40 +0800 2022  | Sat Oct 8 16:01:40 +0800 2022  | Sat Oct 8 16:01:54 +0800 2022  |
| application_0004 |      |           | Insert into user_info partition(year='...B') (Stage-1) |                  |                  | default | 0                    | Thu Sep 29 09:47:50 +0800 2022 | Thu Sep 29 09:47:50 +0800 2022 | Thu Sep 29 09:48:18 +0800 2022 |
| application_0003 |      |           | Insert into user_info partition(year='...A') (Stage-1) |                  |                  | default | 0                    | Thu Sep 29 09:46:07 +0800 2022 | Thu Sep 29 09:46:08 +0800 2022 | Thu Sep 29 09:46:39 +0800 2022 |
| application_0002 |      |           | Insert into table test (name) values ('12') (Stage-1)  |                  |                  | default | 0                    | Thu Sep 29 09:23:28 +0800 2022 | Thu Sep 29 09:23:29 +0800 2022 | Thu Sep 29 09:23:55 +0800 2022 |

**Step 4** Click **ApplicationMaster** or **History** next to **Tracking URL**.



The screenshot shows the Hadoop YARN web UI for application 'application\_0004'. The left navigation menu is visible. The main area displays the application details, including the User (hiveuser), Application Name (Insert into user\_info partition(year='...B') (Stage-1)), Application Type, Application Tags, Application Priority, YarnApplicationState (FINISHED), Queue (default), FinalStatus Reported by AM (SUCCEEDED), Started, Launched, Finished, Elapsed (28sec), Tracking URL (History), Log Aggregation Status (SUCCEEDED), Application Timeout (Remaining Time), Diagnostics, Unmanaged Application (false), Application Node Label expression (<Not set>), and AM container Node Label expression (<DEFAULT\_PARTITION>).

**Step 5** Choose **Configuration** in the left navigation pane and search for the **hive.query.string** parameter in the upper right corner to query the corresponding HiveSQL.



The screenshot shows the Hadoop YARN web UI for the configuration of MapReduce Job 'application\_0004'. The left navigation menu is visible. The main area displays a table of configuration parameters. The parameter 'hive.query.string' is highlighted. The table columns include key, value, and source chain.

| key                                  | value | source chain                |
|--------------------------------------|-------|-----------------------------|
| _hive.hadoop.session.path            |       | job.xml <- programmatically |
| _hive.local.session.path             |       | job.xml <- programmatically |
| _hive.tmp.table.space                |       | job.xml <- programmatically |
| adfs.feature.ownerandgroup.enableupn |       | job.xml <- core-default.xml |
| adfs.http.timeout                    |       | job.xml <- core-default.xml |

----End

## 11.23 How Do I View Logs of a Specified Yarn Task?

- Step 1** Log in to the active Master node as user **root**.
- Step 2** Run the following command to initialize environment variables:
- ```
source Client installation directory/bigdata_env
```
- Step 3** If Kerberos authentication is enabled for the cluster, run the following command to authenticate the user. If Kerberos authentication is disabled, skip this step.
- ```
kinit MRS cluster user
```
- Step 4** Run the following command to obtain the logs of a specified task:
- ```
yarn logs -applicationId application_ID of the job you want to view
```
- End

## 11.24 What Should I Do If a HiveSQL/HiveScript Job Fails to be Submitted After Hive Is Added?

This issue occurs because the **MRS CommonOperations** permission bound to the user group to which the user who submits the job belongs does not include the Hive permission after being synchronized to Manager. To solve this issue, perform the following operations:

1. Add the Hive service.
2. Log in to the IAM console and create a user group. The policy bound to the user group is the same as that of the user group to which the user who submits the job belongs.
3. Add the user who submits the job to the new user group.
4. Refresh the cluster details page on the MRS console. The status of IAM user synchronization is **Not synchronized**.
5. Click **Synchronize** on the right of **IAM User Sync**. Go back to the previous page. In the navigation pane on the left, choose **Operation Logs** and check whether the user is changed.
  - If yes, submit the Hive job again.
  - If no, check whether all the preceding operations are complete.
    - If yes, contact an O&M engineer.
    - If no, submit the Hive job after the preceding operations are complete.

## 11.25 Where Are the Execution Logs of Spark Jobs Stored?

- Logs of unfinished Spark jobs are stored in the **/srv/BigData/hadoop/data1/nm/containerlogs/** directory on the Core node.

- Logs of finished Spark jobs are stored in the `/tmp/logs/Username/logs` directory of the HDFS.

