MapReduce Service

User Guide

Issue 01

Date 2024-05-07

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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

1 Preparing a User

1.1 Configuring Cloud Service Permissions

The MapReduce Service (MRS) console can interact with MRS clusters to provide related functions and monitor cluster status. Permission assignment is required when you use MRS for the first time.

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

Procedure

Step 1 Log in to the management console.

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

Figure 1-1 Assigning permissions

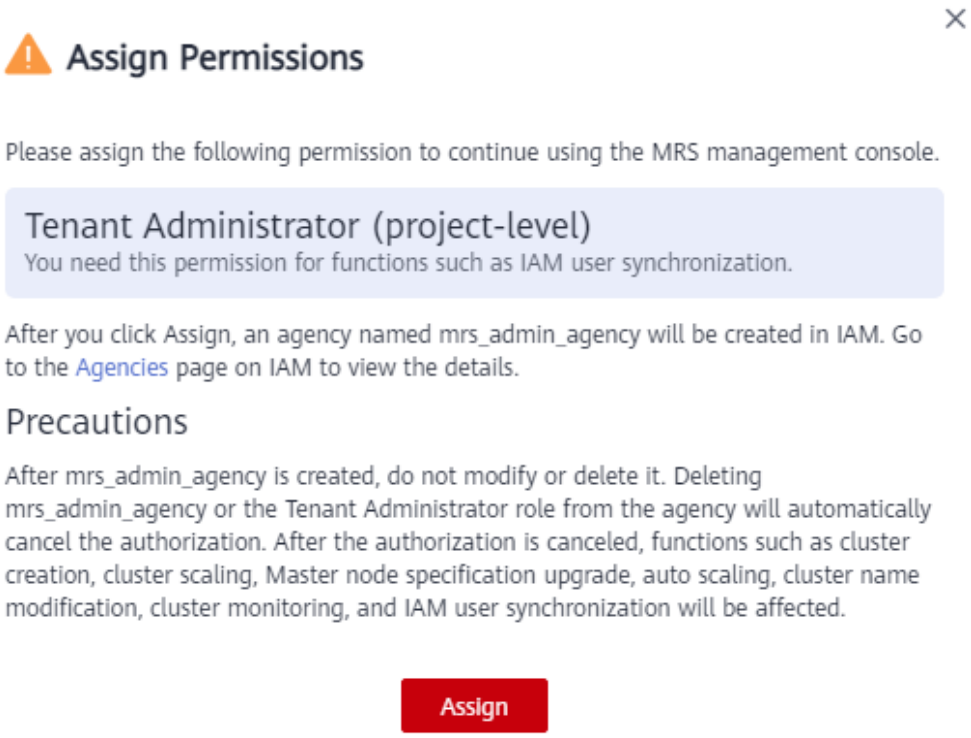

Step 3 Click **Assign**.

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

NOTE

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

----End

1.2 Creating an MRS User

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

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

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

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

Prerequisites

Learn about the permissions supported by MRS by referring to **Permission Management**. For the permissions of other services, see **Permission Description**.

Process Flow

Figure 1-2 Process for granting MRS permissions

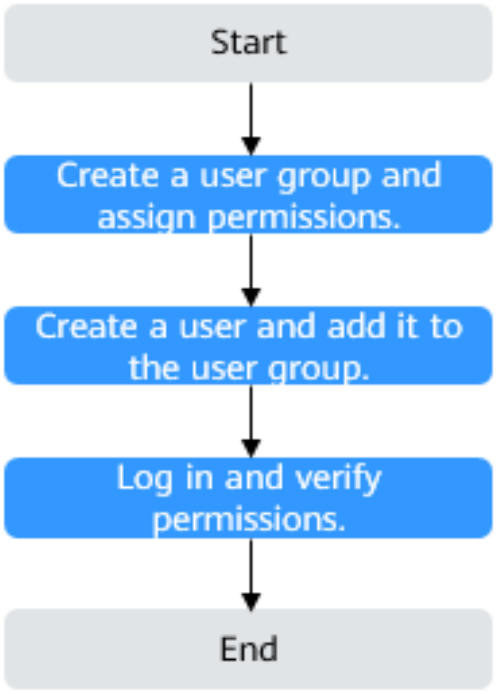

1. **Creating a User Group and Assigning Permissions**

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

2. .

   Create a user on the IAM console and add the user to the group created in **1. Create a user group and assign permissions to it**.

3. **Log in** and verify permissions.

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

   - Choose **Service List** > **Analytics** > **MapReduce Service**. Click **Buy Cluster** on the MRS console. If you fail to buy an MRS cluster (assume that you only have the **MRS ReadOnlyAccess** permission), the **MRS ReadOnlyAccess** policy has taken effect.
   - Choose any other service in **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **MRS ReadOnlyAccess** policy has already taken effect.

MRS Permission Description

By default, new IAM users do not have any permissions. To assign permissions to a user, add the user to one or more groups and assign permissions policies or roles to these groups. The user then inherits permissions from the groups it is a member of and can perform specified operations on cloud services based on the permissions.

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

You can grant permissions by using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to grant permissions, you need to also assign other roles on which the permissions depend to take effect. However, roles are not an ideal choice for fine-grained authorization and secure access control.
- Policies: A type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, you can grant MRS users only the permissions for performing specified operations on MRS clusters, such as creating a cluster and querying a cluster list rather than deleting a cluster. Most policies define permissions based on APIs. For the API actions supported by MRS, see **Permissions Policies and Supported Actions**.

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

Table 1-1 MRS system policies

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

Policy	Description	Type
MRS Administrator	Permissions: • All operations on MRS • Users with permissions of this policy must also be granted permissions of the **Tenant Guest** and **Server Administrator** policies.	RBAC policy

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

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

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

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

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

1.3 Creating a Custom Policy

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

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

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

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

For details, see **Creating a Custom Policy**.

NOTE

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

The following section contains examples of common MRS custom policies.

Example Custom Policies

- Example 1: Allowing users to create MRS clusters only

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

- Example 2: Allowing users to resize an MRS cluster

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

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

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

- Example 4: Allowing users to create an ECS cluster with the minimum permission

NOTE

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

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



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

- Example 5: Allowing users to create a BMS cluster with the minimum permission

  NOTE

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

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



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

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

  NOTE

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

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



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

1.4 Synchronizing IAM Users to MRS

IAM user synchronization is to synchronize IAM users bound with MRS policies to the MRS system and create accounts with the same usernames but different passwords as the IAM users. Then, you can use an IAM username (the password needs to be reset by user **admin** of Manager) to log in to Manager for cluster management, and submit jobs on the GUI in a cluster with Kerberos authentication enabled.

Table 1-3 compares IAM users' permission policies and the synchronized users' permissions on MRS. For details about the default permissions on Manager, see **Default Permission Information**.

Table 1-3 Policy and permission mapping after synchronization

Policy Type	IAM Policy	User's Default Permissions on MRS After Synchronization	Have Permission to Perform the Synchronization	Have Permission to Submit Jobs
Fine-grained	MRS ReadOnlyAccess	Manager_viewer	No	No
	MRS CommonOperations	• Manager_viewer • default • launcher-job	No	Yes
	MRS FullAccess	• Manager_administrator • Manager_auditor • Manager_operator • Manager_tenant • Manager_viewer • System_administrator • default • launcher-job	Yes	Yes

Policy Type	IAM Policy	User's Default Permissions on MRS After Synchronization	Have Permission to Perform the Synchronization	Have Permission to Submit Jobs
RBAC	MRS Administrator	• Manager_administrator • Manager_auditor • Manager_operator • Manager_tenant • Manager_viewer • System_administrator • default • launcher-job	No	Yes
	Server Administrator, Tenant Guest, and MRS Administrator	• Manager_administrator • Manager_auditor • Manager_operator • Manager_tenant • Manager_viewer • System_administrator • default • launcher-job	Yes	Yes

Policy Type	IAM Policy	User's Default Permissions on MRS After Synchronization	Have Permission to Perform the Synchronization	Have Permission to Submit Jobs
	Tenant Administrator	• Manager_administrator • Manager_auditor • Manager_operator • Manager_tenant • Manager_viewer • System_administrator • default • launcher-job	Yes	Yes
Custom	Custom policy	• Manager_viewer • default • launcher-job	• If custom policies use RBAC policies as a template, refer to the RBAC policies. • If custom policies use fine-grained policies as a template, refer to the fine-grained policies. The fine-grained policies are recommended.	Yes

NOTE

To facilitate user permission management, use fine-grained policies rather than RBAC policies. In fine-grained policies, the Deny action takes precedence over other actions.

- A user has permission to synchronize IAM users only when the user has the Tenant Administrator role or has the Server Administrator, Tenant Guest, and MRS Administrator roles at the same time.
- A user with the **action:mrs:cluster:syncUser** policy has permission to synchronize IAM users.

Procedure

Step 1 Create a user and authorize the user to use MRS. For details, see **Creating an MRS User**.

Step 2 Log in to the MRS management console and create a cluster. For details, see **Buying a Custom Cluster**.

Step 3 In the left navigation pane, choose **Clusters** > **Active Clusters**. Click the cluster name to go to the cluster details page.

Step 4 On the **Dashboard** tab page, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.

Step 5 In the **IAM User Sync** dialog box, search for the user group to which the IAM user to be synchronized belongs and click the user group name. In the **User** column, select the desired IAM user and click **Synchronize**.

NOTE

- You can select all users to synchronize them at a time.
- If you select user groups only, users will not be synchronized. You must select specific user names in the user group.
- All user groups are displayed. Those cannot be selected cannot be synchronized.

Step 6 After a synchronization request is sent, choose **Operation Logs** in the navigation tree on the left of the MRS console to check whether the synchronization is successful. For details about the logs, see **Viewing MRS Operation Logs**.

Step 7 After the synchronization is successful, use the user synchronized with IAM to perform subsequent operations.

NOTE

- When the policy of the user group to which the IAM user belongs changes from **MRS ReadOnlyAccess** to **MRS CommonOperations**, **MRS FullAccess**, or **MRS Administrator**, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from **MRS CommonOperations**, **MRS FullAccess**, or **MRS Administrator** to **MRS ReadOnlyAccess**, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.
- After you click **Synchronize** on the right side of **IAM User Sync**, the cluster details page is blank for a short time, because user data is being synchronized. The page will be properly displayed after the data synchronization is complete.

- Submitting jobs in a security cluster: Users can submit jobs using the job management function on the GUI in the security cluster. For details, see **Running a MapReduce Job**.
- All tabs are displayed on the cluster details page, including **Components**, **Tenants**, and **Backups & Restorations**.
- Logging in to Manager
  a. Log in to Manager as user **admin**. For details, see **Accessing Manager**.
  b. Initialize the password of the user synchronized with IAM. For details, see **Initializing the Password of a System User**.
  c. Modify the role bound to the user group to which the user belongs to control user permissions on Manager. For details, see **Related Tasks**. For details about how to create and modify a role, see **Creating a Role**. After the component role bound to the user group to which the user belongs is modified, it takes some time for the role permissions to take effect.
  d. Log in to Manager using the user synchronized with IAM and the password after the initialization in **Step 7.b**.

  NOTE

  If the IAM user's permission changes, go to **Step 4** to perform second synchronization. After the second synchronization, a system user's permissions are the union of the permissions defined in the IAM system policy and the permissions of roles added by the system user on Manager. After the second synchronization, a custom user's permissions are subject to the permissions configured on Manager.
  - System user: If all user groups to which an IAM user belongs are bound to system policies (RABC policies and fine-grained policies belong to system policies), the IAM user is a system user.
  - Custom user: If the user group to which an IAM user belongs is bound to any custom policy, the IAM user is a custom user.

Step 8 Undo IAM user synchronization.

To undo the synchronization of an IAM user, select the user in the **User** column in the **Synchronized** tab and click **Undo Sync**.

To undo the synchronization of all users in an IAM user group, select the user group in the **User Group** column in the **Synchronized** tab and click **Undo Sync**.

----End

2 Configuring a Cluster

2.1 How to Buy an MRS Cluster

This section describes how to buy an MRS cluster.

- **Quickly Buying a Hadoop Analysis Cluster**: In the **Quick Config** tab, you can quickly configure parameters to buy a Hadoop analysis cluster within a few minutes, facilitating analysis and queries of vast amounts of data.
- **Quickly Buying an HBase Query Cluster**: In the **Quick Config** tab, you can quickly configure parameters to buy an HBase query cluster within a few minutes, facilitating storage and distributed computing of vast amounts of data.
- **Quickly Buying a Kafka Streaming Cluster**: In the **Quick Config** tab, you can quickly configure parameters to buy a Kafka streaming cluster within a few minutes, facilitating streaming data ingestion as well as real-time data processing and storage.
- **Quickly Buying a ClickHouse Cluster**: You can quickly buy a ClickHouse cluster. ClickHouse is a columnar database management system used for online analysis. It features optimal compression rate and fast query performance.
- **Quickly Buying a Real-time Analysis Cluster**: You can buy a real-time analysis cluster within a few minutes to quickly collect, analyze, and query a large amount of data.
- **Buying a Custom Cluster**: On the **Custom Config** tab page, you can flexibly configure parameters to buy clusters based on application scenarios, such as the billing mode andECS specifications to better suit your service requirements.

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

2.2 Quick Configuration

2.2.1 Quickly Buying a Hadoop Analysis Cluster

This section describes how to quickly buy a Hadoop analysis cluster for analysis and query of vast amounts of data. In the open source Hadoop ecosystem, Hadoop uses YARN to manage cluster resources, Hive and Spark to provide offline storage and computing of large-scale distributed data, Spark Streaming and Flink to offer streaming data computing, and Presto to enable interactive queries, Tez to provide a distributed computing framework of directed acyclic graphs (DAGs).

The Hadoop analysis cluster consists of the following components:

- MRS 1.9.2: Hadoop 2.8.3, Spark 2.2.2, Hive 2.3.3, Presto 0.216, Tez 0.9.1, Ranger 1.0.1, and Flink 1.7.0.
- MRS 3.1.0: Hadoop 3.1.1, Hive 3.1.0, Spark2x 2.4.5, Flink 1.12.0, ZooKeeper 3.5.6, Ranger 2.0.0, Tez 0.9.2, and Presto 333.
- MRS 3.1.2-LTS.3: Hadoop 3.1.1, Hive 3.1.0, Spark2x 3.1.1, Flink 1.12.2, ZooKeeper 3.6.3, Ranger 2.0.0, and Tez 0.9.2.
- MRS 3.1.5: Hadoop 3.1.1, Hive 3.1.0, Spark2x 3.1.1, Tez 0.9.2, Flink 1.12.2, ZooKeeper 3.6.3, Ranger 2.0.0, and Presto 333.

Quickly Buying a Hadoop Analysis Cluster

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

Step 2 On the displayed page, click the **Quick Config** tab.

Step 3 Configure basic cluster information. For details about the parameters, see **Buying a Custom Cluster**.

- **Region**: Use the default value.
- **Billing Mode**: **Pay-per-use**. If you select this mode, a prepaid balance will be frozen.
- **Cluster Name**: You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing, for example, **mrs_20180321**.
- **Cluster Type**: Use the default value.
- **Version Type**: **Normal** is selected by default. (Components vary depending on the version type. Select a version type as needed.)
- **Cluster Version**: Select the latest version, which is the default value. (The components provided by a cluster vary according to the cluster version. Select a cluster version based on site requirements.)
- **Component**: Select **Hadoop analysis cluster**.
- **AZ**: Use the default value.
- **Enterprise Project**: Retain the default value.
- **VPC**: Use the default value. If there is no available VPC, click **View VPC** to access the VPC console and create a new VPC.
- **Subnet**: Use the default value.
- **CPU Architecture**: Use the default value.
- **Cluster Node**: Select the number of cluster nodes and node specifications based on site requirements. For MRS 3.*x* or later, the memory of the master node must be greater than 64 GB.

- **Kerberos Authentication**: Whether to enable Kerberos authentication. This option cannot be changed after you buy a cluster.
- **Username**: The default value is **root/admin**. User **root** is used to remotely log in to ECSs, and user **admin** is used to access the cluster management page.
- **Password**: Set a password for user **root**/**admin**.
- **Confirm Password**: Enter the password of user **root**/**admin** again.

Figure 2-1 Hadoop Analysis Cluster

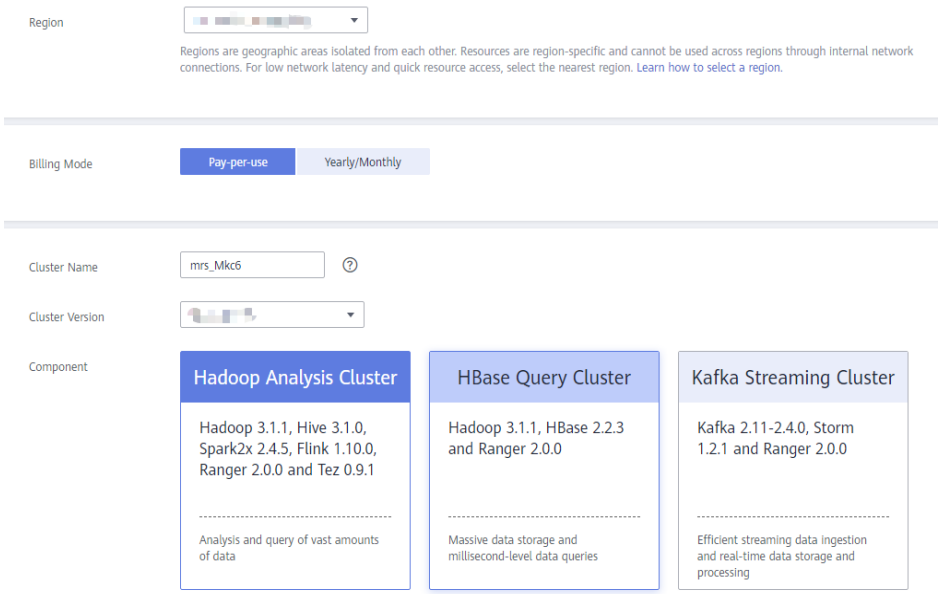

Figure 2-2 Cluster node configurations

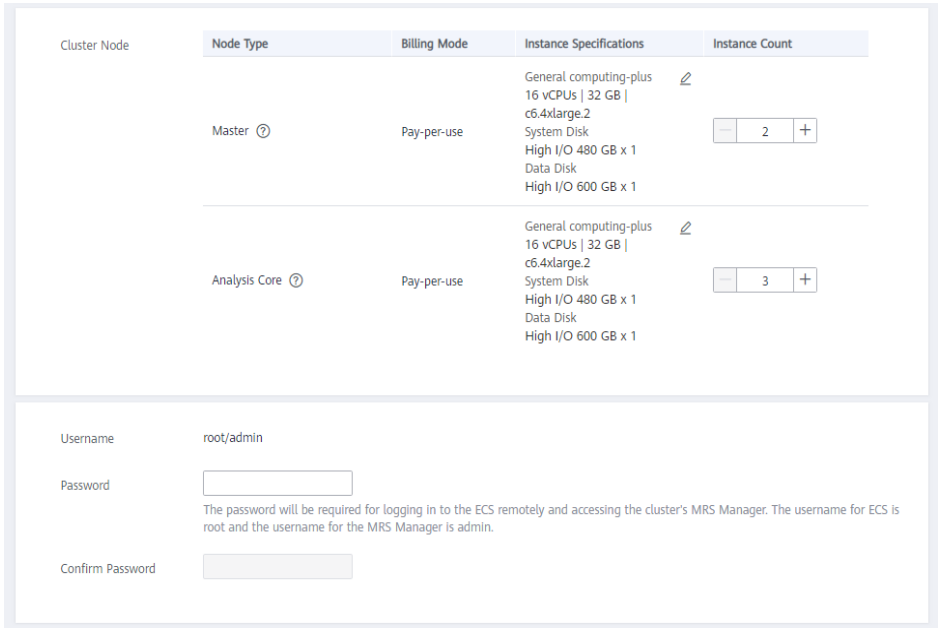

Step 4 Select the checkbox to enable secure communications. For details, see **Communication Security Authorization**.

Step 5 Click **Buy Now**.