## 11.26 What Should I Do If an Alarm Indicating Insufficient Memory Is Reported During Spark Task Execution?

### Symptom

When a Spark task is executed, an alarm indicating insufficient memory is reported. The alarm ID is 18022. As a result, no available memory can be used.

### Procedure

Set the executor parameters in the SQL script to limit the number of cores and memory of an executor.

For example, the configuration is as follows:

```
set hive.execution.engine=spark;  
set spark.executor.cores=2;  
set spark.executor.memory=4G;  
set spark.executor.instances=10;
```

Change the values of the parameters as required.

## 11.27 What Can I Do If an Alarm is Generated Because the NameNode Is not Restarted on Time After the hdfs-site.xml File Is Modified?

**Question:** What can I do if an alarm is generated because the NameNode is not restarted on time after the `hdfs-site.xml` file is modified?

**Answer:** If you do not restart the NameNode process after modifying the `dfs.namenode.checkpoint.period` parameter, an alarm may be falsely reported. In this case, you need to restart the NameNode process as soon as possible.

## 11.28 What Should I Do If It Takes a Long Time for Spark SQL to Access Hive Partitioned Tables Before a Job Starts?

### Symptom

When Spark SQL is used to access Hive partitioned tables stored in OBS, the access speed is slow and a large number of OBS query APIs are called.

Example SQL:

```
select a,b,c from test where b=xxx
```

## Fault Locating

According to the configuration, the task should scan only the partition whose b is xxx. However, the task logs show that the task scans all partitions and then calculates the data whose b is xxx. As a result, the task calculation is slow. In addition, a large number of OBS requests are sent because all files need to be scanned.

By default, the execution plan optimization based on partition statistics is enabled on MRS, which is equivalent to automatic execution of Analyze Table. (The default configuration method is to set **spark.sql.statistics.fallBackToHdfs** to **true**. You can set this parameter to **false**.) After this function is enabled, table partition statistics are scanned during SQL execution and used as cost estimation in the execution plan. For example, small tables identified during cost evaluation are broadcast to each node in the memory for join operations, significantly reducing shuffle time. This function greatly optimizes performance in join scenarios, but increases the number of OBS calls.

## Procedure

Set the following parameter in Spark SQL and then run the SQL statement:

```
set spark.sql.statistics.fallBackToHdfs=false;
```

Alternatively, run the **--conf** command to set this parameter to **false** before startup.

```
--conf spark.sql.statistics.fallBackToHdfs=false
```

## 11.29 How Does the System Select the Queue When an MRS Cluster User Is Associated to Multiple Queues?

### Question

Users A and B share the same user group and permissions in the MRS cluster. They are associated to both queue A and queue B.

Why are tasks of user A always submitted to queue A and tasks of user B always submitted to queue B?

### Answer:

- If a user in the MRS cluster is associated to multiple queues, the system randomly selects a queue for tasks submitted by the user. You can set the queue permission of the user and associate a user to only one queue.
- If you have to associate multiple queues, configure a global user policy to specify the default submission queue for a user.
  - a. Take the MRS 3.x cluster as an example. Log in to Manager and click **Tenant Resources**.
  - b. Choose **Dynamic Resource Plan > Global User Policy > Create Global User Policy**.

**Username:** indicates the user for whom resource scheduling needs to be controlled. Enter an existing username in the current cluster.

**Default Queue:** indicates the queue of the user. Enter the name of an existing queue in the current cluster.

**Figure 11-1** Setting a global user policy

Global User Policy ×

★ Cluster :

MRS Cluster ▼

★ Username :

Max Running Apps :

-1

Max Pending Apps :

-1

Default Queue :

root.default

OK

Cancel

# 12 Performance Tuning

---

## 12.1 How Do I Obtain the Hadoop Pressure Test Tool?

### Question:

How Do I Obtain the Hadoop Pressure Test Tool?

### Answer:

Obtain the Hadoop pressure test tool from the community at <https://github.com/Intel-bigdata/HiBench>.

## 12.2 How Do I Improve the Resource Utilization of Core Nodes in a Cluster?

1. Go to the Yarn service configuration page.
  - For versions earlier than MRS 1.8.10, log in to MRS Manager, choose **Services > Yarn > Service Configuration**, and select **All** from the **Basic** drop-down list.
  - For MRS 1.8.10 or later and MRS 2.x, click the cluster name. On the cluster details page that is displayed, click the **Components** tab and choose **Yarn**. Click the **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 cluster version 3.x or later:  
Log in to FusionInsight Manager. Choose **Cluster > Name of the desired cluster > Services > Yarn > Configurations > All Configurations**.
2. Search for **yarn.nodemanager.resource.memory-mb**, and increase the value based on the actual memory of the cluster nodes.