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

NOTE

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

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

For details about cluster status during creation, see the description of the status parameters in **Table 3-4**.

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

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

----End

2.2.2 Quickly Buying an HBase Query Cluster

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

The HBase analysis cluster consists of the following components:

- MRS 1.9.2: Hadoop 2.8.3, HBase 1.3.1, and Ranger 1.0.1.
- MRS 3.1.0: Hadoop 3.1.1, HBase 2.2.3, ZooKeeper 3.5.6, and Ranger 2.0.0.
- MRS 3.1.2-LTS.3: Hadoop 3.1.1, HBase 2.2.3, ZooKeeper 3.6.3, and Ranger 2.0.0.
- MRS 3.1.5: Hadoop 3.1.1, HBase 2.2.3, ZooKeeper 3.6.3, and Ranger 2.0.0.
- MRS 3.2.0-LTS.1: Hadoop 3.1.1, HBase 2.2.3, ZooKeeper 3.6.3, and Ranger 2.0.0.

Quickly Buying an HBase Query Cluster

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

Step 2 On the displayed page, click the **Quick Config** tab.

Step 3 Configure basic cluster information. For details about the parameters, see **Buying a Custom Cluster**.

- **Region**: Use the default value.

- **Billing Mode**: **Pay-per-use**. If you select this mode, a prepaid balance will be frozen.
- **Cluster Name**: You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing, for example, **mrs_20180321**.
- **Cluster Type**: Use the default value.
- **Cluster Version**: Select the latest version, which is the default value. (The components provided by a cluster vary according to the cluster version. Select a cluster version based on site requirements.)
- **Component**: Select **HBase Query Cluster**.
- **AZ**: Use the default value.
- **Enterprise Project**: Retain the default value.
- **VPC**: Use the default value. If there is no available VPC, click **View VPC** to access the VPC console and create a new VPC.
- **Subnet**: Use the default value.
- **Enterprise Project**: Select the default project.
- **Cluster Node**: Select the number of cluster nodes and node specifications based on site requirements. For MRS 3.*x* or later, the memory of the master node must be greater than 64 GB.
- **Kerberos Authentication**: Whether to enable Kerberos authentication. This option cannot be changed after you buy a cluster.
- **Username**: The default value is **root**/**admin**. User **root** is used to remotely log in to ECSs and user **admin** is used to access FusionInsight Manager.
- **Password**: Set a password for user **root**/**admin**.
- **Confirm Password**: Enter the password of user **root**/**admin** again.

Figure 2-3 HBase Query Cluster

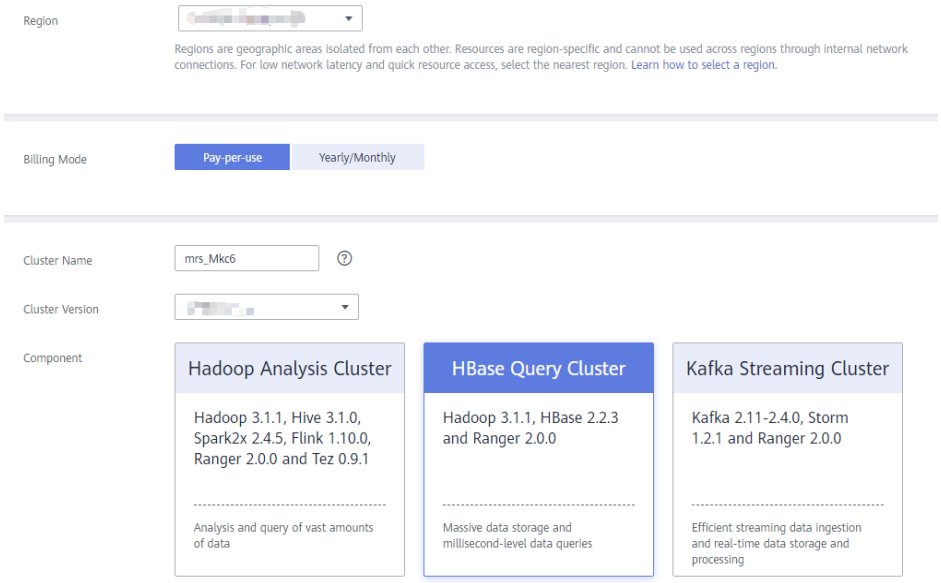

Figure 2-4 Cluster node configurations

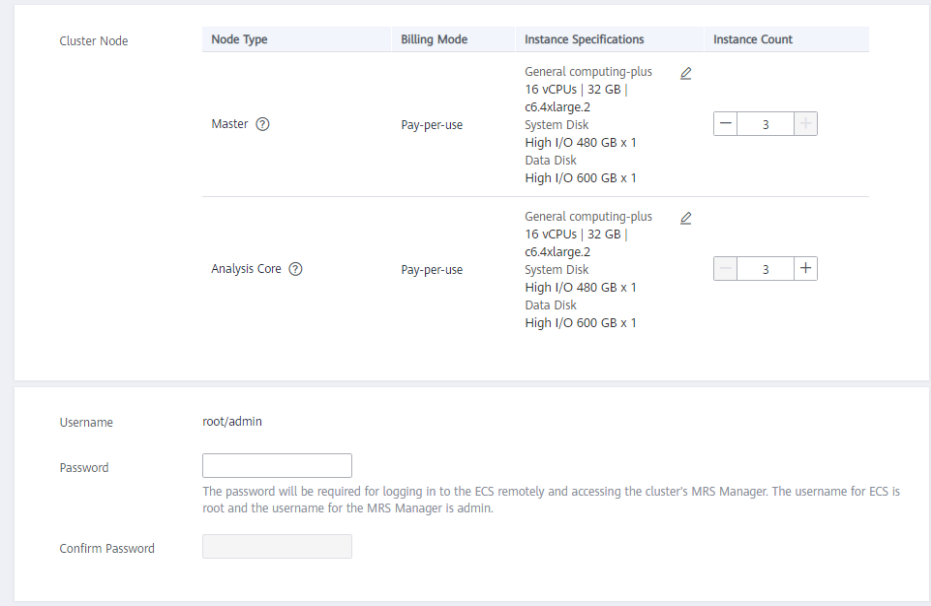

Step 4 Select the checkbox to enable secure communications. For details, see **Communication Security Authorization**.

Step 5 Click **Create Now**.

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

NOTE

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

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

For details about cluster status during creation, see the description of the status parameters in **Table 3-4**.

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

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

----End

2.2.3 Quickly Buying a Kafka Streaming Cluster

This section describes how to quickly buy a Kafka streaming cluster. The Kafka cluster uses the Kafka and Storm components to provide an open source messaging system with high throughput and scalability. It is widely used in scenarios such as log collection and monitoring data aggregation to implement efficient streaming data collection and real-time data processing and storage.

The Kafka streaming cluster consists of the following components:

- MRS 1.9.2: Kafka 1.1.0 and Storm 1.2.1

Quickly Buying a Kafka Streaming Cluster

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

Step 2 On the displayed page, click the **Quick Config** tab.

Step 3 Configure basic cluster information. For details about the parameters, see **Buying a Custom Cluster**.

- **Region**: Use the default value.
- **Billing Mode**: **Pay-per-use**. If you select this mode, a prepaid balance will be frozen.
- **Cluster Name**: You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing, for example, **mrs_20200321**.
- **Cluster Type**: Select **Streaming cluster**.
- Version Type: Select **Normal**.
- **Cluster Version**: Select a cluster version as you need.
- **Component**: Select **Kafka streaming cluster**.
- **AZ**: Use the default value.
- **Enterprise Project**: Retain the default value.
- **VPC**: Use the default value. If there is no available VPC, click **View VPC** to access the VPC console and create a new VPC.
- **Subnet**: Use the default value.
- **CPU Architecture**: Use the default value. This parameter is unavailable in MRS 3.*x*.
- **Cluster Node**: Select the number of cluster nodes and node specifications based on site requirements. For MRS 3.*x* or later, the memory of the master node must be greater than 64 GB.
- LVM: Use the default value. This parameter is not available in MRS 3.*x*.
- **Kerberos Authentication**: Whether to enable Kerberos authentication. This option cannot be changed after you buy a cluster.
- **Username**: The default value is **root/admin**. User **root** is used to remotely log in to ECSs, and user **admin** is used to access the cluster management page.
- **Password**: Set a password for user **root**/**admin**.
- **Confirm Password**: Enter the password of user **root**/**admin** again.

Figure 2-5 Kafka Streaming Cluster

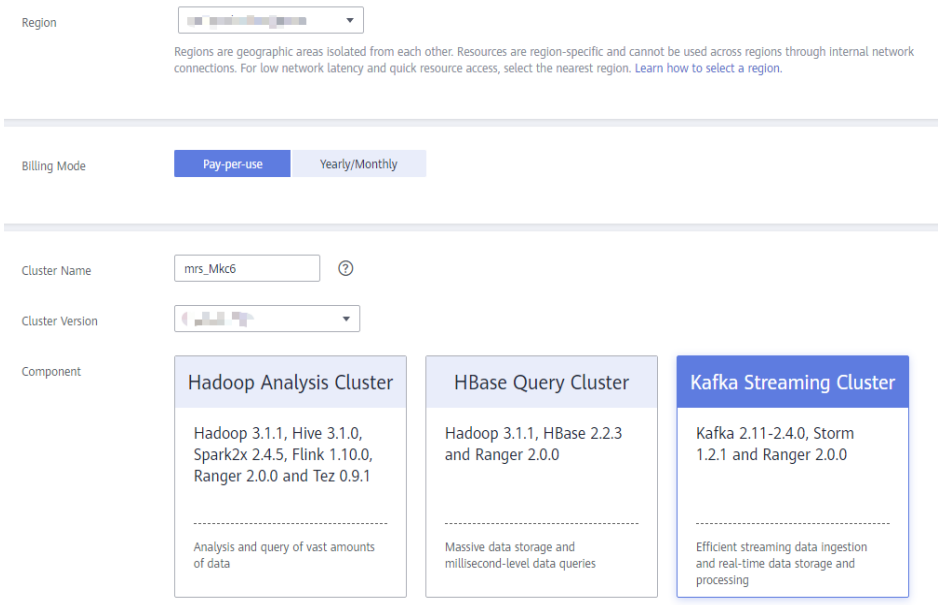

Figure 2-6 Cluster node configurations

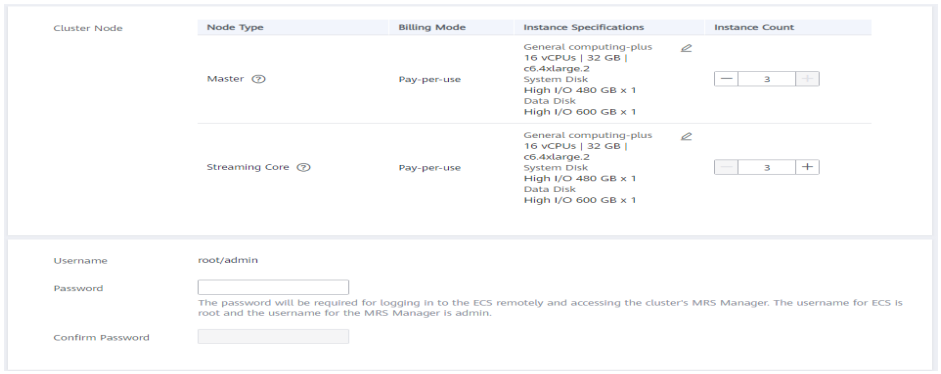

Step 4 Select the checkbox to enable secure communications. For details, see **Communication Security Authorization**.

Step 5 Click **Buy Now**.

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

NOTE

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

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

For details about cluster status during creation, see the description of the status parameters in **Table 3-4**.

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

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

----End

2.2.4 Quickly Buying a ClickHouse Cluster

This section describes how to quickly buy a ClickHouse cluster. ClickHouse is a columnar database management system used for online analysis. It features optimal compression rate and fast query performance. It is widely used in Internet advertisement, app and web traffic analysis, telecom, finance, and IoT fields.

The ClickHouse cluster consists of the following components:

- MRS 3.1.0: ClickHouse 21.3.4.25 and ZooKeeper 3.5.6.
- MRS 3.1.2-LTS.3: .
- MRS 3.1.5: ClickHouse 21.3.4.25 and ZooKeeper 3.6.3.
- MRS 3.2.0-LTS.1: ClickHouse 22.3.2.2 and ZooKeeper 3.6.3.

The ClickHouse cluster table engine that uses Kunpeng as the CPU architecture does not support HDFS and Kafka.

Quickly Buying a ClickHouse Cluster

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

Step 2 On the displayed page, click the **Quick Config** tab.

Step 3 Configure basic cluster information. For details about the parameters, see **Buying a Custom Cluster**.

- **Region**: Use the default value.
- **Billing Mode**: **Pay-per-use**. If you select this mode, a prepaid balance will be frozen.
- **Cluster Name**: You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing, Example: **mrs_20201121**.
- **Cluster Type**: Use the default value.
- **Cluster Version**: Select the latest version, which is the default value. (The components provided by a cluster vary according to the cluster version. Select a cluster version based on site requirements.)
- **Component**: Select **ClickHouse cluster**.
- **AZ**: Use the default value.
- **Enterprise Project**: Retain the default value.
- **VPC**: Use the default value. If there is no available VPC, click **View VPC** to access the VPC console and create a new VPC.
- **Subnet**: Use the default value.
- **CPU Architecture**: Retain the default value.

- **Cluster Node**: Select the number of cluster nodes and node specifications based on site requirements. For MRS 3.*x* or later, the memory of the master node must be greater than 64 GB.
- **Kerberos Authentication**: Whether to enable Kerberos authentication. This option cannot be changed after you buy a cluster.
- **Username**: The default value is **root**/**admin**. User **root** is used to remotely log in to ECSs and user **admin** is used to access FusionInsight Manager.
- **Password**: Set a password for user **root**/**admin**.
- **Confirm Password**: Enter the password of user **root**/**admin** again.

Figure 2-7 ClickHouse Cluster

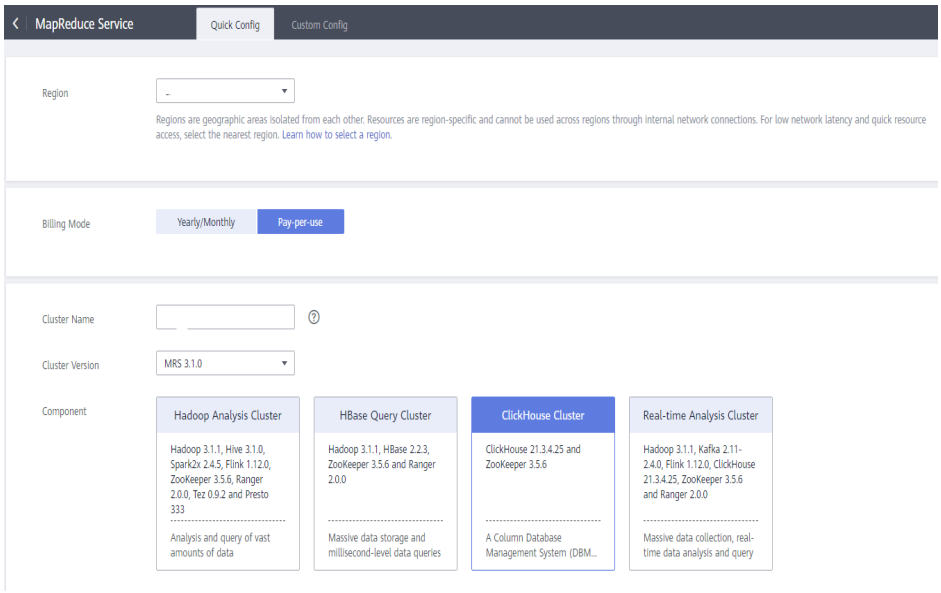

Figure 2-8 Cluster node configurations

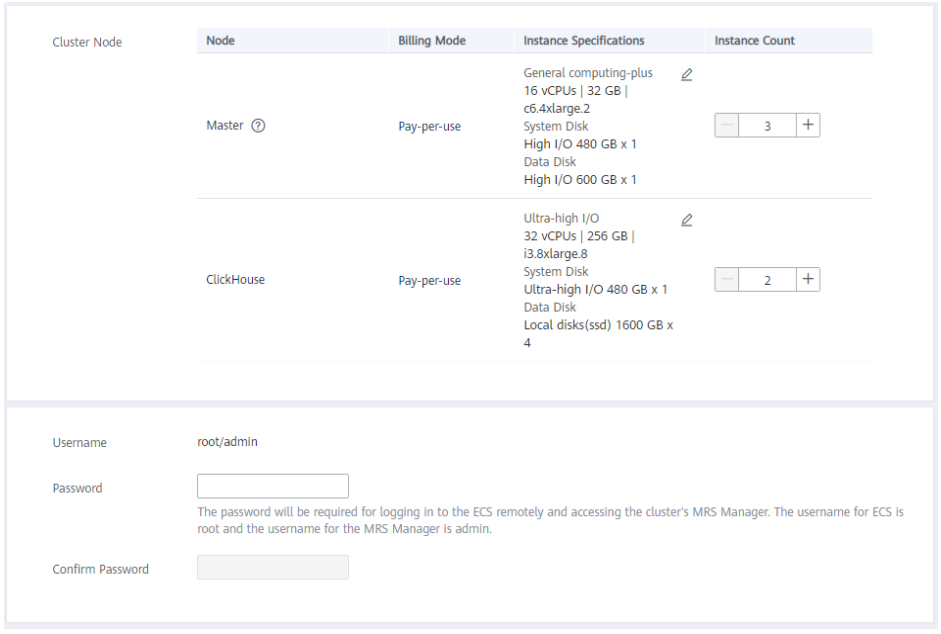

Step 4 Select the checkbox to enable secure communications. For details, see **Communication Security Authorization**.

Step 5 Click **Buy Now**.

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

NOTE

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

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

For details about cluster status during creation, see the description of the status parameters in **Table 3-4**.

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

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

----End

2.2.5 Quickly Buying a Real-time Analysis Cluster

This section describes how to quickly buy a real-time analysis cluster. The real-time analysis cluster uses Hadoop, Kafka, Flink, and ClickHouse to collect, analyze, and query a large amount of data in real time.

The real-time analysis cluster consists of the following components:

- MRS 3.1.0: Hadoop 3.1.1, Kafka 2.11-2.4.0, Flink 1.12.0, ClickHouse 21.3.4.25, ZooKeeper 3.5.6, and Ranger 2.0.0.
- MRS 3.1.2-LTS.3: Hadoop 3.1.1, Kafka 2.11-2.4.0, Flink 1.12.2, ClickHouse 21.3.4.25, ZooKeeper 3.6.3, and Ranger 2.0.0.
- MRS 3.1.5: Hadoop 3.1.1, Kafka 2.11-2.4.0, Flink 1.12.2, ClickHouse 21.3.4.25, ZooKeeper 3.6.3, and Ranger 2.0.0.
- MRS 3.2.0-LTS.1: Hadoop 3.1.1, Kafka 2.11-2.4.0, Flink 1.15.0, ClickHouse 22.3.2.2, ZooKeeper 3.6.3, and Ranger 2.0.0.

Quickly Buying a Real-time Analysis Cluster

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

Step 2 On the displayed page, click the **Quick Config** tab.

Step 3 Configure basic cluster information. For details about the parameters, see **Buying a Custom Cluster**.

- **Region**: Use the default value.
- **Billing Mode**: **Pay-per-use**. If you select this mode, a prepaid balance will be frozen.

- **Cluster Name**: You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing, Example: **mrs_20201130**.
- **Cluster Type**: Use the default value.
- **Cluster Version**: Select the latest version, which is the default value. (The components provided by a cluster vary according to the cluster version. Select a cluster version based on site requirements.)
- **Component**: Select **Real-time Analysis Cluster**.
- **AZ**: Use the default value.
- **VPC**: Use the default value. If there is no available VPC, click **View VPC** to access the VPC console and create a new VPC.
- **Subnet**: Use the default value.
- **Enterprise Project**: Use the default value.
- **CPU Architecture**: Retain the default value.
- **Cluster Node**: Select the number of cluster nodes and node specifications based on site requirements. For MRS 3.*x* or later, the memory of the master node must be greater than 64 GB.
- **Kerberos Authentication**: Whether to enable Kerberos authentication. This option cannot be changed after you buy a cluster.
- **Username**: The default value is **root**/**admin**. User **root** is used to remotely log in to ECSs and user **admin** is used to access FusionInsight Manager.
- **Password**: Set a password for user **root**/**admin**.
- **Confirm Password**: Enter the password of user **root**/**admin** again.

Figure 2-9 Real-time Analysis Cluster

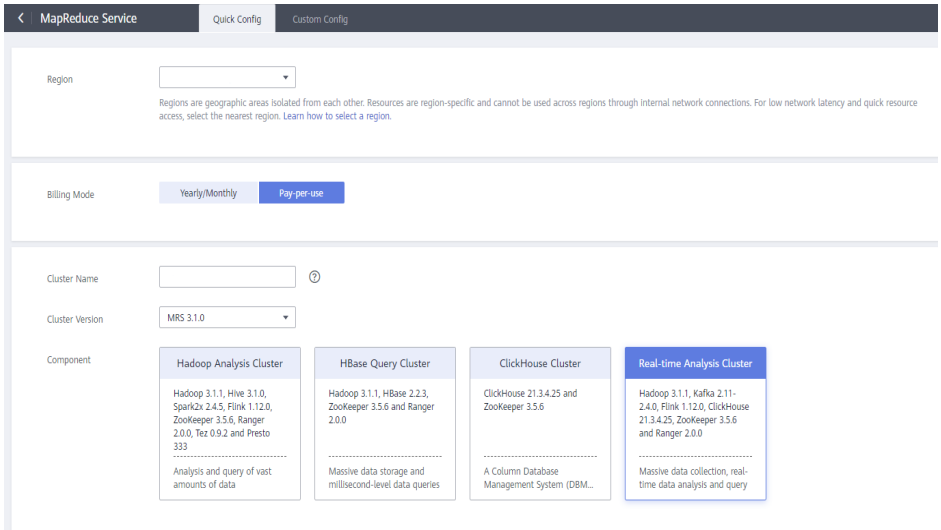

Figure 2-10 Cluster node configurations

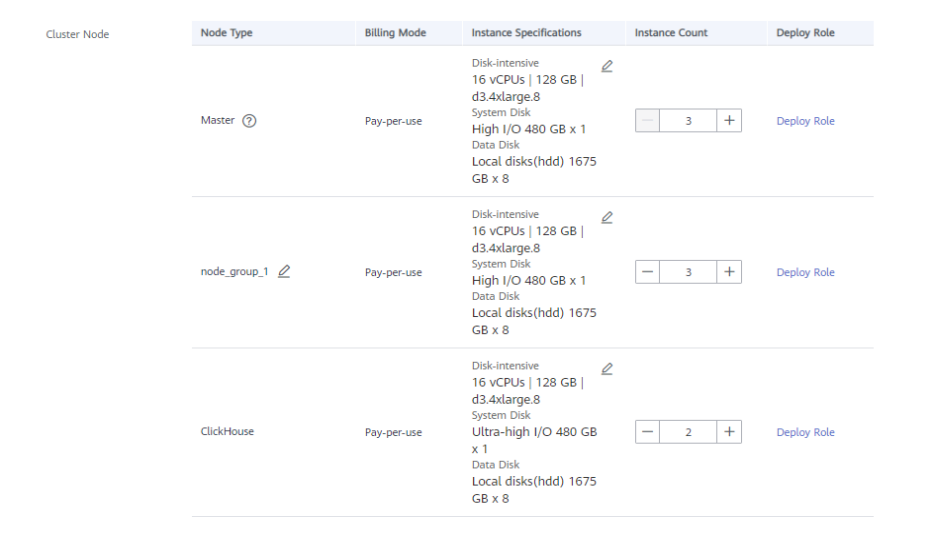

Step 4 Select the checkbox to enable secure communications. For details, see **Communication Security Authorization**.

Step 5 Click **Buy Now**.

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

NOTE

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

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

For details about cluster status during creation, see the description of the status parameters in **Table 3-4**.

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

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

----End

2.3 Buying a Custom Cluster

To use MRS, buy a cluster on the MRS management console.

You can create an IAM user or user group on the IAM management console and grant it specific operation permissions, to perform refined resource management after registering an account. For details, see **Creating an MRS User**.

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

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

NOTE

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

Step 3 Configure cluster information by referring to **Software Configurations** and click **Next**.

Figure 2-11 Configure Software

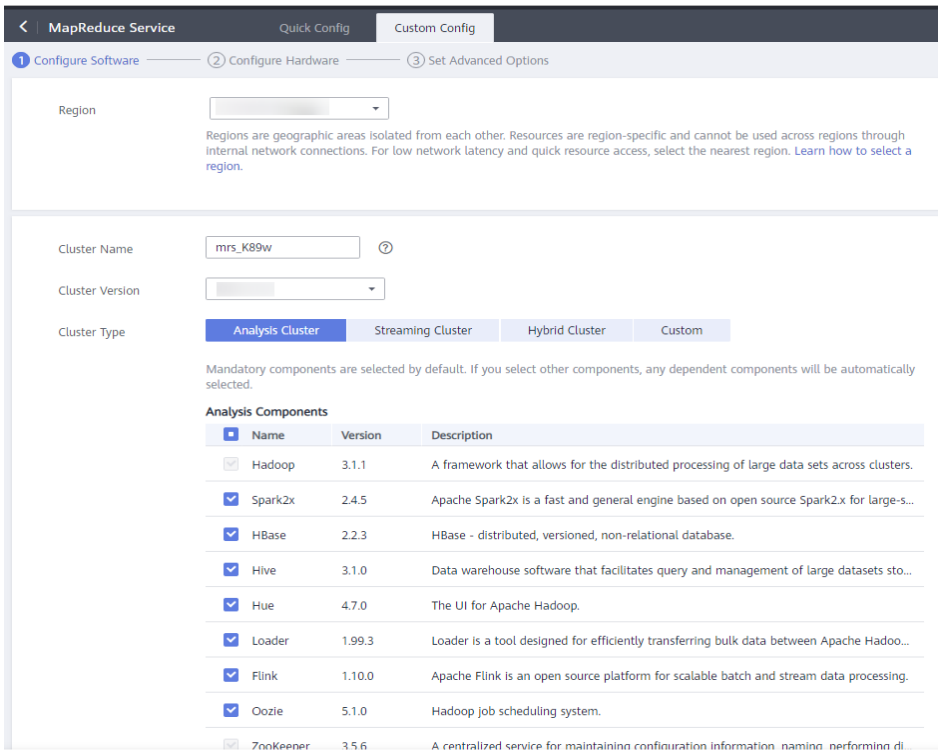

NOTE

Only one billing mode is supported in some regions. For details, see the management console.

Step 4 Configure cluster information by referring to **Hardware Configurations** and click **Next**.

Figure 2-12 Configure Hardware

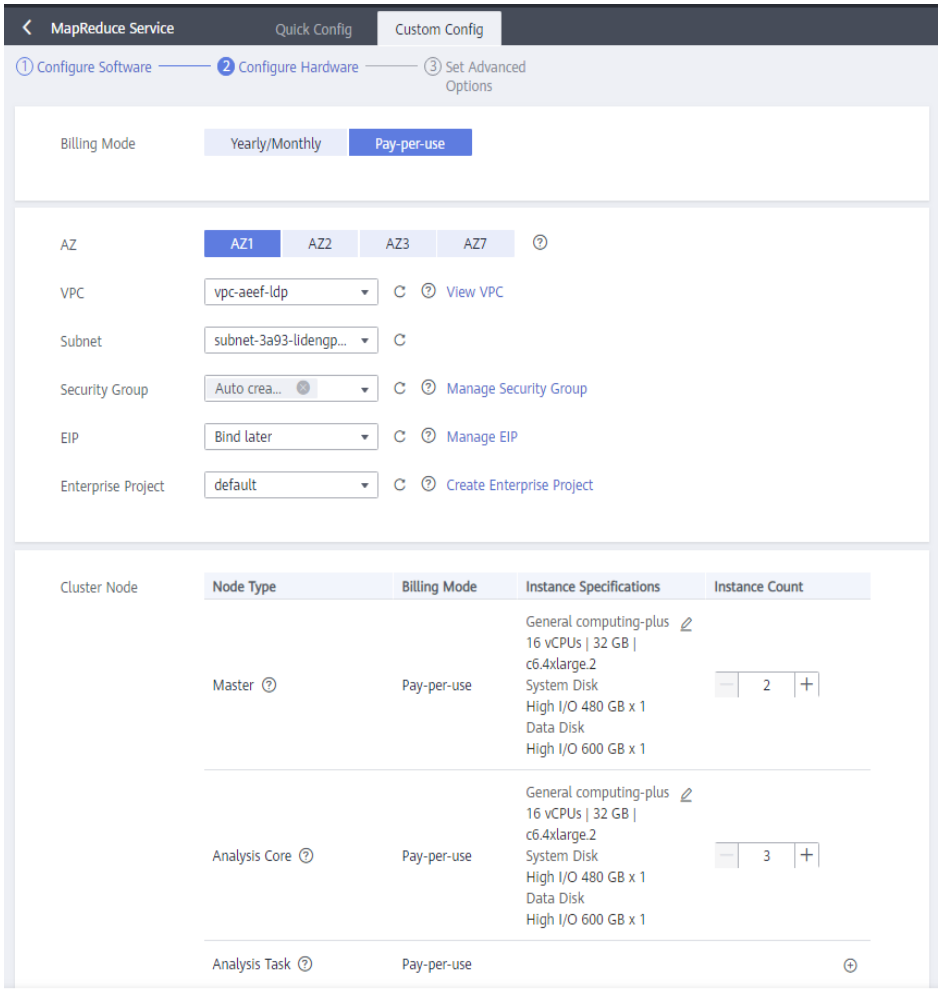

Step 5 Set advanced options by referring to **Advanced Options** and click **Buy Now**.

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

Figure 2-13 Set Advanced Options

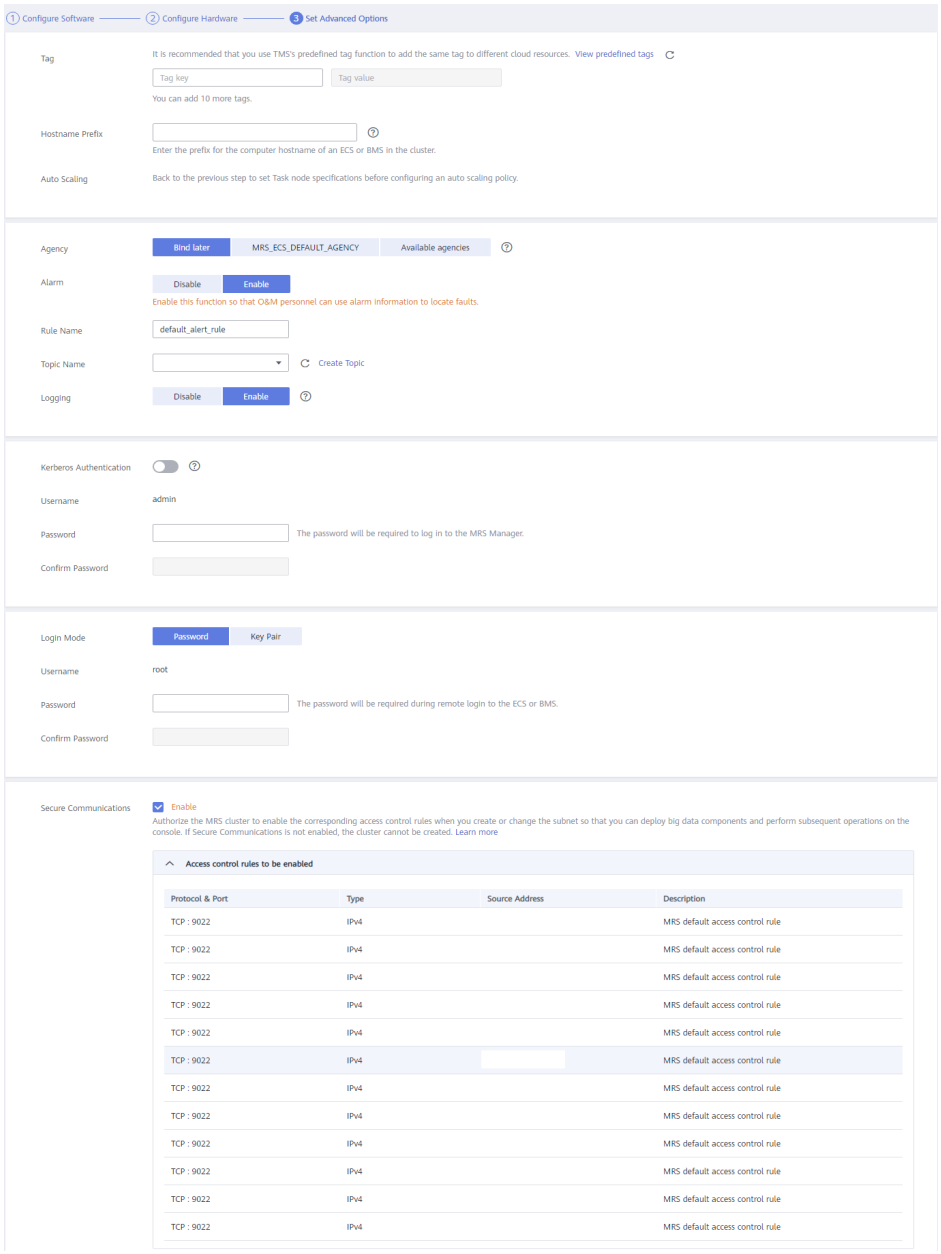

NOTE

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

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

For details about cluster status during creation, see the description of the status parameters in **Table 3-4**.

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

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

----End

Software Configurations

Table 2-1 MRS cluster software configuration

Parameter	Description
Region	Select a **region**. Cloud service products in different regions cannot communicate with each other over an intranet. For low network latency and quick access, select the nearest region.
Cluster Name	The cluster name must be unique. A cluster name can contain 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. The default name is **mrs**_*xxxx*. *xxxx* is a random collection of letters and digits.
Cluster Type	The cluster types are as follows: • **Analysis cluster**: is used for offline data analysis and provides Hadoop components. • **Streaming cluster**: is used for streaming tasks and provides stream processing components. • **Hybrid cluster**: is used for both offline data analysis and streaming processing and provides Hadoop components and streaming processing components. You are advised to use a hybrid cluster to perform offline data analysis and streaming processing tasks at the same time. • **Custom**: You can adjust the cluster service deployment mode based on service requirements. For details, see **Configuring Custom Topology**. (This type is currently available only in MRS 3.*x*.) **NOTE** • MRS streaming clusters do not support job and file management functions. • To install all components in a cluster, select **Custom**.

Parameter	Description
Version Type	The following version types are available: • Normal: – Supports basic cluster operations, such as configuration, management, and O&M. – Supports components such as Presto, Impala, Kudu, and Sqoop. • LTS: – In addition to basic cluster operations, the LTS version supports version upgrade. – Supports multi-AZ deployment. – Supports HetuEngine, IoTDB, and CDL. The default version type is **Normal**.
Cluster Version	Currently, MRS 1.9.2, 3.1.0, 3.1.5, 3.3.0-LTS.1, 3.2.0-LTS.1, and 3.1.2-LTS.3, and are supported. The latest version of MRS is used by default.
Component	MRS cluster components. For details about component versions supported by different versions of MRS clusters, see **List of MRS Component Versions**.
Metadata	Whether to use external data sources to store metadata. • **Local**: Metadata is stored in the local cluster. • **External data connection**: Metadata of external data sources is used. If the cluster is abnormal or deleted, metadata is not affected. This mode applies to scenarios where storage and compute are decoupled. Clusters that support the Hive or Ranger component support this function.
Component	This parameter is available only when **Metadata** is set to **External data connection**. It indicates the type of an external data source. • Hive • Ranger
Data Connection Type	This parameter is available only when **Metadata** is set to **External data connection**. It indicates the type of an external data source. When you create a cluster, **Data Connection Type** can only be set to **Local database**.

Parameter	Description
Component port (supported only for the LTS version)	Policy of the default communication port of each component in the MRS cluster. • Open source: Use the port provided by the open source component. • Custom: Customize a port for the component. For details about the differences between default open source port and default custom port, see **Web UIs of Open Source Components**.

Hardware Configurations

Table 2-2 MRS cluster billing configuration

Parameter	Description
Billing Mode	MRS provides two billing modes. • Yearly/Monthly • Pay-per-use
Required Duration	This parameter is valid in Yearly/Monthly billing mode and indicates a cluster subscription duration. The minimum cluster duration is 1 month and the maximum available cluster duration is 1 year. If **Auto-renew** is selected, monthly subscriptions are automatically renewed every month and yearly subscriptions are automatically renewed every year.

Table 2-3 MRS cluster hardware configuration

Parameter	Description
AZ	Select the **AZ** associated with the region of the cluster. An AZ is a physical area that uses independent power and network resources. AZs are physically isolated but interconnected through the internal network. This improves the availability of applications. You are advised to create clusters in different AZs.
VPC	A VPC is a secure, isolated, and logical network environment. Select the VPC for which you want to create a cluster and click **View VPC** to view the name and ID of the VPC. If no VPC is available, create one.

Parameter	Description
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. Select the subnet for which you want to create a cluster. Click **View Subnet** to view details about the selected subnet. If no subnet is created in the VPC, go to the VPC console and choose **Subnets** > **Create Subnet** to create one. For details about how to configure network ACL outbound rules, see **How Do I Configure a Network ACL Outbound Rule?** **NOTE** The number of IP addresses required by creating an MRS cluster depends on the number of cluster nodes and selected components, but not the cluster type. In MRS, IP addresses are automatically assigned to clusters during cluster creation basically based on the following formula: Quantity of IP addresses = Number of cluster nodes + 2 (one for Manager; one for the DB). In addition, if the Hadoop, Hue, Sqoop, and Presto or Loader and Presto components are selected during cluster deployment, one IP address is added for each component. To buy a ClickHouse cluster independently, the number of IP addresses required is calculated as follows: Number of IP addresses = Number of cluster nodes + 1 (for Manager).
Security Group	A security group is a set of ECS access rules. It provides access policies for ECSs that have the same security protection requirements and are mutually trusted in a VPC. When you create a cluster, you can select **Auto create** from the drop-down list of **Security Group** to create a security group or select an existing security group. **NOTE** When you select a security group created by yourself, ensure that the inbound rule contains a rule in which **Protocol** is set to **All**, **Port** is set to **All**, and **Source** is set to a trusted accessible IP address range. Do not use **0.0.0.0/0** as a source address. Otherwise, security risks may occur. If you do not know the trusted accessible IP address range, select **Auto create**.
EIP	After binding an EIP to an MRS cluster, you can use the EIP to access the Manager web UI of the cluster. When creating a cluster, you can select an available EIP from the drop-down list and bind it. If no EIP is available in the drop-down list, click **Manage EIP** to access the **EIPs** service page to buy one. **NOTE** The EIP must be in the same region as the cluster.

Parameter	Description
Enterprise Project	Select the enterprise project to which the cluster belongs. To use an enterprise project, create one on the **Enterprise** > **Project Management** page. The **Enterprise Management** console of the enterprise project is designed for resource management. It helps enterprises manage cloud-based personnel, resources, permissions, and finance in a hierarchical manner, such as management of companies, departments, and projects.

Table 2-4 Cluster node information

Parameter	Description
CPU Architecture	CPU architecture supported by MRS. This parameter is not available for MRS 3.1.0 and 3.1.5. • **x86**: The x86-based CPU architecture uses Complex Instruction Set Computing (CISC). Each instruction can be used to execute low-level hardware operations. The number of instructions is large, and the length of each instruction is different. Therefore, executing such an instruction is complex and time-consuming. • **Kunpeng**: The Kunpeng-based CPU architecture uses Reduced Instruction Set Computing (RISC). RISC is a microprocessor that executes fewer types of computer instructions but at a higher speed than CISC. RISC simplifies the computer architecture and improves the running speed. Compared with the x86-based CPU architecture, the Kunpeng-based CPU architecture has a more balanced performance and power consumption ratio. Kunpeng features high density, low power consumption, high cost-effectiveness.
Common Node Configurations	This parameter is available only when **Cluster Type** is set to **Custom**. Value options include **Compact**, **Full-size**, and **OMS-separate**. For details, see **Custom Cluster Template Description**.