3. Save the change and restart the affected services or instances.

## 12.3 How Do I Configure the Knox Memory?

**Step 1** Log in to a Master node of the cluster as user **root**.

**Step 2** Run the following command on the Master node to open the **gateway.sh** file:

```
su omm
```

```
vim /opt/knox/bin/gateway.sh
```

**Step 3** Change **APP\_MEM\_OPTS=""** to **APP\_MEM\_OPTS="-Xms256m -Xmx768m"**, save the file, and exit.

**Step 4** Run the following command on the Master node to restart the Knox process:


```
sh /opt/knox/bin/gateway.sh stop
```

```
sh /opt/knox/bin/gateway.sh start
```

**Step 5** Repeat the preceding steps on each Master node.

**Step 6** Run the **ps -ef |grep Knox** command to check the configured memory.

Figure 12-1 Knox memory



```
omm@node-master1E3H1 ~$ ps -ef |grep Knox
omm      11688      1  0 15:42 pts/0    00:00:00 /opt/bigdata/jdk1.8.0_212/bin/java -Djava.library.path=/opt/knox/ext/native -Xms256m -Xmx768m -jar /opt/knox/bin/gateway.jar
omm      29369 11354  0 15:52 pts/0    00:00:00 grep --color=auto Knox
omm@node-master1E3H1 ~$
```

----End

## 12.4 How Do I Adjust the Memory Size of the manager-executor Process?

### Symptom

The **manager-executor** process runs either on the Master1 or Master2 node in the MRS cluster in active/standby mode. This process is used to encapsulate the MRS management and control plane's operations on the MRS cluster, such as job submission, heartbeat reporting, certain alarm reporting, as well as cluster creation, scale-out, and scale-in. When you submit jobs on the MRS management and control plane, the Executor memory may become insufficient as the tasks increase or the number of concurrent tasks increases. As a result, the CPU usage is high and the Executor process experiences out-of-memory (OOM) errors.

### Procedure

1. Log in to either the Master1 or Master2 node as user **root** and run the following command to switch to user **omm**:  
**su - omm**
2. Run the following command to modify the **catalina.sh** script. Specifically, search for **JAVA\_OPTS** in the script, find the configuration items similar to

**JAVA\_OPTS="-Xms1024m -Xmx4096m"**, and change the values of the items to desired ones, and save the modification.

**vim /opt/executor/bin/catalina.sh**

```
JAVA_OPTS="-Xms1024m -Xmx4096m"
JAVA_HOME="${JAVA_HOME}"
LOG4J_PROPERTIES_PATH="${CATALINA_HOME}/lib/log4j.properties"
CATALINA_OPTS="-XX:+PrintGC -XX:+PrintGCDetails -XX:+PrintGCTimeStamps -XX:+PrintGCApplicationStoppedTime \
-XX:+PrintHeapAtGC -Xloggc:/var/log/executor/logs/gc.log -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10 \
-XX:GCLogFileSize=20M -XX:OnOutOfMemoryError='kill -9 %p' -XX:+HeapDumpOnOutOfMemoryError \
-XX:HeapDumpPath=/var/log/executor/logs/executor-dump.hprof"
```

3. The **manager-executor** process only runs on either the Master1 or Master2 node in active/standby mode. Check whether it exists on the node before restarting it.
  - a. Log in to the Master1 and Master2 nodes and run the following command to check whether the process exists. If any command output is displayed, the process exists.

**ps -ef | grep "/opt/executor" | grep -v grep**

```
root@master1:~# ps -ef | grep "/opt/executor" | grep -v grep
root      16854      1  S   java      0 0:00:00 /opt/executor/bin/java -Xms1024m -Xmx4096m -XX:+PrintGC -XX:+PrintGCDetails -XX:+PrintGCTimeStamps -XX:+PrintGCApplicationStoppedTime \
-XX:+PrintHeapAtGC -Xloggc:/var/log/executor/logs/gc.log -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=20M -XX:OnOutOfMemoryError='kill -9 %p' -XX:+HeapDumpOnOutOfMemoryError \
-XX:HeapDumpPath=/var/log/executor/logs/executor-dump.hprof -Dlog4j.configuration=file:/opt/executor/lib/log4j.properties org.apache.catalina.startup.Bootstrap start
```

- b. Run the following command to restart the process:

**sh /opt/executor/bin/shutdown.sh**

**sh /opt/executor/bin/startup.sh**

## 12.5 How Do I Configure Spark Jobs to Automatically Obtain More Resources During Execution?

### Question:

How Do I Configure Spark Jobs to Automatically Obtain More Resources During Execution?