Parameter	Description
Node Group	Name of a node group An MRS cluster consists of multiple ECS nodes. The system manages the nodes based on node groups. Nodes in a cluster are classified into the following types based on the roles of components deployed on the nodes: ● Master: manages the cluster and allocates cluster executable files to core nodes. traces the execution status of each job, and monitors the DataNode running status. ● Core: cluster worker node, which processes and analyzes data and stores process data. The system automatically creates a core node group based on the components contained in the cluster. For example, if you select the ClickHouse component, the system adds the ClickHouse node group and deploys the ClickHouseServer role in the node group by default. ● Task: provides compute resources, on which Yarn and Storm (supported only by MRS 1.9.2) are installed. Task nodes do not store persistent data. When compute resources in a cluster are insufficient, you can configure auto scaling policies to automatically increase task nodes. When the data volume change is small in a cluster but the cluster's service processing capabilities need to be remarkably and temporarily improved, add Task nodes to address the following situations: For clusters whose **Cluster Type** is **Analysis cluster**, **Streaming cluster** (supported only by MRS 1.9.2), and **Hybrid cluster**, the system automatically adds the corresponding task node groups. You can delete the task node groups if they are not required.
Node Type	Type of the nodes in the group. Options include **Core** and **Task**. **NOTE** If the node group type is set to **Task**, only the **NodeManager** role (except mandatory roles) can be deployed in the node group.
Billing Mode	Billing mode of nodes in a cluster. ● The billing mode of Master and Core nodes is the same as that of the cluster. ● The billing mode of Task nodes is fixed to Pay-per-use, that is, Task nodes in a yearly/monthly subscribed cluster are still billed in a pay-per-use basis.

Parameter	Description
Instance Specifications	Instance specifications of Master or Core nodes. MRS supports host specifications determined by CPU, memory, and disk space. Click  to configure the instance specifications, system disk, and data disk parameters of the cluster node. MRS supports BMS specifications only when the billing mode of a cluster is **Yearly/Monthly**. MRS supports the following hybrid deployment of **ECSs** and **BMSs**: • Master, Core, and Task nodes are deployed on ECSs. • Master and Core nodes are deployed on BMSs, and Task nodes are deployed on ECSs. • Master nodes are deployed on either ECSs or BMSs, Core nodes are deployed on either ECSs or BMSs, and Task nodes are deployed on ECSs. Tenants share physical resources of ECSs, but can exclusively use resources of BMSs. BMSs can better meet your requirements for deploying key applications and services that require high performance (such as big data clusters and enterprise middleware systems) and a secure and reliable running environment. If BMS specifications are used, Master node specifications cannot be scaled up. **NOTE** • More advanced instance specifications provide better data processing. However, they require higher cluster cost. • Instance specifications may vary in different AZs. If no instance specifications in the current AZ can meet your requirements, switch to another AZ. • If you select HDDs for Core nodes, there is no billing information for data disks. The fees are charged with ECSs. • If you select non-HDD disks for Core nodes, the disk types of Master and Core nodes are determined by **Data Disk**. • If **Sold out** appears next to an instance specification of a node, the node of this specification cannot be bought. You can only buy nodes of other specifications. • The Master node specification (4 vCPUs 8 GB) is not within the SLA after-sales scope. It is applicable only to the test environment and is not recommended for the production environment. • For MRS 3.*x* or later, the memory of the master node must be greater than 64 GB.

Parameter	Description
System Disk	Storage type and storage space of the system disk on a node. Storage type can be any of the following: ● **SAS**: high I/O ● **SSD**: ultra-high I/O ● **GPSSD**: general-purpose SSD
Data Disk	Data disk storage space of a node. For more data storage, you can add disks when creating a cluster. A maximum of 10 disks can be added to each Core or Task node. ● Data storage and computing are separated. Data is stored in OBS, which features low cost and unlimited storage capacity. The clusters can be deleted at any time in OBS. The computing performance is determined by OBS access performance and is lower than that of HDFS. This configuration is recommended if data computing is infrequent. ● Data storage and computing are not separated. Data is stored in HDFS, which features high cost, high computing performance, and limited storage capacity. Before deleting clusters, you must export and store the data. This configuration is recommended if data computing is frequent. The storage type can be any of the following: ● SAS: high I/O ● SSD: ultra-high I/O ● GPSSD: general-purpose SSD **NOTE** More nodes in a cluster require higher disk capacity of Master nodes. To ensure stable cluster running, set the disk capacity of the Master node to over 600 GB if the number of nodes is 300 and increase it to over 1 TB if the number of nodes reaches 500.

Parameter	Description
Instance Count	Number of Master and Core nodes. • Master Node Groups: The number of Master instances ranges from 3 to 9. • At least one Core node must exist and the total number of Core and Task nodes cannot exceed 10,000. Click ⊕ to add a node group, click ✎ to modify the node instance specifications, and click 🗑 to delete the added node group. **NOTE** A small number of nodes may cause clusters to run slowly while a large number of nodes may be unnecessarily costly. Set an appropriate value based on data to be processed.
LVM	This parameter is valid when a streaming Core node is created only. Click this parameter to enable or disable the disk LVM management function. This parameter is not available in MRS 3.*x* and later versions. If LVM is enabled, all disks on a node are mounted as logical volumes. This delivers more proper disk planning to avoid data skew, thereby improving system stability.
Topology Adjustment	If the deployment mode in the **Common Node** does not meet the requirements, set **Topology Adjustment** to **Enable** and adjust the instance deployment mode based on service requirements. For details, see **Topology Adjustment for a Custom Cluster**. This parameter is valid only when **Cluster Type** is set to **Custom**.

Advanced Options

Table 2-5 MRS cluster advanced configuration topology

Parameter	Description
Tag	For details, see **Adding a Tag to a Cluster/Node**.
Hostname Prefix	Enter the prefix for the computer hostname of an ECS or BMS in the cluster.
Auto Scaling	Auto scaling can be configured only after you specify task node specifications in the **Configure Hardware** step by referring to **Configuring Auto Scaling Rules**.
Bootstrap Action	For details, see **Adding a Bootstrap Action**.

Parameter	Description
Agency	By binding an agency, ECSs or BMSs can manage some of your resources. Determine whether to configure an agency based on the actual service scenario. For example, you can configure an agency of the ECS type to automatically obtain the AK/SK to access OBS. For details, see **Configuring a Storage-Compute Decoupled Cluster (Agency)**. The **MRS_ECS_DEFAULT_AGENCY** agency has the OBSOperateAccess permission of OBS and the CESFullAccess (for users who have enabled fine-grained policies), CES Administrator, and KMS Administrator permissions in the region where the cluster is located.
Metric Sharing	Monitoring metrics of big data components are collected. If a fault occurs when you use a cluster, share the monitoring metrics with Huawei Cloud technical support for troubleshooting. This parameter is not available in MRS 3.*x*.
System Disk Encryption	Whether to encrypt data in the system disks mounted to the cluster. This function is disabled by default. To use this function, you must have the Security Administrator and KMS Administrator permissions. The keys used to encrypt system disks are provided by the Key Management Service (KMS). You do not need to build or maintain the key management infrastructure. You can choose whether to enable system disk encryption by configuring this parameter.
System Disk Key ID	Key ID corresponding to the selected key name. This parameter is available only when **System Disk Encryption** is enabled.
System Disk Key Name	Name of the key used to encrypt the system disk. This parameter is mandatory when **System Disk Encryption** is enabled. By default, the default master key named **evs/default** is selected. You can select another master key from the drop-down list. If cloud disks are encrypted using a master key and it is then disabled or scheduled for deletion, the cloud disks can no longer be read from or written to, and data on these disks may never be restored. Exercise caution when performing this operation. Click **View Key List** to enter a page where you can create and manage keys.

Parameter	Description
Data Disk Encryption	Whether to encrypt data in the data disk mounted to the cluster. This function is disabled by default. To use this function, you must have the Security Administrator and KMS Administrator permissions. The keys used to encrypt data disks are provided by Key Management Service (KMS). You do not need to build and maintain the key management infrastructure. Click **Data Disk Encryption** to enable or disable the data disk encryption function.
Data Disk Key ID	This parameter is displayed only when the **Data Disk Encryption** function is enabled. This parameter indicates the key ID corresponding to the selected key name.
Data Disk Key Name	This parameter is mandatory when the **Data Disk Encryption** function is enabled. Select the name of the key used to encrypt the data disk. By default, the default master key named **evs/default** is selected. You can select another master key from the drop-down list. If disks are encrypted using a CMK, which is then disabled or scheduled for deletion, the disks can no longer be read from or written to, and data on these disks may never be restored. Exercise caution when performing this operation. Click **View Key List** to enter a page where you can create and manage keys.
Alarm	If the alarm function is enabled, the cluster maintenance personnel can be notified in a timely manner to locate faults when the cluster runs abnormally or the system is faulty.
Rule Name	Name of the rule for sending alarm messages. The value can contain only digits, letters, hyphens (-), and underscores (_).
Topic Name	Select an existing topic or click **Create Topic** to create a topic. To deliver messages published to a topic, you need to add a subscriber to the topic. For details, see **Adding Subscriptions to a Topic**. A topic serves as a message sending channel, where publishers and subscribers can interact with each other.

Parameter	Description
Logging	Whether to collect logs when cluster creation fails. If this function is enabled, system logs and component run logs will be automatically recorded and saved to OBS when a cluster fails to be created or to be scaled. The log files are used by O&M engineers to quickly locate faults. The log information is retained for a maximum of seven days.
Kerberos Authentication	Whether to enable Kerberos authentication when logging in to Manager. After buying a cluster, you cannot enable or disable Kerberos authentication. •  : If **Kerberos Authentication** is disabled, common users can use all functions of an MRS cluster. You are advised to disable Kerberos authentication in single-user scenarios. If Kerberos authentication is disabled, you can follow instructions in **Security Configuration Suggestions for Clusters with Kerberos Authentication Disabled** to perform security configuration. •  : If **Kerberos Authentication** is enabled, common users cannot use the file and job management functions of an MRS cluster and cannot view cluster resource usage or the job records for Hadoop and Spark. To use more cluster functions, the users must contact the Manager administrator to assign more permissions. You are advised to enable Kerberos authentication in multi-user scenarios. • Currently, Presto does not support Kerberos authentication.
Username	Name of the administrator of Manager. **admin** is used by default.

Parameter	Description
Password	Password of the Manager administrator The following requirements must be met: ● Must contain 8 to 26 characters. ● Must contain at least four of the following: – Lowercase letters – Uppercase letters – Digits – At least one of the following special characters: `~!@#$%^&*()-_=+|[{}];:'",<.>/? ● Cannot be the same as the username or the username spelled backwards. Password Strength: The colorbar in red, orange, and green indicates weak, medium, and strong password, respectively.
Confirm Password	Enter the password of the Manager administrator again.

Parameter	Description
Login Mode	• Password You can log in to ECS nodes using a password. A password must meet the following requirements: 1. Must be a string and 8 to 26 characters long. 2. Must contain at least four of the following: uppercase letters, lowercase letters, digits, and special characters (`~!@#$%^&*()-_=+|[{}];:',<.>/?). 3. The password cannot be the username or the reverse username. • Key Pair Key pairs are used to log in to ECS nodes of the cluster. Select a key pair from the drop-down list. Select "I acknowledge that I have obtained private key file *SSHkey-xxx* and that without this file I will not be able to log in to my ECS." If you have never created a key pair, click **View Key Pair** to create or import a key pair. And then, obtain a private key file. A key pair, also called an SSH key, consists of a public key and a private key. You can create an SSH key and download the private key for authenticating remote login. For security, a private key can only be downloaded once. Keep it secure. Use an SSH key in either of the following two methods: 1. Creating an SSH key: After you create an SSH key, a public key and a private key are generated. The public key is stored in the system, and the private key is stored in the local ECS. When you log in to an ECS, the public and private keys are used for authentication. 2. Importing an SSH key: If you have obtained the public and private keys, import the public key into the system. When you log in to an ECS, the public and private keys are used for authentication.

Parameter	Description
Key Pair	Key pairs are used to log in to ECS nodes of the cluster. Select a key pair from the drop-down list. Select "I acknowledge that I have obtained private key file *SSHkey-xxx* and that without this file I will not be able to log in to my ECS." If you have never created a key pair, click **View Key Pair** to create or import a key pair. And then, obtain a private key file. A key pair, also called an SSH key, consists of a public key and a private key. You can create an SSH key and download the private key for authenticating remote login. For security, a private key can only be downloaded once. Keep it secure. Use an SSH key in either of the following two methods: 1. Creating an SSH key: After you create an SSH key, a public key and a private key are generated. The public key is stored in the system, and the private key is stored in the local ECS. When you log in to an ECS, the public and private keys are used for authentication. 2. Importing an SSH key: If you have obtained the public and private keys, import the public key into the system. When you log in to an ECS, the public and private keys are used for authentication.
Secure Communications	MRS clusters provision, manage, and use big data components through the management console. Big data components are deployed in a user's VPC. If the MRS management console needs to directly access big data components deployed in the user's VPC, you need to enable the corresponding security group rules after you have obtained user authorization. This authorization process is called secure communications. For details, see **Communication Security Authorization**. If the secure communications function is not enabled, MRS clusters cannot be created.

Failed to Create a Cluster

If a cluster fails to be created, the failed task will be managed on the **Manage Failed Tasks** page. Choose **Clusters** > **Active Clusters**. Click  in **Figure 2-14** to go to the **Manage Failed Tasks** page. In the **Task Status** column, hover your cursor over  to view the failure cause. See **Figure 2-15**. You can delete failed tasks by referring to **Viewing Failed MRS Tasks**.

Figure 2-14 Failed task management

Figure 2-15 Failure cause

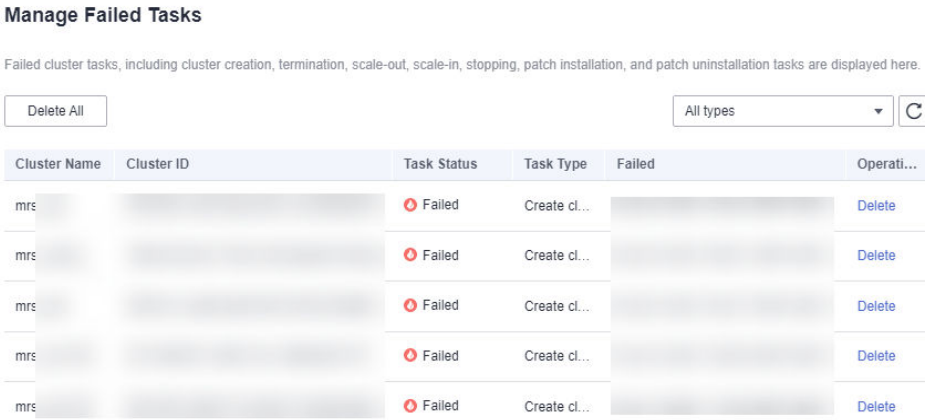

Table 2-6 lists the error codes of MRS cluster creation failures.

Table 2-6 Error codes

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

2.4 Configuring Custom Topology

The analysis cluster, streaming cluster, and hybrid cluster provided by MRS use fixed templates to deploy cluster processes. Therefore, you cannot customize service processes on management nodes and control nodes. If you want to customize the cluster deployment, set **Cluster Type** to **Custom** when creating a cluster. In this way, you can customize the deployment mode of process instances

on the management nodes and control nodes in the cluster. Only MRS 3.x and later versions support the creation of clusters in a custom topology.

A custom cluster provides the following functions:

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

Roles in an MRS cluster:

- Management Node (MN): is the node to install Manager (the management system of the MRS cluster). It provides a unified access entry. Manager centrally manages nodes and services deployed in the cluster.
- Control Node (CN): controls and monitors how data nodes store and receive data, and send process status, and provides other public functions. Control nodes of MRS include HMaster, HiveServer, ResourceManager, NameNode, JournalNode, and SlapdServer.
- Data Node (DN): A data node executes the instructions sent by the management node, reports task status, stores data, and provides other public functions. Data nodes of MRS include DataNode, RegionServer, and NodeManager.

Customizing a Cluster

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

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

Step 3 Configure basic cluster information. For details about the parameters, see **Software Configurations**.

- **Region**: Retain the default value.
- **Billing Mode**: Select **Yearly/Monthly** or **Pay-per-use**.
- **Cluster Name**: You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing, for example, **mrs_20180321**.
- **Cluster Version**: Currently, only MRS 3.*x* is available.
- **Cluster Type**: Select **Custom** and select components as required.

Step 4 Click **Next**. Configure hardware information.

- **AZ**: Retain the default value.
- **Enterprise Project**: Retain the default value.
- **VPC**: Retain the default value. If there is no available VPC, click **View VPC** to access the VPC console and create a new VPC.
- **Subnet**: Retain the default value.
- **Security Group**: Select **Auto create**.
- **EIP**: Select **Bind later**.

- **Common Node**: For details, see **Custom Cluster Template Description**.
- **Node Count**: Adjust the number of cluster instances based on the service volume. For details, see **Table 2-8**.
- **Instance Specifications**: Click  to configure the instance specifications, system disk and data disk storage types, and storage space.
- **Topology Adjustment**: If the deployment mode in the **Common Node** does not meet the requirements, you need to manually install some instances that are not deployed by default, or you need to manually install some instances, set **Topology Adjustment** to **Enable** and adjust the instance deployment mode based on service requirements. For details, see **Topology Adjustment for a Custom Cluster**.

Step 5 Click **Next** and set advanced options.

For details about the parameters, see **Advanced Options**.

Step 6 Click **Buy Now**.

If Kerberos authentication is enabled for a cluster, check whether Kerberos authentication is required. If yes, click **Continue**. If no, click **Back** to disable Kerberos authentication and then create a cluster.

Step 7 Click **Back to Cluster List** to view the cluster status.

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

----End

Custom Cluster Template Description

Table 2-7 Common templates for custom clusters

Common Node	Description	Node Range
Compact	The management role and control role are deployed on the Master node, and data instances are deployed in the same node group. This deployment mode applies to scenarios where the number of control nodes is less than 100, reducing costs.	• The number of Master nodes is greater than or equal to 3 and less than or equal to 11. • The total number of node groups is less than or equal to 10, and the total number of nodes in non-Master node groups is less than or equal to 10,000.

Common Node	Description	Node Range
OMS-separate	The management role and control role are deployed on different Master nodes, and data instances are deployed in the same node group. This deployment mode is applicable to a cluster with 100 to 500 nodes and delivers better performance in high-concurrency load scenarios.	• The number of Master nodes is greater than or equal to 5 and less than or equal to 11. • The total number of node groups is less than or equal to 10, and the total number of nodes in non-Master node groups is less than or equal to 10,000.
Full-size	The management role and control role are deployed on different Master nodes, and data instances are deployed in different node groups. This deployment mode is applicable to a cluster with more than 500 nodes. Components can be deployed separately, which can be used for a larger cluster scale.	• The number of Master nodes is greater than or equal to 9 and less than or equal to 11. • The total number of node groups is less than or equal to 10, and the total number of nodes in non-Master node groups is less than or equal to 10,000.

Table 2-8 Node deployment scheme of a customized MRS cluster

Node Deployment Principle		Applicable Scenario	Networking Rule
Management nodes, control nodes, and data nodes are deployed separately. (This scheme requires at least eight nodes.)	MN × 2 + CN × 9 + DN × n	(Recommended) This scheme is used when the number of data nodes is 500–2000.	• If the number of nodes in a cluster exceeds 200, the nodes are distributed to different subnets and the subnets are interconnected with each other in Layer 3 using core switches. Each subnet can contain a maximum of 200 nodes and the allocation of nodes to different subnets must be balanced. • If the number of nodes is less than 200, the nodes in the cluster are deployed in the same subnet and the nodes are interconnected with each other in Layer 2 using aggregation switches.
	MN × 2 + CN × 5 + DN × n	(Recommended) This scheme is used when the number of data nodes is 100–500.	
	MN × 2 + CN × 3 + DN × n	(Recommended) This scheme is used when the number of data nodes is 30–100.	
The management nodes and control nodes are deployed together, and the data nodes are deployed separately.	(MN+CN) × 3 + DN × n	(Recommended) This scheme is used when the number of data nodes is 3–30.	Nodes in the cluster are deployed in the same subnet and are interconnected with each other at Layer 2 through aggregation switches.

Node Deployment Principle	Applicable Scenario	Networking Rule
The management nodes, control nodes, and data nodes are deployed together.	• This scheme is applicable to a cluster having fewer than 6 nodes. • This scheme requires at least three nodes. **NOTE** This template is not recommended in the production environment or commercial environment. • If management, control, and data nodes are co-deployed, cluster performance and reliability are greatly affected. • If the number of nodes meet the requirements, deploy data nodes separately. • If the number of nodes is insufficient to support separately deployed data nodes, use the dual-plane networking mode for this scenario. The traffic of the management network is isolated from that of the service network to prevent excessive data volumes on the service plane, ensuring correct delivery of management operations.	Nodes in the cluster are deployed in the same subnet and are interconnected with each other at Layer 2 through aggregation switches.

Topology Adjustment for a Custom Cluster

Table 2-9 Topology adjustment

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

Service	Dependency	Role	Role Deployment Suggestions	Description
ClickHouse	Depends on ZooKeeper.	CHS (ClickHouseServer)	This role can be deployed on all nodes. Number of role instances to be deployed: an even number ranging from 2 to 256	A non-Master node group with this role assigned is considered as a Core node.
		CLB (ClickHouseBalancer)	This role can be deployed on all nodes. Number of role instances to be deployed: 2 to 256	-
ZooKeeper	-	QP(quorumpeer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 3 to 9, with the step size of 2	-
Hadoop	Depends on ZooKeeper.	NN(NameNode)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2	-
		HFS (HttpFS)	This role can be deployed on the Master node only. Number of role instances to be deployed: 0 to 10	-
		JN(JournalNode)	This role can be deployed on the Master node only. Number of role instances to be deployed: 3 to 60, with the step size of 2	-

Service	Dependency	Role	Role Deployment Suggestions	Description
		DN(DataNode)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000	A non-Master node group with this role assigned is considered as a Core node.
		RM(ResourceManager)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2	-
		NM(NodeManager)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000	-
		JHS(JobHistoryServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 2	-
		TLS(TimelineServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 0 to 1	-
Presto	Depends on Hive.	PCD(Coordinator)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2	-

Service	Dependency	Role	Role Deployment Suggestions	Description
		PWK(Worker)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000	-
Spark2x	• Depends on Hadoop. • Depends on Hive. • Depends on ZooKeeper.	JS2X(JDBCServer2x)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 10	-
		JH2X(JobHistory2x)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2	-
		SR2X(SparkResource2x)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 50	-
		IS2X(IndexServer2x)	(Optional) This role can be deployed on the Master node only. Number of role instances to be deployed: 0 to 2, with the step size of 2	-
HBase	Depends on Hadoop.	HM(HMaster)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2	-

Service	Dependency	Role	Role Deployment Suggestions	Description
		TS(ThriftServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 10,000	-
		RT(RESTServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 10,000	-
		RS(RegionServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000	-
		TS1(Thrift1Server)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 10,000	If the Hue service is installed in a cluster and HBase needs to be used on the Hue web UI, install this instance for the HBase service.
Hive	• Depends on Hadoop. • Depends on DBService.	MS(MetaStore)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 10	-
		WH(WebHCat)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 10	-

Service	Dependency	Role	Role Deployment Suggestions	Description
		HS(HiveServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2 to 80	-
Hue	Depends on DBService.	H(Hue)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2	-
Sqoop	Depends on Hadoop.	SC(Sqoop Client)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000	-
Kafka	Depends on ZooKeeper.	B(Broker)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000	-
Flume	-	MS(MonitorServer)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 2	-
		F(Flume)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000	A non-Master node group with this role assigned is considered as a Core node.

Service	Dependency	Role	Role Deployment Suggestions	Description
Tez	• Depends on Hadoop. • Depends on DBService. • Depends on ZooKeeper.	TUI(TezUI)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 2	-
Flink	• Depends on ZooKeeper. • Depends on KrbServer. • Depends on DBService. • Depends on Hadoop.	FR(FlinkResource)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000	-
		FS(FlinkServer)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 2	-
Oozie	• Depends on Hadoop. • Depends on DBService. • Depends on ZooKeeper.	O(oozie)	This role can be deployed on the Master node only. Number of role instances to be deployed: 2	-

Service	Dependency	Role	Role Deployment Suggestions	Description
Impala	• Depends on Hadoop. • Depends on Hive. • Depends on DBService. • Depends on ZooKeeper.	StateStore	This role can be deployed on the Master node only. Number of role instances to be deployed: 1	-
		Catalog	This role can be deployed on the Master node only. Number of role instances to be deployed: 1	-
		Impalad	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 10,000	-
Kudu	-	KuduMaster	This role can be deployed on the Master node only. Number of role instances to be deployed: 3 or 5	-
		KuduTserver	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 10,000	-
Ranger	Depends on DBService.	RA(RangerAdmin)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1 to 2	-

Service	Dependency	Role	Role Deployment Suggestions	Description
		USC(User Sync)	This role can be deployed on the Master node only. Number of role instances to be deployed: 1	-
		TSC (TagSync)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 1	-
HetuEngine (available for MRS 3.1.2-LTS.3 and later versions only)	• Depends on Hadoop. • Depends on DBService. • Depends on Hive. • Depends on ZooKeeper. • Depends on KrbServer. • Depends on Yarn. • Depends on HDFS.	HSB(HSBroker)	This role can be deployed on all nodes. Number of role instances to be deployed: 2 to 50	-
		HSC(HSConsole)	This role can be deployed on all nodes. Number of role instances to be deployed: 2	-
		HSF(HSFabric)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 50	-
		QAS (available for MRS 3.2.0-LTS.1 and later versions only)	This role can be deployed on all nodes. Number of role instances to be deployed: 0 to 2	-

Service	Dependency	Role	Role Deployment Suggestions	Description
IoTDB (available for MRS 3.2.0-LTS.1 and later versions only)	Depends on KrbServer.	ConfigNode (CN)	This role can be deployed on Master nodes only. Number of role instances to be deployed: 3 to 9, with the step size of 2	-
		IoTDBServer (IoTDBS)	This role can be deployed on all nodes. Number of role instances to be deployed: 3 to 256	-
CDL (available for MRS 3.2.0-LTS.1 and later versions only)	• Depends on DBService. • Depends on HDFS. • Depends on Hive. • Depends on KrbServer. • Depends on Kafka. • Depends on Spark. • Depends on ZooKeeper. • Depends on Yarn.	CDLConnector (CC)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 256	-
		CDLService (CS)	This role can be deployed on all nodes. Number of role instances to be deployed: 1 to 2	-

2.5 Adding a Tag to a Cluster/Node

Tags are used to identify clusters/nodes. Adding tags to clusters/nodes can help you identify and manage your resources.

- Cluster tags: You can add up to 10 tags to a cluster during cluster creation or add them on the details page of a created cluster. Updating a cluster tag will synchronize the tag to all nodes in the cluster.
- Node tags: You can use the default tag or add tags to nodes in an MRS cluster when you configure an auto scaling policy. Node tags take the tag quotas. You can view the tags of a node in the **Nodes** tab on MRS console.
- Default tags: An MRS cluster contains multiple nodes, and each node is an ECS and contains EVS disks. After the default tag is enabled, the system automatically creates a cluster tag and a tag for each node. The default tag is automatically synchronized to the corresponding ECS or EVS instances.

  To view node tags, go to the **Nodes** tab on the MRS console, and move the cursor to the tag icon of a node in the node list.

  NOTE

  - MRS tag updates are synchronized to the ECSs or EVS disks in the cluster. However, if you modify MRS cluster tags on the ECS or EVS console, the modification will not be synchronized to MRS. To ensure tag consistency, do not modify MRS cluster tags on the ECS or EVS console.
  - You can add a maximum of 10 tags to a cluster. If the number of tags of a node in the cluster reaches the upper limit, no more tags can be added to the cluster.
  - If default tags are enabled, a default tag is added to the cluster and each node, which takes two quotas. That is, a maximum of 10 tags can be added by default. In this case, a maximum of eight more tags can still be added.

If your organization has configured tag strategies for MRS, add tags to clusters/nodes based on the strategies. If a tag does not comply with the tag strategies, the cluster/node may be failed to be created. Contact the organization administrator to learn more about the tag strategies.

A tag consists of a tag key and a tag value. **Table 2-10** provides tag key and value requirements.

Table 2-10 Tag key and value requirements

Parameter	Requirement	Example
Key	A tag key cannot be left blank. A tag key must be unique in a cluster. A tag key contains a maximum of 36 characters. A tag value cannot contain special characters (=*<>\,|/) or start or end with spaces.	Organization
Value	A tag value contains a maximum of 43 characters. A tag value cannot contain special characters (=*<>\,|/) or start or end with spaces. This parameter can be left blank.	Apache

Adding Tags to a Cluster

- Adding cluster tags during cluster creation
  a. Log in to the MRS console.
  b. Click **Buy Cluster**. The page for buying a cluster is displayed.
  c. Click the **Custom Config** tab.
  d. Configure the cluster software and hardware by referring to **Buying a Custom Cluster**.
  e. Select **Configure** on the right of **Set Advanced Options** and enter the key and value of a new tag.

  Figure 2-16 Adding a tag to a cluster

  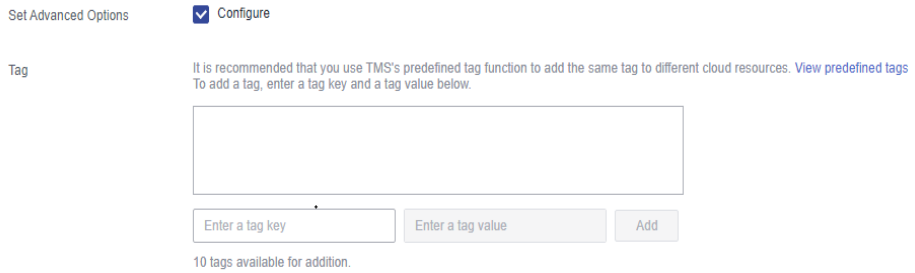

- Adding tags to an existing cluster

a. Log in to the MRS management console.

b. In the navigation pane on the left, choose **Clusters** > **Active Clusters**. Click the name of a running cluster. The basic information page of the cluster is displayed.

c. Click the **Tags** tab.

d. Click **Add/Edit Tag**. If this is your first time adding a tag, click **Add Tag**. In the displayed dialog box, enter the key and value of a tag, and click **Add**.

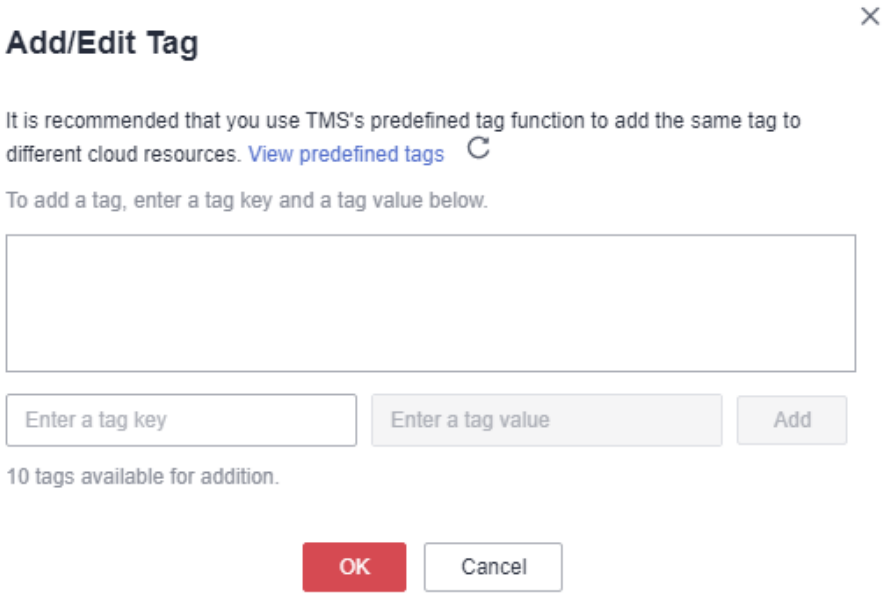

NOTE

You can also add cluster tags by enabling default tags. All nodes will be tagged with the cluster ID and node IDs, which takes two quotas.

e. Click **OK**.

Adding Tags to a Node

- Node tags are automatically added when a default tag is added to a cluster. For details, see **Adding tags to an existing cluster**.
- Adding node tags for auto scaling

  If you add a tag when configuring an auto scaling policy, MRS automatically adds the tag to the new nodes and synchronizes the tag to the ECSs and EVS disks.

  a. Log in to the MRS management console.

  b. In the navigation pane on the left, choose **Clusters** > **Active Clusters**. Click the name of a running cluster. The basic information page of the cluster is displayed.

  c. On the page that is displayed, click the **Auto Scaling** tab.

  d. Click **Edit** on the right of an existing auto scaling policy. In the displayed dialog box, enter the key and value of the tag you want to add, and click **Add**.

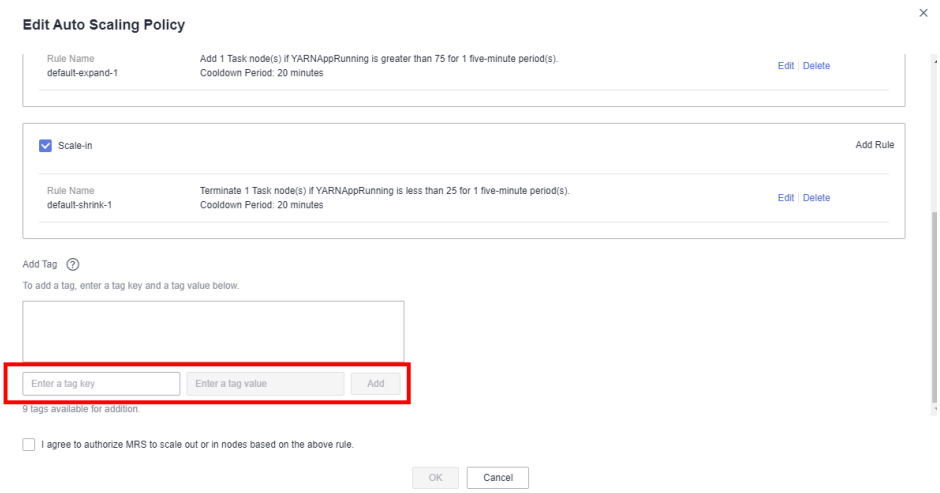

NOTE

- You need to enable the auto scaling policy and configure scale-out rules. Otherwise, the node tags will not take effect.
- If tag quotas are insufficient, delete the cluster tag or modify existing a tag of the auto scaling policy, and then enable the policy.
- Tags cannot be added to auto scaling policies of resource pools.

e. Click **OK**.

Searching for Target Clusters by Tags

On the **Active Clusters** page, search for the target cluster by tag key or tag value.

1. Log in to the MRS console.
2. In the upper right corner of the **Active Clusters** page, click **Search by Tag** to access the search page.
3. Enter the tag of the cluster to be searched.

   You can select a tag key or tag value from their drop-down lists. When the tag key or tag value is exactly matched, the system can automatically locate the target cluster. If you enter multiple tags, their intersections are used to search for the cluster.
4. Click **Search**.

   The system searches for the target cluster by tag key or value.

Managing Tags

You can view, add, and delete tags on the **Tags** tab page of the cluster.

1. Log in to the MRS console.
2. On the **Active Clusters** page, click the name of a cluster for which you want to manage tags.

   The cluster details page is displayed.
3. Click the **Tags** tab and view, add, and delete tags on the tab page.
   - View

     On the **Tags** tab page, you can view details about tags of the cluster, including the number of tags and the key and value of each tag.

- Add

  Click **Add/Edit Tag** in the upper left corner. In the displayed **Add/Edit Tag** dialog box, enter the key and value of the tag to be added, and click **OK**.

- Delete

  Locate the row that contains the tag you want to delete and click **Delete** in the **Operation** column. In the displayed **Delete Tag** dialog box, enter **DELETE**, and click **Yes**.

2.6 Communication Security Authorization

MRS clusters provision, manage, and use big data components through the management console. Big data components are deployed in a user's VPC. If the MRS management console needs to directly access big data components deployed in the user's VPC, you need to enable the corresponding security group rules after you have obtained user authorization. This authorization process is called secure communications.

If the secure communications function is not enabled, MRS clusters cannot be created. If you disable the communication after a cluster is created, the cluster status will be **Network channel is not authorized** and the following functions will be affected:

- Functions, such as big data component installation, cluster scale-out/scale-in, and Master node specification upgrade, are unavailable.
- The cluster running status, alarms, and events cannot be monitored.
- The node management, component management, alarm management, file management, job management, patch management, and tenant management functions on the cluster details page are unavailable.
- The Manager page and the website of each component cannot be accessed.

After the secure communications function is enabled again, the cluster status is restored to **Running**, and the preceding functions become available. For details, see **Enabling Secure Communications for Clusters with This Function Disabled**.

If the security group rules authorized in the cluster are insufficient for you to provision, manage, and use big data components, is displayed on the right of **Secure Communications**. In this case, click **Update** to update the security group rules. For details, see **Update**.

Enabling Secure Communications During Cluster Creation

Step 1 Log in to the MRS console.

Step 2 Click **Buy Cluster**. The page for buying a cluster is displayed.

Step 3 On the displayed page, select **Quick Config**.

Step 4 Configure cluster information by referring to **Quick Configuration** or **Buying a Custom Cluster**.

Step 5 Select the check box for **Secure Communications**.

Figure 2-17 Secure communications

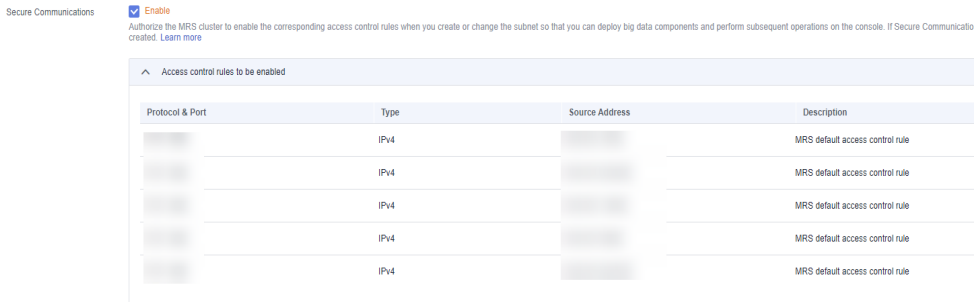

Step 6 Click **Buy Now**.

If Kerberos authentication is enabled for a cluster, check whether Kerberos authentication is required. If yes, click **Continue**. If no, click **Back** to disable Kerberos authentication and then create a cluster.

----End

Disabling Secure Communications After a Cluster Is Created

Step 1 Log in to the MRS console.

Step 2 In the active cluster list, click the name of the cluster for which you want to disable secure communications.

The cluster details page is displayed.

Figure 2-18 Secure communications

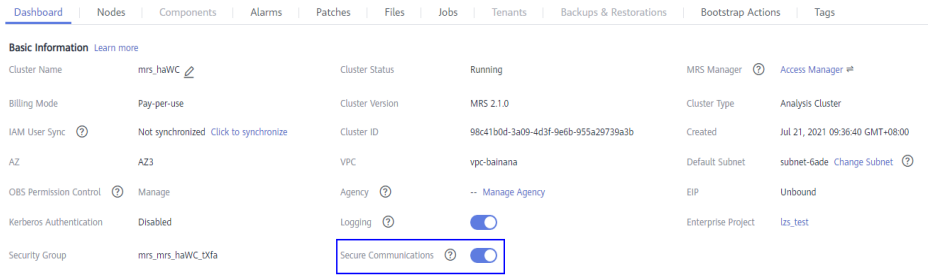

Step 3 Click the switch on the right of **Secure Communications** to disable authorization. In the dialog box that is displayed, click **OK**.