### Answer:

Resources are a key factor that affects Spark execution efficiency. When multiple executors but no tasks are allocated to a long-running service (such as the JDBCServer) and resources of other applications are insufficient, resources are wasted and cannot be scheduled properly.

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.

You can use the following method to enable dynamic resource scheduling. For other related configurations, see [Configuring Dynamic Resource Scheduling in Yarn Mode](#).

Log in to FusionInsight Manager, choose **Cluster > Services > Spark** and choose **Configurations > All Configurations**. Enter **spark.dynamicAllocation.enabled** in the search box and set its value to **true** to enable dynamic resource scheduling.



## 12.6 What Should I Do If the `spark.yarn.executor.memoryOverhead` Setting Does Not Take Effect?

### Symptom

The overhead memory of the executor needs to be adjusted for Spark tasks. The `spark.yarn.executor.memoryOverhead` parameter is set to **4096**. However, the default value **1024** is used to apply for resources during actual computation.

### Fault Locating

In Spark 2.3 and later versions, use the new parameter `spark.executor.memoryOverhead` to set the overhead memory of the executor. If both old and new parameters are set, the value of `spark.yarn.executor.memoryOverhead` does not take effect, and the value of `spark.executor.memoryOverhead` is used.

Same thing happens if you use `spark.driver.memoryOverhead` to set the overhead memory of driver.

### Procedure

Use the new parameter:

```
spark.executor.memoryOverhead=4096
```

## 12.7 How Do I Improve Presto Resource Usage?

### Symptom

Presto tasks take too long to finish due to uneven node resource utilization.

### Procedure

1. Decrease the memory available for the YARN node.
  - a. Log in to Manager, choose **Cluster > Services > Presto > Instances**. Check and record the host names of the nodes where all Presto instances are deployed.
  - b. Choose **Cluster > Services > Yarn > Instances**, click the names of the NodeManger instances on the same node as the Presto instance, click **Instance Configurations** and **All Configurations**, and search for `yarn.nodemanager.resource.memory-mb`. Decrease the value of this parameter. Click **Save**.
  - c. Go back to the **Instances** tab, select the instances whose configurations have been changed, and choose **More > Instance Rolling Restart**.

 NOTE

- **yarn.nodemanager.resource.memory-mb** indicates the total physical memory that can be used by YARN on the node. When the workload is high, YARN and Presto contend for available memory. Set this parameter to a smaller value based on service requirements.
- YARN tasks may behave abnormally during a rolling restart due to repeated attempts. Restart the service during off-peak hours.

2. Increase the value of **xmx** for Presto Work.

Log in to FusionInsight Manager, choose **Cluster > Services > Presto > Configurations > All Configurations**, search for the **JAVA\_OPTS** parameter, and set the **xmx** value of it under Presto Work. Increase the value based on service requirements. Click **Save** and restart the Presto service.

 NOTE

Restarting Presto will cause service unavailability. Restart Presto during off-peak hours.

# 13 Application Development

---

## 13.1 How Do I Get My Data into OBS or HDFS?

MRS can process data in OBS and HDFS. You can get your data into OBS or HDFS as follows:

1. Upload local data to OBS.
  - a. Log in to the OBS console.
  - b. Create a parallel file system named **userdata** on OBS and create the **program**, **input**, **output**, and **log** folders in the file system.
    - i. Choose **Parallel File System** > **Create Parallel File System**, and create a file system named **userdata**.
    - ii. In the OBS file system list, click the file system name **userdata**, choose **Files** > **Create Folder**, and create the **program**, **input**, **output**, and **log** folders.
  - c. Upload data to the **userdata** file system.
    - i. Go to the **program** folder and click **Upload File**.
    - ii. Click **add file** and select a user program.
    - iii. Click **Upload**.
    - iv. Upload the user data file to the **input** directory using the same method.
2. Import OBS data to HDFS.

You can import OBS data to HDFS only when **Kerberos Authentication** is disabled and the cluster is running.

  - a. Log in to the MRS console.
  - b. Click the name of the cluster.
  - c. On the page displayed, select the **Files** tab page and click **HDFS File List**.
  - d. Select a data directory, for example, **bd\_app1**.

The **bd\_app1** directory is only an example. You can use any directory on the page or create a new one.
  - e. Click **Import Data** and click **Browse** to select an OBS path and an HDFS path.

- f. Click **OK**.

You can view the file upload progress on the **File Operation Records** tab page.

## 13.2 Can MRS Write Data to HBase Through an HBase External Table of Hive?

**Question:** Can MRS write data to HBase through an HBase external table of Hive?

No. Hive on HBase supports only data query.

## 13.3 Where Can I Download the Dependency Package (com.huawei.gaussc10) in the Hive Sample Project?

**Question:** Where can I download the dependency package (com.huawei.gaussc10) in the Hive sample project?

**Answer:** MRS does not have the **com.huawei.gaussc10** dependency package, which is a GaussDB dependency package and does not need to be configured. Exclude this package when building a Maven project.

## 13.4 Does MRS Support Python Code?

- The sample project and application development in MRS are the same. You can select either of them.
- MRS supports Python code. For details about component sample projects programmed using Python, see the *Developer Guide*.

## 13.5 Does OpenTSDB Support Python APIs?

**Question:** Does OpenTSDB support Python APIs?

OpenTSDB supports Python APIs. OpenTSDB provides HTTP-based RESTful APIs that are language-independent. Any language that supports HTTP requests can interconnect to OpenTSDB.

## 13.6 How Do I Obtain a Spark JAR File?

Huawei provides Huawei Mirrors for you to download all dependency JAR files of sample projects, but you need to download the open source dependency JAR files from the Maven central repository or other custom repositories.

 NOTE

Ensure the following conditions are met before you use a development tool in a local environment to download dependency JAR files:

- The local network is normal.  
Open a browser and enter the Huawei Mirrors URL in the address box to check whether your access is normal. If the access is abnormal, check whether your local network is accessible.
- Check whether the proxy is enabled for the development tool. If it is enabled, disable it by performing the following operations.  
Take IntelliJ IDEA 2020.2 as an example. Choose **File > Settings > Appearance & Behavior > System Settings > HTTP Proxy**, select **No proxy**, and click **OK** to save the configuration.

For details about how to configure an open source mirror, see [Configuring Huawei Open Source Mirrors](#).

## 13.7 How Do I Configure the node\_id Parameter When Calling the API for Adjusting Cluster Nodes?

### Question

How Do I Configure the Request Parameter node\_id When Calling the API for Adjusting Cluster Nodes (/v1.1/{project\_id}/cluster\_infos/{cluster\_id})?

### Answer

When the API is called, the value of the request parameter **node\_id** is fixed to **node\_orderadd**.

## 13.8 How Do I Manage and Use Third-Party JAR Packages for MRS Cluster Components?

Some components within an MRS cluster support the use of custom third-party JAR packages to meet specific customization needs. When a user application invokes a third-party JAR package while using a component, it can be stored in a designated directory on each node. Upon restarting the component, the JAR package in that directory will be automatically loaded by the corresponding component program.

The directory is as follows: *<Cluster installation directory>/third\_lib/<Component name>* of the component installation node.

Here are two examples of saving JAR files to their respective directories:

- Save the JAR package file used by the HBase service to the **\$ {BIGDATA\_HOME}/third\_lib/HBase** directory of the HBase node.
- Save the JAR package file used by the Hive service to the **\$ {BIGDATA\_HOME}/third\_lib/Hive** directory of each HiveServer instance node.

 **NOTE**

- Currently, custom JAR files are supported on HDFS, YARN, MapReduce, Spark, HBase, and Hive.
- Hive supports only Hive custom JAR files developed by users for UDF implementation. Do not store other open-source JAR files in the directory, including **org.apache.\***, **com.google.\***, and **javax.\***.
- This directory is supported only by MRS 3.x clusters.

# 14 Peripheral Service Interconnection

---

## 14.1 Does MRS Support Read and Write Operations on DLI Service Tables?

**Question:** Does MRS support read and write operations on DLI service tables?

**Answer:** If you have stored data on OBS, you can use Spark in MRS to read Data Lake Insight (DLI) tables, flexibly process table data, and save the result to another DLI table. If you have not stored data on OBS, you cannot use MRS to read or write DLI tables.

## 14.2 Does OBS Support the ListObjectsV2 Protocol?

**Question:** Does OBS support the ListObjectsV2 protocol?

**Answer:** No it doesn't.

## 14.3 Can a Crawler Service Be Deployed on Nodes in an MRS Cluster?

**Question:** Can a crawler service be deployed on nodes in an MRS cluster?

**Answer:** To ensure stable running, the crawler service or other third-party services cannot be deployed on nodes in the MRS cluster.