After the authorization is disabled, the cluster status changes to **Network channel unauthorized**, and some functions of the cluster are unavailable. Exercise caution when performing this operation.

Figure 2-19 Disabling secure communications

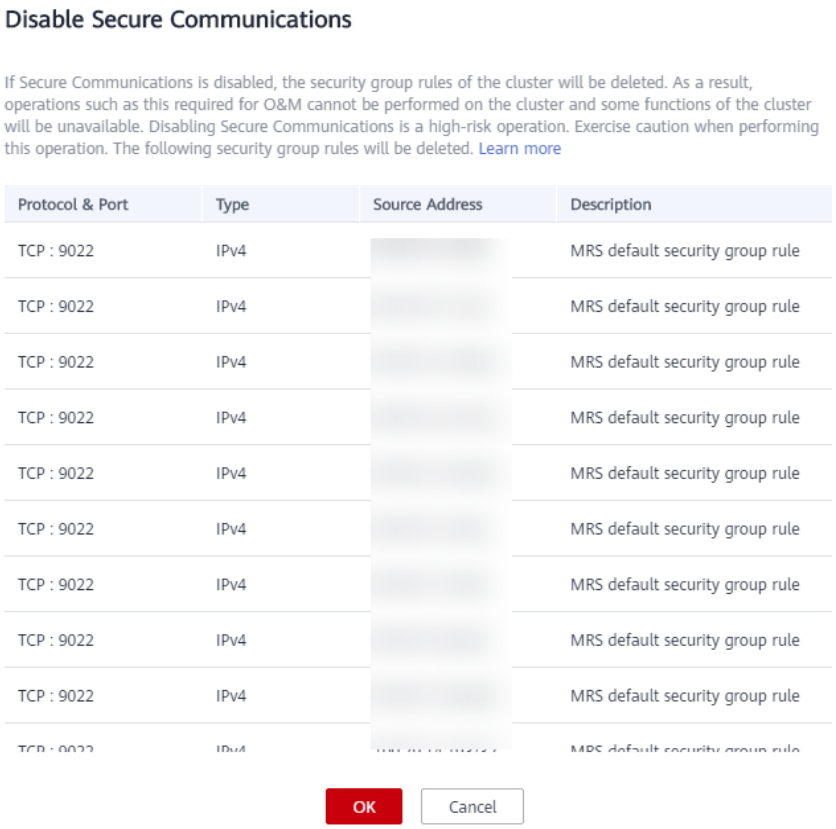

Step 4 If you have enabled the critical operation protection function (for details, see **Critical Operation Protection** on IAM), enter the verification code obtained in the corresponding verification mode to avoid risks and losses caused by misoperations.

Figure 2-20 Identity verification

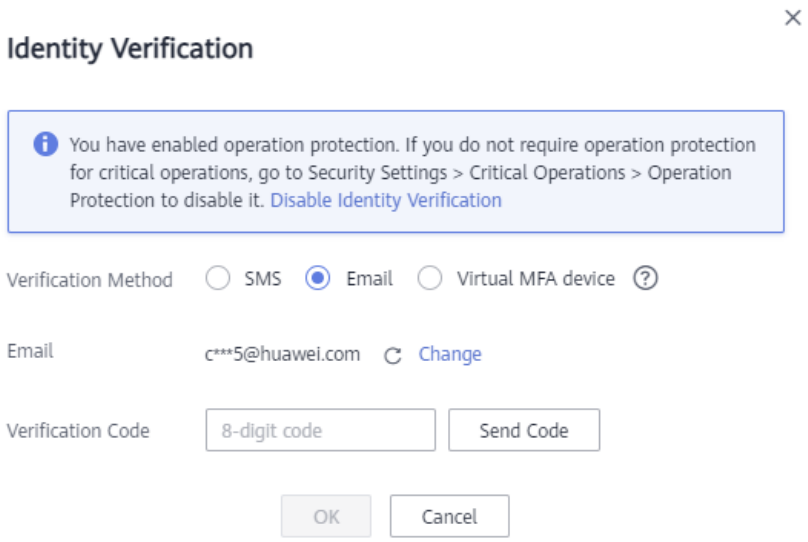

----End

Enabling Secure Communications for Clusters with This Function Disabled

Step 1 Log in to the MRS console.

Step 2 In the active cluster list, click the name of the cluster for which you want to enable secure communications.

The cluster details page is displayed.

Step 3 Click the switch on the right of **Secure Communications** to enable the function.

After the function is enabled, the cluster status changes to **Running**.

Figure 2-21 Enabling secure communications

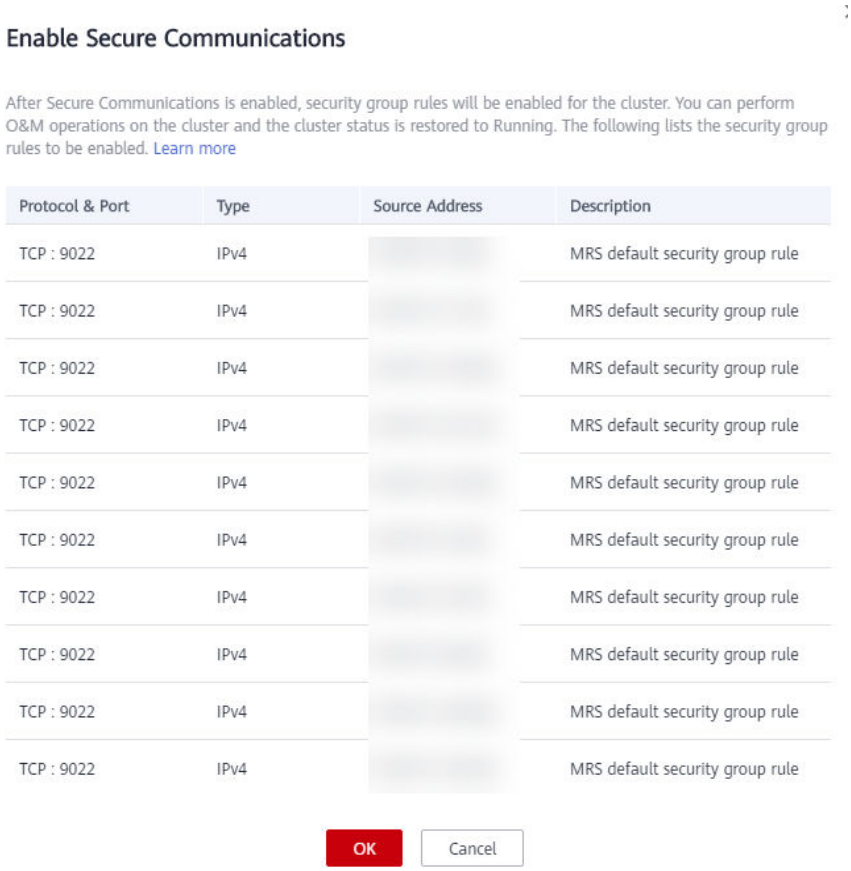

----End

Update

If the security group rules authorized in the cluster are insufficient for you to provision, manage, and use big data components, is displayed on the right of **Secure Communications**. In this case, click **Update** to update the security group rules. For details, see **Update**.

Step 1 Log in to the MRS console.

Step 2 In the active cluster list, click the name of the cluster for which you want to update secure communications.

The cluster details page is displayed.

Step 3 Click **Update** on the right of **Secure Communications**.

Figure 2-22 Update

Step 4 Click **OK**.

----End

2.7 Configuring Auto Scaling Rules

2.7.1 Overview

In big data application scenarios, especially real-time data analysis and processing, the number of cluster nodes needs to be dynamically adjusted according to data volume changes to provide the required number of resources. The auto scaling function of MRS enables the task nodes of a cluster to be automatically scaled to match cluster loads. If the data volume changes periodically, you can configure an auto scaling rule so that the number of task nodes can be automatically adjusted in a fixed period of time before the data volume changes.

- Auto scaling rules: You can increase or decrease task nodes based on real-time cluster loads. Auto scaling will be triggered with a certain delay when the data volume changes.
- Resource plans: Set the task node quantity based on the time range. If the data volume changes periodically, you can create resource plans to resize the cluster before the data volume changes, thereby avoiding delays in increasing or decreasing resources.

You can configure either auto scaling rules or resource plans or both to trigger auto scaling. Configuring both resource plans and auto scaling rules improves the cluster node scalability to cope with occasionally unexpected data volume peaks.

In some service scenarios, resources need to be reallocated or service logic needs to be modified after cluster scale-out or scale-in. If you manually scale out or scale in a cluster, you can log in to cluster nodes to reallocate resources or modify service logic. If you use auto scaling, MRS enables you to customize automation scripts for resource reallocation and service logic modification. Automation scripts can be executed before and after auto scaling and automatically adapt to service load changes, all of which eliminates manual operations. In addition, automation scripts can be fully customized and executed at various moments, meeting your personalized requirements and improving auto scaling flexibility.

- Auto scaling rules:
  - You can set a maximum of five rules for scaling out or in a cluster, respectively.
  - The system determines the scale-out and then scale-in based on your configuration sequence. Important policies take precedence over other

policies to prevent repeated triggering when the expected effect cannot be achieved after a scale-out or scale-in.

  - Comparison factors include greater than, greater than or equal to, less than, and less than or equal to.
  - Cluster scale-out or scale-in can be triggered only after the configured metric threshold is reached for consecutive 5*n* (the default value of *n* is 1) minutes.
  - After each scale-out or scale-in, there is a cooling duration that is greater than 0 and lasts 20 minutes by defaults.
  - In each cluster scale-out or scale-in, at least one node and at most 100 nodes can be added or reduced.
  - The number of task nodes in a cluster is limited to the default number of nodes configured by users or the node quantity range in the resource plan that takes effect in the current time period. The node quantity range in the resource plan that takes effect in the current time period has a higher priority.

- Resource plans (setting the number of Task nodes by time range):
  - You can specify a Task node range (minimum number to maximum number) in a time range. If the number of Task nodes is beyond the Task node range in a resource plan, the system triggers cluster scale-out or scale-in.
  - You can set a maximum of five resource plans for a cluster.
  - A resource plan cycle is by day. The start time and end time can be set to any time point between 00:00 and 23:59. The start time must be at least 30 minutes earlier than the end time. Time ranges configured for different resource plans cannot overlap.
  - After a resource plan triggers cluster scale-out or scale-in, there is 10-minute cooling duration. Auto scaling will not be triggered again within the cooling time.
  - When a resource plan is enabled, the number of Task nodes in the cluster is limited to the default node range configured by you in other time periods except the time period configured in the resource plan.
- Automation scripts:
  - You can set an automation script so that it can automatically run on cluster nodes when auto scaling is triggered.
  - You can set a maximum number of 10 automation scripts for a cluster.
  - You can specify an automation script to be executed on one or more types of nodes.
  - Automation scripts can be executed before or after scale-out or scale-in.
  - Before using automation scripts, upload them to a cluster VM or OBS file system in the same region as the cluster. The automation scripts uploaded to the cluster VM can be executed only on the existing nodes. If you want to make the automation scripts run on the new nodes, upload them to the OBS file system.

2.7.2 Configuring Auto Scaling During Cluster Creation

When you create a cluster, you can configure the auto scaling function in advanced configuration parameters.

NOTE

Auto scaling policies can be configured during cluster creation only for analysis, streaming, and hybrid clusters.

Procedure

Step 1 Log in to the MRS management console.

Step 2 When you buy a cluster containing task nodes, configure the cluster software and hardware information by referring to **Buying a Custom Cluster**. Then, on the **Set Advanced Options** page, enable **Analysis Task** and configure or modify auto scaling rules and resource plans.

Figure 2-23 Configuring auto scaling rules when creating a cluster

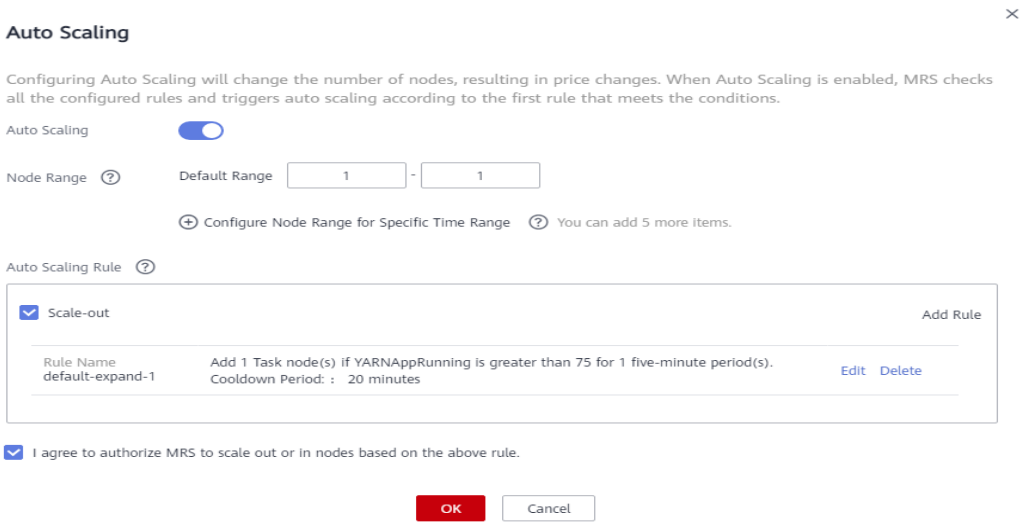

NOTE

You can configure the auto scaling rules by referring to the following scenarios:

- **Scenario 1: Using Auto Scaling Rules Alone**
- **Scenario 2: Using Resource Plans Alone**
- **Scenario 3: Using Both Auto Scaling Rules and Resource Plans**

----End

2.7.3 Creating an Auto Scaling Policy for an Existing Cluster

After a cluster is created, you can configure rules for the task node group in a cluster by node group or resource pool.

The node group policy and resource pool policy are mutually exclusive. You can configure either of them as needed.

MRS 3.1.5 or later supports the specified resource pool policy.

Item	By Node Group	By Resource Pool
Auto scaling object	All nodes in the task node group	Task nodes in the resource pool specified by an auto scaling policy
Resource pool ownership of added nodes	Default resource pool	Resource pool specified by the auto scaling policy
Scale-in object	Random scale-in of nodes in the task node group	Random scale-in of nodes in a resource pool specified by an auto scaling policy

Prerequisites

- A task node group has been configured by referring to **Adding a Task Node**.
- A resource pool has been added by referring to **Creating a Resource Pool** if you plan to configure auto scaling policies by resource pool.

Procedure

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**. In the cluster list, select a running cluster to access its details page.

Step 3 On the page that is displayed, click the **Auto Scaling** tab.

You can configure policies by resource pool or node group as needed.

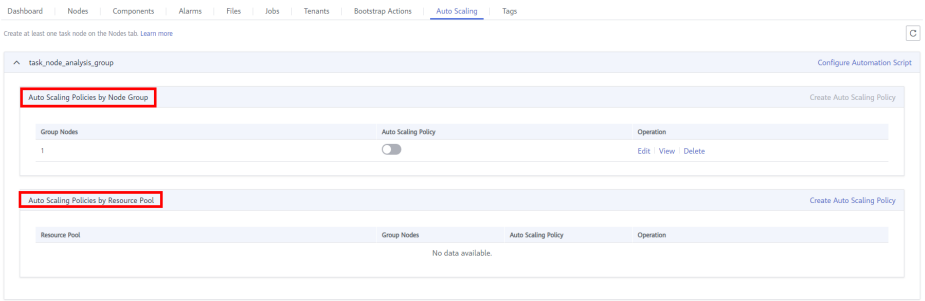

NOTE

- Auto scaling policies of different node groups are mutually exclusive. That is, you can enable auto scaling policies only for one node group.
- An auto scaling rule adjusts the number of nodes, but also affects the actual price. Exercise caution when adding an auto scaling rule.

Step 4 Click **Create Auto Scaling Policy** to create an auto scaling policy.

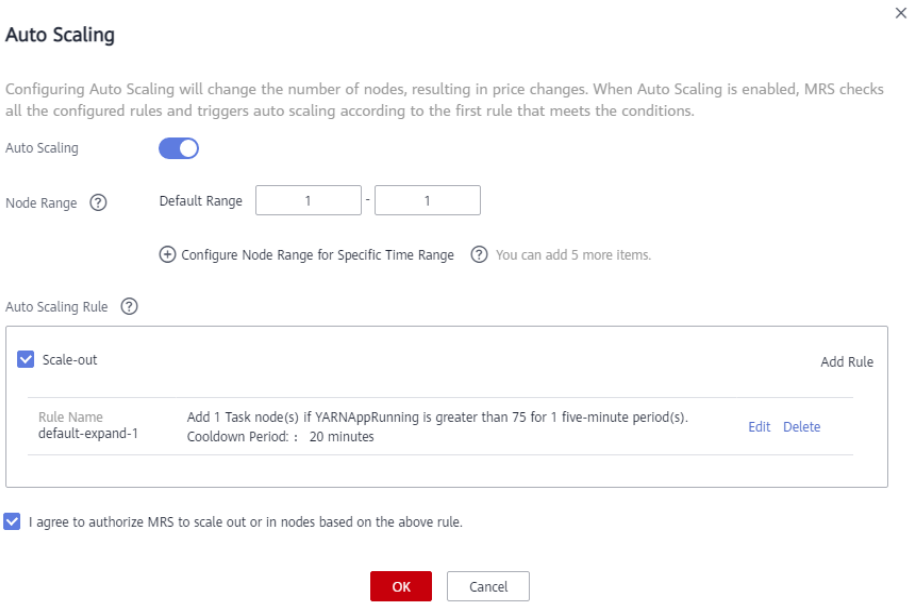

NOTE

You can configure the auto scaling rules by referring to the following scenarios:

- **Scenario 1: Using Auto Scaling Rules Alone**
- **Scenario 2: Using Resource Plans Alone**
- **Scenario 3: Using Both Auto Scaling Rules and Resource Plans**

----End

2.7.4 Scenario 1: Using Auto Scaling Rules Alone

Scenario where only auto scaling rules are configured: The number of nodes needs to be dynamically adjusted based on the YARN resource usage. When the available YARN memory is less than 20%, five nodes need to be added. When the available YARN memory is greater than 70%, five nodes need to be reduced. The number of nodes in a task node group ranges from 1 to 10.

Procedure

Step 1 Go to the **Auto Scaling** page to configure auto scaling rules.

- Configure the **Default Range** parameter.

  Enter a task node range, in which auto scaling is performed. This constraint applies to all scale-in and scale-out rules. The maximum value range allowed is 0 to 500.

  The value range in this example is 1 to 10.

- Configure an auto scaling rule.

  To enable **Auto Scaling**, you must configure a scale-out or scale-in rule.

  a. Select **Scale-Out** or **Scale-In**.

  b. Click **Add Rule**.

Figure 2-24 Add Rule dialog box

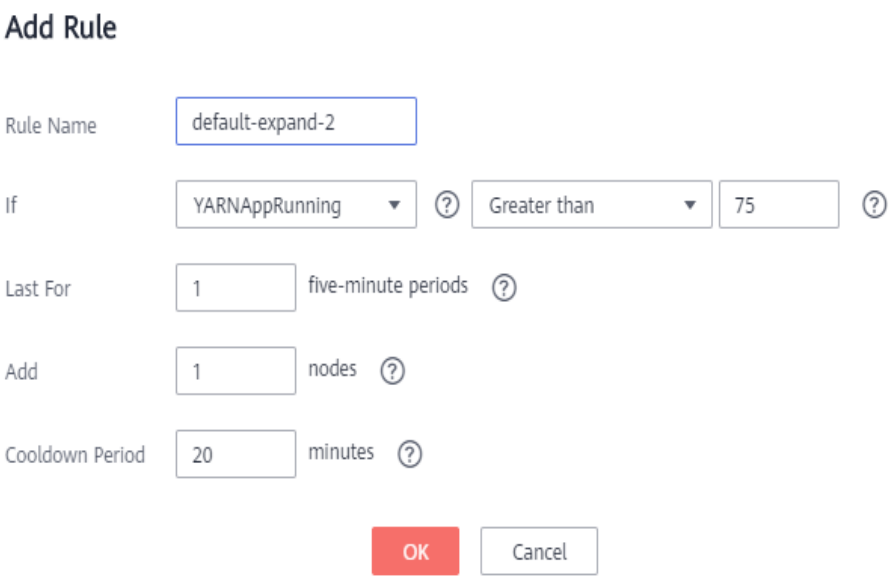

c. Configure **Rule Name**, **If**, **Last For**, **Add**, and **Cooldown Period**. For details about auto scaling metrics, see **Configuring Auto Scaling Metrics**.

d. Click **OK**.