## 14.4 Does MRS Support Secure Deletion?

**Question:** Does MRS support secure deletion (preventing invalid restoration after deletion)?

**Answer:** Currently, MRS clusters can only restore data backed up by users. That is, MRS supports only secure deletion.

## 14.5 How Do I Use PySpark to Connect MRS Spark?

**Question:** How do I use pyspark on an ECS to connect an MRS Spark cluster with Kerberos authentication enabled on the Intranet?

**Answer:** Change the value of `spark.yarn.security.credentials.hbase.enabled` in the `spark-defaults.conf` file of Spark to `true` and use `spark-submit --master yarn --keytab keytabfile --principal principal` command to specify the Kerberos authentication file.

## 14.6 Why Mapped Fields Do not Exist in the Database After HBase Synchronizes Data to CSS?

**Question:** Why mapped fields do not exist in the database after HBase synchronizes data to CSS?

**Answer:** After data is synchronized from HBase of MRS to CSS, the database will not have mapped fields. Only a single table has mapped fields.

## 14.7 Can MRS Connect to an External KDC?

**Question:** Can MRS Connect to an External Key Distribution Center (KDC)?

**Answer:** MRS supports only the built-in KDC.

## 14.8 What Can I Do If Jetty Is Incompatible with MRS when Open-Source Kylin 3.x Is Interconnected with MRS 1.9.3?

For security purposes, MRS has upgraded some open-source third-party components that have serious security vulnerabilities. The upgrade causes the Jetty version compatibility issue when the open-source Kylin interconnects with MRS 1.9.3.

Perform the following operations to solve the issue:

1. Install the MRS client on an ECS node. For details, see [Using an MRS Client on Nodes Outside a Cluster](#). This operation uses the MRS client installation path `/srv/client/` as an example.
2. After the installation is complete, run the following commands to import the MRS client environment variable `bigdata_env` and the environment variables `HIVE_CONF` and `HCAT_HOME` required by Kylin.

```
source /srv/client/bigdata_env
export HIVE_CONF=/srv/client/Hive/config/
export HCAT_HOME=/srv/client/Hive/HCatalog
```



3. Install Kylin on the node where the MRS client is installed and specify **KYLIN\_HOME**. For details, see the [Kylin official website](#). For MRS 1.9.3, select Kylin for HBase 1.x for interconnection.  
**export KYLIN\_HOME=/srv/client/apache-kylin-3.0.2-bin-hbase1x**
4. Remove Jetty .jar packages from the **/srv/client/Hive/Beeline/lib/** directory in the Hive client directory to prevent version conflicts.  
  
Jetty .jar packages:  
javax-websocket-server-impl-9.4.26.v20200117.jar  
websocket-server-9.4.26.v20200117.jar  
jetty-all-9.4.26.v20200117-uber.jar  
jetty-runner-9.4.26.v20200117.jar  
apache-jsp-9.4.26.v20200117.jar
5. Start the Kylin service and check Kylin logs. In normal cases, the logs do not contain compatibility errors, such as **java.lang.NoSuchMethodException** and **java.lang.ClassNotFoundException**.  
**\$KYLIN\_HOME/bin/kylin.sh start**
6. Access the native Kylin page at **http://<hostname>:7070/kylin** and run the sample Cube script **\${KYLIN\_HOME}/bin/sample.sh** to check whether Kylin is running properly.

## 14.9 What Should I Do If Data Failed to Be Exported from MRS to an OBS Encrypted Bucket?

**Question:** What should i do if data failed to be exported from MRS to an obs encrypted bucket?

**Answer:** Only MRS 1.9.x and later versions support OBS bucket encryption. If you want to use encrypted OBS buckets, use MRS cluster version 1.9.x or later.

## 14.10 How Do I Interconnect MRS with LTS?

### Prerequisites

You have obtained the account AK and SK. For details, see [How Do I Obtain Access Keys \(AK/SK\)?](#)

### Procedure

1. Install ICAgent on the MRS node. For details, see [Installing ICAgent \(Cloud Hosts\)](#).

#### NOTE

For the first installation, install one server first, and then batch install all other hosts in the same way.

2. Create a host group and add the ICAgent host installed in [1](#) to the host group.
  - a. Log in to the Log Tank Service (LTS) console, and choose **Host Management** in the left navigation pane. On the displayed page, click **Create Host Group** in the upper right corner.