You can view, edit, or delete the rules you configured in the **Scale-out** or **Scale-in** area on the **Auto Scaling** page. You can click **Add Rule** to configure multiple rules.

Step 2 Click **OK**.

NOTE

If you want to configure an auto scaling rule for an existing cluster, select **I agree to authorize MRS to scale out or in nodes based on the above rule**.

----End

2.7.5 Scenario 2: Using Resource Plans Alone

If the data volume changes regularly every day and you want to scale out or scale in a cluster before the data volume changes, you can create resource plans to adjust the number of Task nodes as planned in the specified time range.

Background

A real-time processing service sees a sharp increase in data volume from 7:00 to 13:00 on Monday, Tuesday, and Saturday. Assume that an MRS streaming cluster is used to process the service data. Five task nodes are required from 7:00 to 13:00 on Monday, Tuesday, and Saturday, while only two are required at other time.

Procedure

Step 1 Go to the **Auto Scaling** page to configure a resource plan.

Step 2 For example, the **Default Range** of node quantity is set to **2-2**, indicating that the number of task nodes is fixed to 2 except the time range specified in the resource plan.

Step 3 Click **Configure Node Range for Specific Time Range** under **Default Range** or **Add Resource Plan**.

Step 4 Configure **Effective On**, **Time Range**, and **Node Range**.

For example, set **Effective On** to **Monday**, **Tuesday**, and **Saturday**, **Time Range** to **07:00-13:00**, and **Node Range** to **5-5**. This indicates that the number of task nodes is fixed at 5 from 07:00 to 13:00.

You can click **Configure Node Range for Specific Time Range** to configure multiple resource plans.

NOTE

- **Effective On** is set to **Daily** by default. You can also select one or multiple days from Monday to Sunday.
- If you do not set **Node Range**, its default value will be used.
- If you set both **Node Range** and **Time Rang**e, the node range you set will be used during the time range you set, and the default node range will be used beyond the time range you set. If the time is not within the configured time range, the default range is used.

----End

2.7.6 Scenario 3: Using Both Auto Scaling Rules and Resource Plans

If the data volume is not stable and the expected fluctuation may occur, the fixed Task node range cannot guarantee that the requirements in some service scenarios are met. In this case, it is necessary to adjust the number of Task nodes based on the real-time loads and resource plans.

Background

A real-time processing service sees an unstable increase in data volume from 7:00 to 13:00 on Monday, Tuesday, and Saturday. For example, 5 to 8 task nodes are required from 7:00 to 13:00 on Monday, Tuesday, and Saturday, and 2 to 4 are required beyond this period. Therefore, you can set an auto scaling rule based on a resource plan. When the data volume exceeds the expected value, the number of Task nodes can be adjusted if resource loads change, without exceeding the node range specified in the resource plan. When a resource plan is triggered, the number of nodes is adjusted within the specified node range with minimum affect. That is, increase nodes to the upper limit and decrease nodes to the lower limit.

Procedure

Step 1 Go to the **Auto Scaling** page to configure auto scaling rules.

An auto scaling rule adjusts the number of nodes, but also affects the actual price. Exercise caution when adding an auto scaling rule.

- **Default Range**

  Enter a task node range, in which auto scaling is performed. This constraint applies to all scale-in and scale-out rules.

  For example, this parameter is set to **2-4** in this scenario.

- **Auto Scaling**

  To enable **Auto Scaling**, you must configure a scale-out or scale-in rule.

  a. Select **Scale-Out** or **Scale-In**.

  b. Click **Add Rule**. The **Add Rule** page is displayed.

  Figure 2-25 Adding a rule

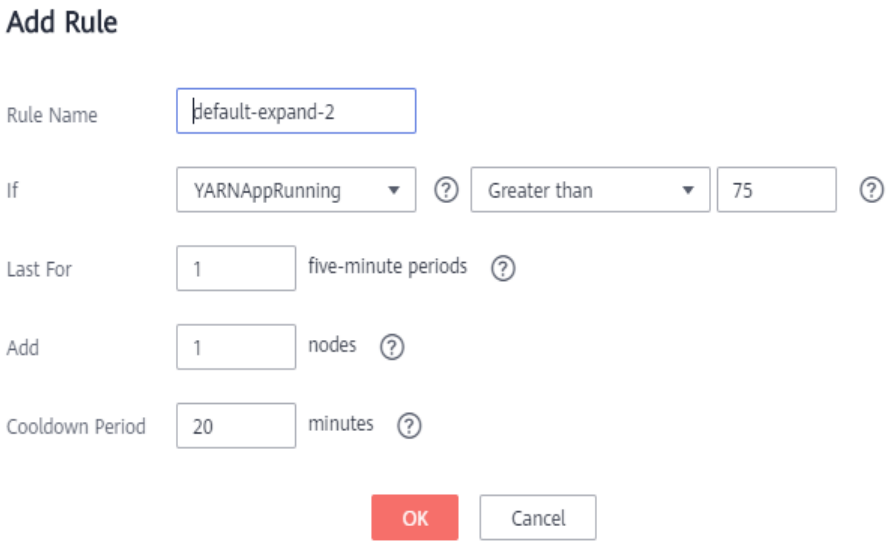

  c. Configure the **Rule Name**, **If**, **Last for**, **Add**, and **Cooldown Period** parameters.

  d. Click **OK**.

  You can view, edit, or delete the rules you configured in the **Scale-out** or **Scale-in** area on the **Auto Scaling** page.

Step 2 Configure a resource plan.

1. Click **Configure Node Range for Specific Time Range** under **Default Range** or **Add Resource Plan**.
2. Configure **Effective On**, **Time Range**, and **Node Range**.

   For example, set **Effective On** to **Monday**, **Tuesday**, and **Saturday**, **Time Range** to **07:00-13:00**, and **Node Range** to **5-8**.

   You can click **Configure Node Range for Specific Time Range** or **Add Resource Plan** to configure multiple resource plans.

NOTE

- **Effective On** is set to **Daily** by default. You can also select one or multiple days from Monday to Sunday.
- If you do not set **Node Range**, its default value will be used.
- If you set both **Node Range** and **Time Range**, the node range you set will be used during the time range you set, and the default node range will be used beyond the time range you set. If the time is not within the configured time range, the default range is used.

----End

2.7.7 Modifying an Auto Scaling Policy

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**. In the cluster list, select a running cluster to access its details page.

Step 3 Click the **Auto Scaling** tab.

Step 4 Click **Edit** on the right of the target auto scaling policy.

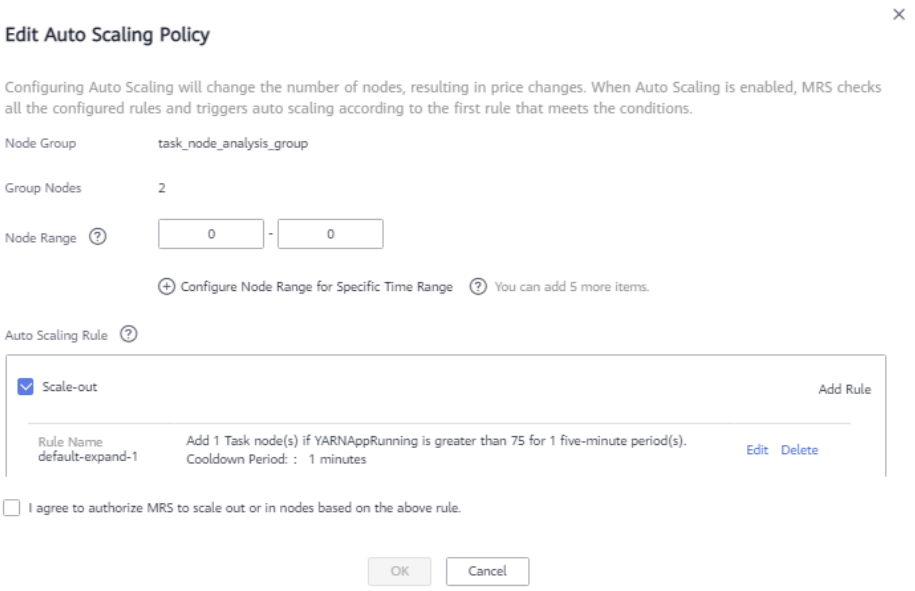

----End

2.7.8 Deleting an Auto Scaling Policy

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**. In the cluster list, select a running cluster to access its details page.

Step 3 Click the **Auto Scaling** tab.

Step 4 Click **Delete** on the right of an existing AS policy.In the displayed dialog box, click **OK**.

----End

2.7.9 Enabling or Disabling an Auto Scaling Policy

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**. In the cluster list, select a running cluster to access its details page.

Step 3 Click the **Auto Scaling** tab.

Step 4 Toggle **Auto Scaling Policy** on or off to enable or disable an auto scaling policy.

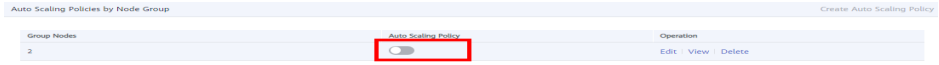

----End

2.7.10 Viewing an Auto Scaling Policy

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**. In the cluster list, select a running cluster to access its details page.

Step 3 Click the **Auto Scaling** tab.

Step 4 Click **View** on the right of the target auto scaling policy to view it.

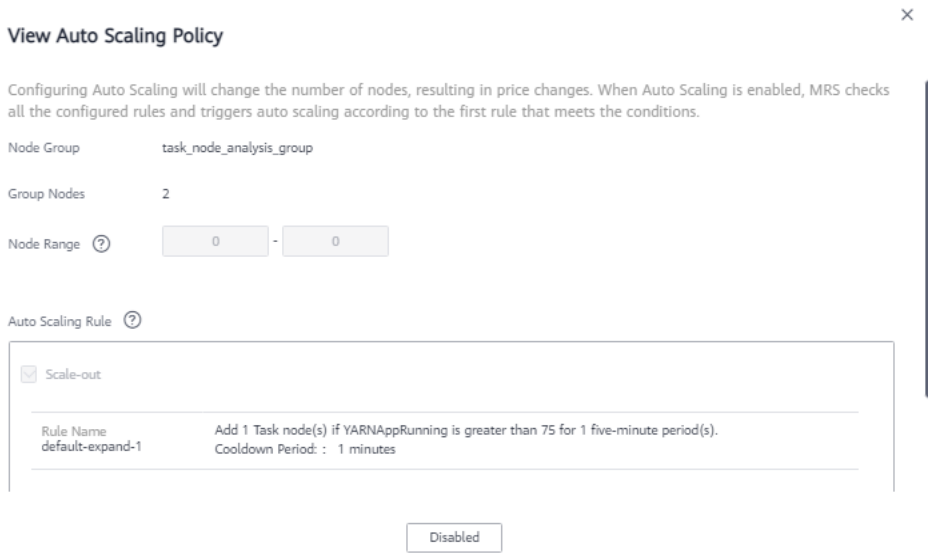

----End

2.7.11 Configuring Automation Scripts

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**. In the cluster list, select a running cluster to access its details page.

Step 3 Click the **Auto Scaling** tab.

Step 4 Click **Configure Automation Script**.

Step 5 Click **Add**.

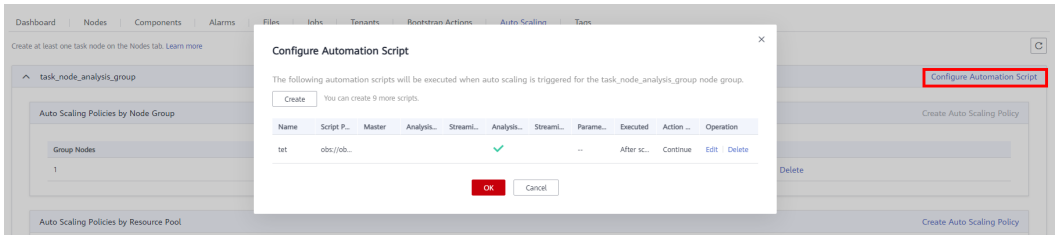

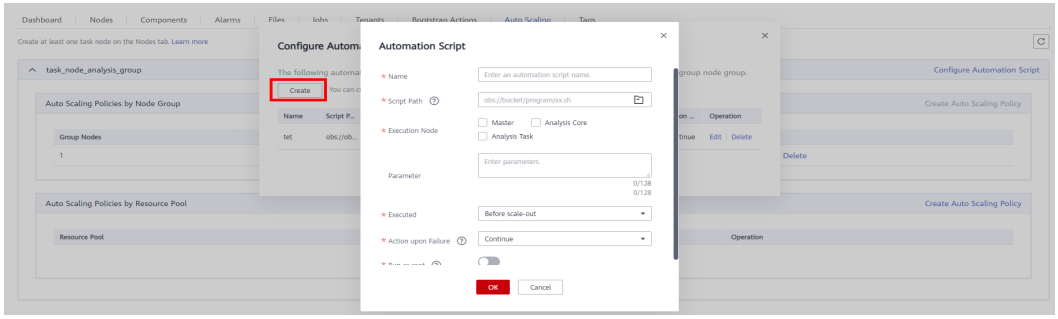

Step 6 Configure **Name**, **Script Path**, **Execution Node**, **Parameter**, **Executed**, and **Action upon Failure**. For details about the parameters, see **Table 2-14**.

Step 7 Click **OK** to save the automation script configurations.

----End

2.7.12 Configuring Auto Scaling Metrics

Auto Scaling Policies by Node Group

When you add a rule, you can refer to **Table 2-11** to configure the corresponding metrics.

Table 2-11 Auto scaling metrics

Cluster Type	Metric	Value Type	Description
Streaming cluster	StormSlotAvailable	Integer	Number of available Storm slots Value range: 0 to 2147483646
	StormSlotAvailablePercentage	Percentage	Percentage of available Storm slots, that is, the proportion of the available slots to total slots Value range: 0 to 100
	StormSlotUsed	Integer	Number of used Storm slots Value range: 0 to 2147483646

Cluster Type	Metric	Value Type	Description
	StormSlotUsedPercentage	Percentage	Percentage of the used Storm slots, that is, the proportion of the used slots to total slots Value range: 0 to 100
	StormSupervisorMemAverageUsage	Integer	Average memory usage of the Supervisor process of Storm Value range: 0 to 2147483646
	StormSupervisorMemAverageUsagePercentage	Percentage	Average percentage of the used memory of the Supervisor process of Storm to the total memory of the system Value range: 0 to 100
	StormSupervisorCPUAverageUsagePercentage	Percentage	Average percentage of the used CPUs of the Supervisor process of Storm to the total CPUs Value range: 0 to 6000
Analysis cluster	YARNAppPending	Integer	Number of pending tasks on YARN Value range: 0 to 2147483646
	YARNAppPendingRatio	Ratio	Ratio of pending tasks on YARN, that is, the ratio of pending tasks to running tasks on YARN Value range: 0 to 2147483646
	YARNAppRunning	Integer	Number of running tasks on YARN Value range: 0 to 2147483646
	YARNContainerAllocated	Integer	Number of containers allocated to YARN Value range: 0 to 2147483646
	YARNContainerPending	Integer	Number of pending containers on YARN Value range: 0 to 2147483646
	YARNContainerPendingRatio	Ratio	Ratio of pending containers on Yarn, that is, the ratio of pending containers to running containers on YARN Value range: 0 to 2147483646
	YARNCPUAllocated	Integer	Number of virtual CPUs (vCPUs) allocated to YARN Value range: 0 to 2147483646

Cluster Type	Metric	Value Type	Description
	YARNCPUAvailable	Integer	Number of available vCPUs on YARN Value range: 0 to 2147483646
	YARNCPUAvailablePercentage	Percentage	Percentage of available vCPUs on YARN that is, the proportion of available vCPUs to total vCPUs Value range: 0 to 100
	YARNCPUPending	Integer	Number of pending vCPUs on YARN Value range: 0 to 2147483646
	YARNMemoryAllocated	Integer	Memory allocated to YARN, in MB Value range: 0 to 2147483646
	YARNMemoryAvailable	Integer	Available memory on YARN in MB Value range: 0 to 2147483646
	YARNMemoryAvailablePercentage	Percentage	Percentage of available memory on YARN that is, the proportion of available memory to total memory on YARN Value range: 0 to 100
	YARNMemoryPending	Integer	Pending memory on YARN Value range: 0 to 2147483646

NOTE

- When the value type is percentage or ratio in **Table 2-11**, the valid value can be accurate to percentile. The percentage metric value is a decimal value with a percent sign (%) removed. For example, 16.80 represents 16.80%.
- Hybrid clusters support all metrics of analysis and streaming clusters.

Auto Scaling Policies by Resource Pool

When adding a rule, you can refer to **Table 2-12** to configure the corresponding metrics.

NOTE

Auto scaling policies can be configured for a cluster by resource pool in MRS 3.1.5 or later.

Table 2-12 Rule configuration description

Cluster Type	Metric	Value Type	Description
Analysis/Custom cluster	ResourcePoolMemoryAvailable	Integer	Available memory on YARN in the resource pool, in MB Value range: 0 to 2147483646
	ResourcePoolMemoryAvailablePer-centage	Percentage	Percentage of available memory on YARN in the resource pool, that is, the proportion of available memory to total memory on YARN Value range: 0 to 100
	ResourcePoolCPUAvailable	Integer	Number of available vCPUs on YARN in the resource pool Value range: 0 to 2147483646
	ResourcePoolCPUAvailablePercent-age	Percentage	Percentage of available vCPUs on YARN in the resource pool. that is, the proportion of available vCPUs to total vCPUs Value range: 0 to 100

When you add a resource plan, you can configure parameters by referring to **Table 2-13**.

Table 2-13 Resource plan configuration items

Parameter	Description
Effective On	The effective date of a resource plan. **Daily** is selected by default. You can also select one or multiple days from Monday to Sunday.

Parameter	Description
Time Range	Start time and end time of a resource plan are accurate to minutes, with the value ranging from **00:00** to **23:59**. For example, if a resource plan starts at 8:00 and ends at 10:00, set this parameter to **8:00-10:00**. The end time must be at least 30 minutes later than the start time.
Node Range	The number of nodes in a resource plan ranges from **0** to **500**. In the time range specified in the resource plan, if the number of task nodes is less than the specified minimum number of nodes, it will be increased to the specified minimum value of the node range at a time. If the number of task nodes is greater than the maximum number of nodes specified in the resource plan, the auto scaling function reduces the number of task nodes to the maximum value of the node range at a time. The minimum number of nodes must be less than or equal to the maximum number of nodes.

NOTE

- When a resource plan is enabled, the **Default Range** value on the auto scaling page forcibly takes effect beyond the time range specified in the resource plan. For example, if **Default Range** is set to **1-2**, **Time Range** is between **08:00-10:00**, and **Node Range** is **4-5** in a resource plan, the number of Task nodes in other periods (0:00-8:00 and 10:00-23:59) of a day is forcibly limited to the default node range (1 to 2). If the number of nodes is greater than 2, auto scale-in is triggered; if the number of nodes is less than 1, auto scale-out is triggered.
- When a resource plan is not enabled, the **Default Range** takes effect in all time ranges. If the number of nodes is not within the default node range, the number of Task nodes is automatically increased or decreased to the default node range.
- Time ranges of resource plans cannot be overlapped. The overlapped time range indicates that two effective resource plans exist at a time point. For example, if resource plan 1 takes effect from **08:00** to **10:00** and resource plan 2 takes effect from **09:00** to **11:00**, the time range between **09:00** to **10:00** is overlapped.
- The time range of a resource plan must be on the same day. For example, if you want to configure a resource plan from **23:00** to **01:00** in the next day, configure two resource plans whose time ranges are **23:00-00:00** and **00:00-01:00**, respectively.