- b. In the **Create Host Group** slide-out panel that is displayed, enter a host group name in the **Host Group** field and set **Host OS** to **Linux** or **Windows**.
    - c. In the host list, select one or more hosts to add to the group and click **OK**.
      - You can filter hosts by host name or host IP address. You can also click **Search by Host IP Address** and enter multiple host IP addresses in the displayed search box to search for matches.
      - If your desired hosts are not in the list, click **Install ICAgent**. In the **Install ICAgent** slide-out panel displayed, install ICAgent on the hosts as prompted.
  3. Create a log group.
    - a. Log in to the LTS console. On the **Log Management** page, click **Create Log Group**.
    - b. In the **Create Log Group** slide-out panel displayed, enter a log group name. A log group name:
      - Can contain only letters, numbers, hyphens (-), underscores (\_), and periods (.).
      - Cannot start with a period (.) or underscore (\_), and cannot end with a period (.).
      - Can contain 1 to 64 characters.
    - c. Select an enterprise project. You can click **View Enterprise Projects** to view all enterprise projects.
    - d. Set **Log Retention Duration**. The system stores logs for 30 days by default. You can change the log retention duration after the log group is created.

LTS offers a free quota of 500 MB per month for log read/write, retention, and indexing. When the free quota is used up, you will be billed for excess usage on a pay-per-use basis.
    - e. Click **OK**.
      - On the **Log Management** page, you can view log group details, including name, retention duration, creation time, creation type, and tags. You can also modify retention duration and tags. Click **Modify** in the **Operation** column of the log group to modify the name and retention duration.
      - Click the log group name to access the log stream details page.
  4. Create a log stream.
    - a. On the LTS console, click the expansion button of a log group.
    - b. Click **Create Log Stream**. In the **Create Log Stream** slide-out panel displayed, enter a log stream name. A log stream name:
      - Can contain only letters, numbers, hyphens (-), underscores (\_), and periods (.).

- Cannot start with a period (.) or underscore (\_), and cannot end with a period (.).
  - Can contain 1 to 64 characters.
- c. Select an enterprise project. You can click **View Enterprise Projects** to view all enterprise projects.
  - d. Click **OK**.

On the log stream page, you can view details of the log stream, including log stream name, enterprise project, log retention duration, creation time, and creation type.

 **NOTE**

You can configure different log streams for different components.

5. Add hosts for log ingestion.
  - a. In the navigation pane of the LTS management console, choose **Log Ingestion > Ingestion Center**.
  - b. On the **All** tab page displayed, click **Elastic Cloud Server (ECS)**.
  - c. In the **Select Log Stream** step, set **Log Group** and **Log Stream** to the created log group and log stream name, and click **Next**.
  - d. In the **(Optional) Select Host Group** step, select the host group created in 2 and click **Next**.
  - e. In the **Collections** step, set **Collection Configuration Name** and **Collection Paths**, and click **Next**.

 **NOTE**

Path configurations

- You can configure multiple collection paths by clicking **Add Collection Path**.
- Each collection path must be unique. That is, the same path of the same host cannot be configured for different log groups and log streams.
- Collection paths support recursion. You can use double asterisks (\*\*) to collect logs from up to five directory levels.
- Collection paths support fuzzy match. You can use an asterisk (\*) to represent one or more characters of a directory or file name.
- If the collection path is set to a directory (such as **/var/logs/**), only **.log**, **.trace**, and **.out** files in the directory are collected.

For example, configure the following collection paths:

YARN task log path:

```
/srv/BigData/*/nm/containerlogs/**/container-localizer-syslog
/srv/BigData/*/nm/containerlogs/**/directory.info
/srv/BigData/*/nm/containerlogs/**/launch_container.sh
/srv/BigData/*/nm/containerlogs/**/prelaunch.err
/srv/BigData/*/nm/containerlogs/**/prelaunch.out
/srv/BigData/*/nm/containerlogs/**/stderr
/srv/BigData/*/nm/containerlogs/**/stdout
/srv/BigData/*/nm/containerlogs/**/syslog*
/srv/BigData/*/nm/containerlogs/**/*.*.log
```

All HDFS service logs:

/var/log/Bigdata/audit/hdfs/jn/hdfs-audit-journalnode.log

/var/log/Bigdata/audit/hdfs/jn/SecurityAuth.audit

/var/log/Bigdata/hdfs/\*/

Configure log collection paths for other components by referring to HDFS service logs.

- f. Skip the **Index Settings** step and submit.

#### NOTE

For details about how to configure index settings, see the *Log Tank Service User Guide*.

6. View the logs.
  - a. Log in to the LTS console and choose **Log Management** in the navigation pane.
  - b. In the **Log Group Name** column, click the name of the created log group to view logs.

## 14.11 How Do I Install HSS on MRS Cluster Nodes?

### Scenario

Nodes in an MRS cluster can be managed by the Host Security Service (HSS) for risk prevention and intrusion detection. You can identify and manage information assets on MRS nodes and build a security system that defends against server attacks.

### Constraints and Restrictions


HSS can be installed for MRS nodes run by OSs supported by the HSS agent.

### Procedure

- Step 1** Check whether the OSs of the MRS nodes are supported by HSS agents.

1. Log in to the MRS console.
2. In the navigation pane on the left, choose **Active Clusters**. Click a cluster name to go to the cluster details page.
3. Click the **Nodes** tab. Click the target node group and click any node in the list to access the basic information page of the server.

**Figure 14-1** Nodes tab



| Node Group | Node Type | Node Count | Billing Mode | Operation                     |
|------------|-----------|------------|--------------|-------------------------------|
| master_0   | Master    | AZ: 3      | Pay-per-use  | Upgrade Specifications   More |


  

| Node            | IP          | Operating Status | Health Status | CPU Usage | Memory Usage | Disk Usage | Network Speed                    | Specifications                                        | AZ  |
|-----------------|-------------|------------------|---------------|-----------|--------------|------------|----------------------------------|-------------------------------------------------------|-----|
| node-master1091 | 192.168.1.1 | Running          | Good          | 9%        | 18.91%       | 3.25%      | R: 28.1 Bytes/s; W: 14.7 Bytes/s | m3.2xlarge.8<br>8 vCPUs   64 GB   600 GB High I/O * 1 | AZ2 |
| node-master2002 | 192.168.1.2 | Running          | Good          | 7%        | 8.41%        | 3.23%      | R: 10.8 Bytes/s; W: 6.9 Bytes/s  | m3.2xlarge.8<br>8 vCPUs   64 GB   600 GB High I/O * 1 | AZ2 |

4. On the displayed ECS page, click the **Summary** tab. Check the EulerOS version in the **ECS Information** area.

If **EulerOS\_2.10\_x86\_64** is displayed in the image, the OS version is EulerOS 2.10 64-bit for Linux (x86).

**Figure 14-2** Viewing the OS version



5. Check whether the OS version of the MRS node is supported by HSS agents.  
For details about the OSs supported by HSS agents, see [Supported OSs](#).
  - If the OS version is supported, go to [Step 2](#) to install HSS agents.
  - If it is not supported, HSS cannot be installed.

**Step 2** Purchase HSS quotas and log in to each node in the MRS cluster to install HSS agents. After the agents are installed, you can enable HSS.

**NOTE**

Ensure you have purchased HSS in your MRS cluster node region and have used the installation package or installation command in the region to install HSS agents on your nodes.

For details about how to purchase HSS quotas, install agents, and enable HSS, see [HSS Getting Started](#).

----End

## 14.12 How Do I Connect to HBase of MRS Through HappyBase?

ThriftServer1 and ThriftServer2 cannot coexist. The HBase service of the MRS cluster uses ThriftServer2. However, HappyBase can connect to HBase only through the ThriftServer1 interface. Therefore, use Python to directly connect to HBase. For details, see <https://github.com/huaweicloud/huaweicloud-mrs-example/blob/mrs-1.8/src/hbase-examples/hbase-python-example/DemoClient.py>.

## 14.13 Can the Hive Driver Be Interconnected with DBCP2?

**Question:** Can the Hive driver be interconnected with DBCP2?

**Answer:** The Hive driver cannot be interconnected with the DBCP2 database connection pool. The DBCP2 database connection pool invokes the **isValid** method to check whether a connection is available. However, Hive directly throws an exception when implementing this method.

# 15 Upgrade and Patching

---

## 15.1 How Do I Upgrade an MRS Cluster?

- Non-LTS clusters  
An existing non-LTS cluster cannot be upgraded to a later version. You can create a cluster of the new version and migrate data from the old cluster to the new cluster.
- MRS LTS clusters  
MRS LTS clusters can be upgraded to a target LTS version. Contact our O&M engineers to confirm the upgrade path and buy professional services for cluster upgrade.

## 15.2 Does MRS Support Kernel Version Upgrade of Components in a Cluster?

**Question:** Does MRS support kernel version upgrade of components in a cluster?

**Answer:** MRS does not support independent upgrade of kernel versions of components in a cluster. To use components of new versions, delete the current cluster and create a cluster again. For details about the components supported by MRS, see [List of MRS Component Versions](#).