Automation Script

When you add an automation script, you can configure related parameters by referring to **Table 2-14**.

Table 2-14 Automation script configuration description

Parameter	Description
Name	Name of an automation script The value can contain only numbers, letters, spaces, hyphens (-), and underscores (_) and must not start with a space. The value can contain 1 to 64 characters. **NOTE** A name must be unique in the same cluster. You can configure the same name for different clusters.
Script Path	Script path. The value can be an OBS file system path or a local VM path. ● An OBS file system path must start with **obs://** and end with **.sh**, for example, **obs://mrs-samples/*xxx*.sh**. ● A local VM path must start with a slash (/) and end with **.sh**. For example, the path of the example script for installing the Zepelin is **/opt/bootstrap/zepelin/zepelin_install.sh**.
Execution Node	Select a type of the node where an automation script is executed. **NOTE** ● If you select **Master** nodes, you can choose whether to run the script only on the active Master nodes by enabling or disabling the **Active Master** switch. ● If you enable it, the script runs only on the active Master nodes. If you disable it, the script runs on all master nodes. This function is disabled by default.
Parameter	Automation script parameter. The following predefined variables can be imported to obtain auto scaling information: ● **${mrs_scale_node_num}**: Number of auto scaling nodes. The value is always positive. ● **${mrs_scale_type}**: Scale-out/in type. The value can be **scale_out** or **scale_in**. ● **${mrs_scale_node_hostnames}**: Host names of the auto scaling nodes. Use commas (,) to separate multiple host names. ● **${mrs_scale_node_ips}**: IP address of the auto scaling nodes. Use commas (,) to separate multiple IP addresses. ● **${mrs_scale_rule_name}**: Name of the triggered auto scaling rule. For a resource plan, this parameter is set to **resource_plan**.

Parameter	Description
Executed	Time for executing an automation script. The following four options are supported: **Before scale-out**, **After scale-out**, **Before scale-in**, and **After scale-in**. **NOTE** Assume that the execution nodes include Task nodes. • The automation script executed before scale-out cannot run on the Task nodes to be added. • The automation script executed after scale-out can run on the added Task nodes. • The automation script executed before scale-in can run on Task nodes to be deleted. • The automation script executed after scale-in cannot run on the deleted Task nodes.
Action upon Failure	Whether to continue to execute subsequent scripts and scale-out/in after the script fails to be executed. **NOTE** • You are advised to set this parameter to **Continue** in the commissioning phase so that the cluster can continue the scale-out/in operation no matter whether the script is executed. • If the script fails to be executed, view the log in **/var/log/Bootstrap** on the cluster VM. • The scale-in operation cannot be rolled back. Therefore, the **Action upon Failure** can only be set to **Continue** after scale-in.

NOTE

The automation script is triggered only during auto scaling. It is not triggered when the cluster node is manually scaled out or in.

2.8 Managing Data Connections

2.8.1 Configuring Data Connections

MRS data connections are used to manage external source connections used by components in a cluster. For example, if Hive metadata uses an external relational database, a data connection can be used to associate the external relational database with the Hive component.

- **Local**: Metadata is stored in the local GaussDB of a cluster. When the cluster is deleted, the metadata is also deleted. To retain the metadata, manually back up the metadata in the database in advance.
- **External data connection**: After the cluster is created, you can select **RDS PostgreSQL database** or **RDS MySQL database** that is associated with the same VPC and subnet as the current cluster. Metadata is stored in the associated database and is not deleted when the current cluster is deleted. Multiple MRS clusters can share the same metadata.

NOTE

When Hive metadata is switched between different clusters, MRS synchronizes only the permissions in the metadata database of the Hive component. The permission model on MRS is maintained on MRS Manager. Therefore, when Hive metadata is switched between clusters, the permissions of users or user groups cannot be automatically synchronized to MRS Manager of another cluster.

Creating a Data Connection

Step 1 Log in to the MRS management console, and choose **Data Connections** in the left navigation pane.

Step 2 Click **Create Data Connection**.

For details about how to configure an RDS data connection, see **Creating an RDS Data Connection**.

NOTE

- **RDS PostgreSQL database**. Clusters that support Hive can connect to this type of database.
- **RDS MySQL database**. Clusters that support Hive or Ranger can connect to this type of database.
- Currently, MRS supports **PostgreSQL 14** on RDS.
- Currently, MRS supports only **MySQL 5.7.x/MySQL 8.0***x* on RDS.

Step 3 Click **OK**.

----End

Viewing Data Connection Details

Step 1 Log in to the MRS management console, and choose **Data Connections** in the left navigation pane.

Step 2 In the data connection list, click the desired data connection. On the page that is displayed, view its details.

For example, the data connection information of the RDS MySQL database is as follows:

Figure 2-26 Viewing the data connection information of the RDS MySQL database

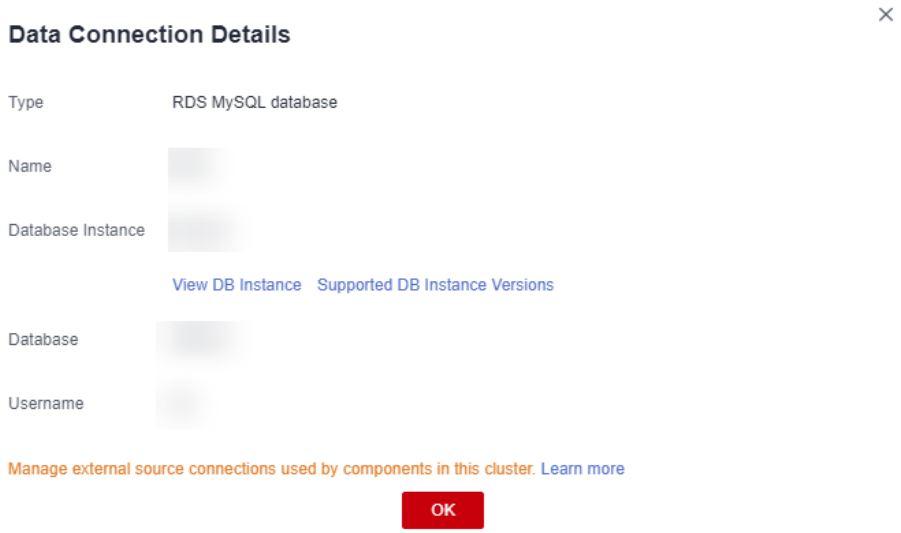

----End

Deleting a Data Connection

Step 1 Log in to the MRS management console, and choose **Data Connections** in the left navigation pane.

Step 2 In the **Operation** column of the data connection list, click **Delete** in the row where the data connection to be deleted is located. Enter **DELETE** in the **Delete Data Connection** dialog box and click **OK**.

If the selected data connection has been associated with a cluster, the deletion does not affect the cluster.

----End

Configuring a Data Connection During Cluster Creation

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

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

Step 3 When you create a cluster, **Data Connection Type** can only be set to **Local database**. For details about how to configure other parameters, see **Buying a Custom Cluster**.

----End

2.8.2 Configuring an RDS Data Connection

2.8.2.1 Configuring an RDS Data Connection

Preparations

Step 1 Log in to the RDS management console.

Step 2 Buy an RDS DB instance.

Step 3 In the left navigation pane of the RDS management console, choose **Instances**. Locate the row containing the RDS DB instance used by MRS data connections, click **More** in the **Operation** column, and select **Log In** to log in to the DB instance as user **root**.

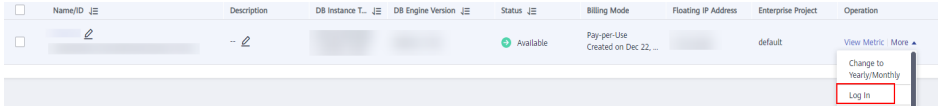

Step 4 On the home page of the instance, click **Create Database** to create a database.

NOTE

If no new database is created, the MRS data connections will fail to configure.

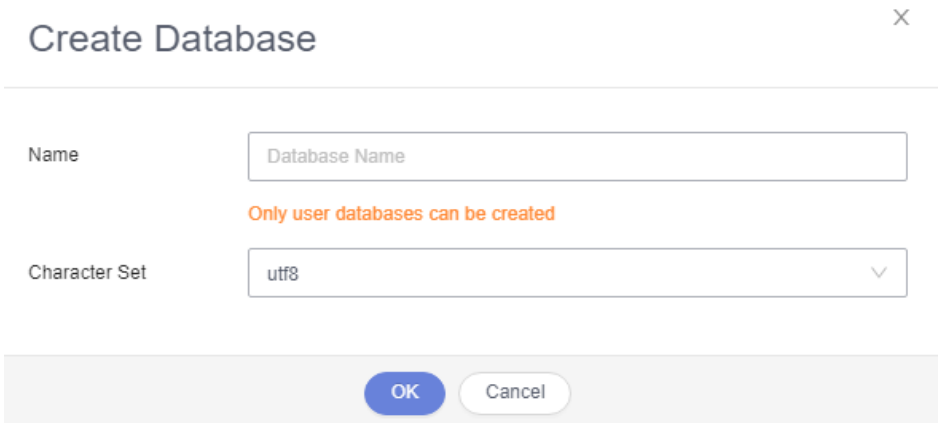

Step 5 On the top of the page, choose **Account Management** > **User Management**.

NOTE

If the selected data connection is **RDS MySQL database**, ensure that the database user is user **root**. If the user is not **root**, perform **Step 5** to **Step 7**.

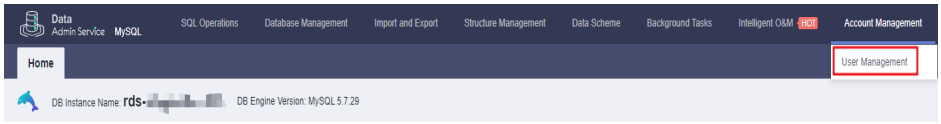

Step 6 Click **Create User** to create a non-root user.

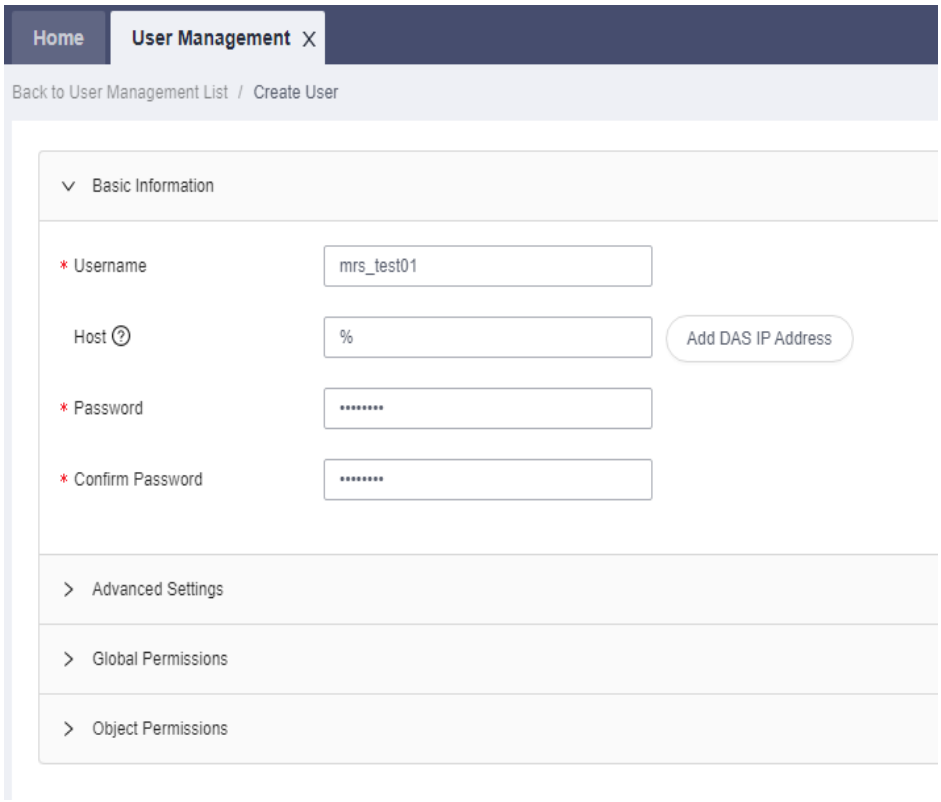

Step 7 On the top of the page, choose **SQL Operations** > **SQL Query**, switch to the target database by database name, and run the following SQL statements to grant permissions to the database user. In the following statements, *${db_name}* and *${db_user}* indicate the name of the database to be connected to MRS and the name of the new user, respectively.

```
grant SELECT, INSERT on mysql.* to '${db_user}'@'%' with grant option;
grant all privileges on ${db_name}.* to '${db_user}'@'%' with grant option;
grant reload on *.* to '${db_user}'@'%' with grant option;
flush privileges;
```

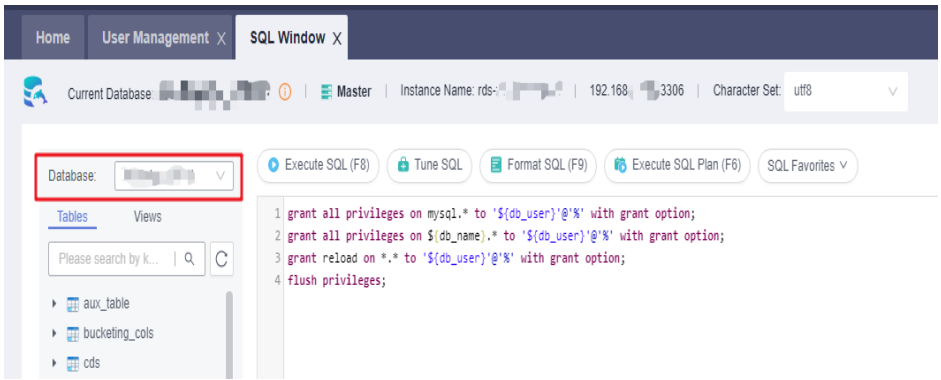

Step 8 Create a data connection by referring to **Creating an RDS Data Connection**.

----End

Creating an RDS Data Connection

Create an RDS data connection for an existing MRS cluster.

Step 1 Log in to the MRS management console, and choose **Data Connections** in the left navigation pane.

Step 2 Click **Create Data Connection**.

Step 3 Configure parameters according to **Table 2-15**.

Table 2-15 Parameters for creating a data connection

Parameter	Description
Type	The type of an external source connection. Value options are as follows: • **RDS PostgreSQL database**. Clusters that support Hive can connect to this type of database. • **RDS MySQL database**. Clusters that support Hive or Ranger can connect to this type of database.
Name	The name of a data connection.
Database Instance	The RDS database instance. This instance must be created in RDS before being referenced here, and the database must have been created. For details, see **Preparations**. Click **View DB Instance** to view the created DB instances. **NOTE** • To ensure network communications between the cluster and the PostgreSQL database, create the instance in the same VPC and subnet as the cluster. • The inbound rule of the security group of the RDS DB instance must allow access of the instance to port 3306. To configure that, click the instance name on the RDS console to go to the instance management page. In the **Connection Information** area, click the name next to **Security Group**. On the page that is displayed, click the **Inbound Rules** tab, and click **Add Rule**. In the displayed **Add Inbound Rule** dialog box, in the **Protocol & Port** area, select **TCP** and enter port number **3306**. In the **Source** area, select **IP address** and enter the IP addresses of all nodes where the MetaStore instances of Hive are located. • Currently, MRS supports **PostgreSQL 14** on RDS. • Currently, MRS supports only **MySQL 5.7.x/MySQL 8.0***x* on RDS.
Database	The name of the database to be connected to.
Username	The username for logging in to the database to be connected.
Password	The password for logging in to the database to be connected.

Figure 2-27 Parameters for creating an RDS database connection

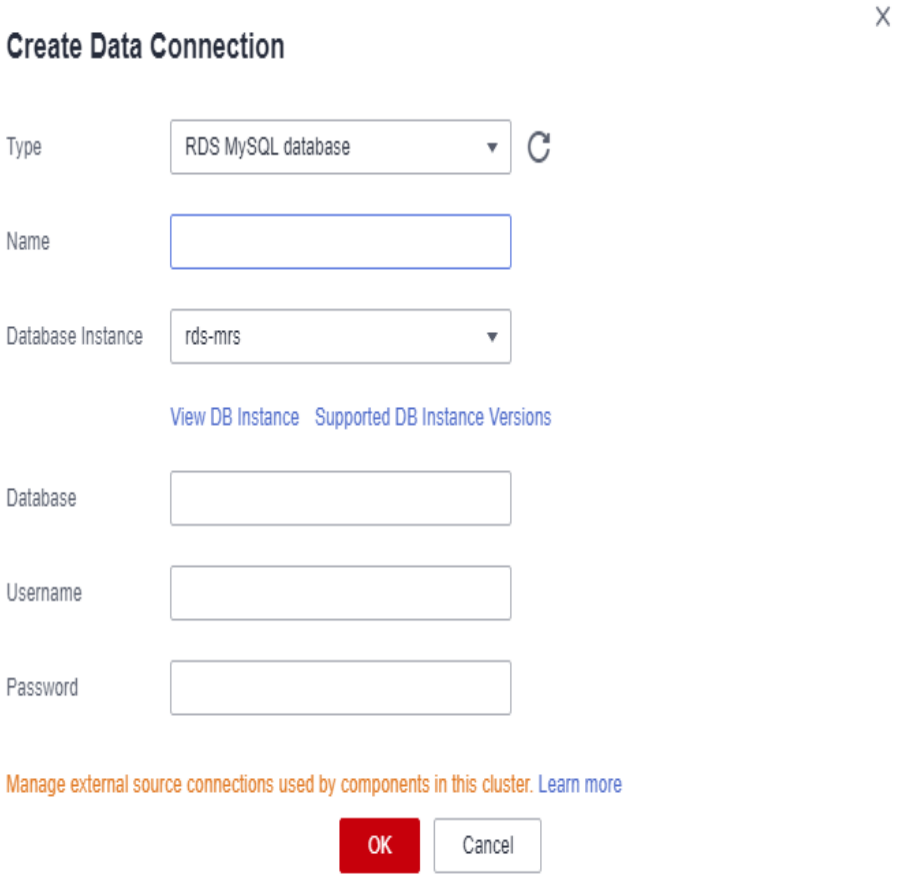

NOTE

When **Type** is set to **RDS MySQL database** or **GaussDB(for MySQL)**, **Username** must be **root**. If the user is not **root**, perform operations by referring to **Preparations**.

Step 4 Click **OK**.

----End

2.8.2.2 Configuring a Ranger Data Connection

Switch the Ranger metadata of the existing cluster to the metadata stored in the RDS database. This operation enables multiple MRS clusters to share the same metadata, and the metadata will not be deleted when the clusters are deleted. In this way, Ranger metadata migration is not required during cluster migration.

Prerequisites

You have created an RDS MySQL database instance by referring to **Creating an RDS Data Connection**.

NOTE

- For versions earlier than MRS 3.*x*, if **Type** is set to **RDS MySQL database**, **Username** must be **root**. If the user is not **root**, create a user and grant permissions to the user by referring to **Preparations**.
- For MRS 3.*x* or later clusters, when **Type** is set to **RDS MySQL database**, **Username** must not be **root**. In this case, create a user and grant permissions to the user by referring to **Preparations**.

Preparing for MySQL Database Ranger Metadata Configuration

This operation is required only for **MRS 3.1.0 or later**.

Step 1 Log in to FusionInsight Manager. For details, see **Accessing FusionInsight Manager (MRS 3.x or Later)**. Choose **Clusters** > **Services** > *Service name*.

Currently, the following components in an MRS 3.1.*x* cluster support Ranger authentication: HDFS, HBase, Hive, Spark, Impala, Storm, and Kafka.

Step 2 In the upper right corner of the **Dashboard** page, click **More** and select **Disable Ranger**. If **Disable Ranger** is dimmed, Ranger authentication is disabled, as shown in **Figure 2-28**.

Figure 2-28 Disabling Ranger authentication

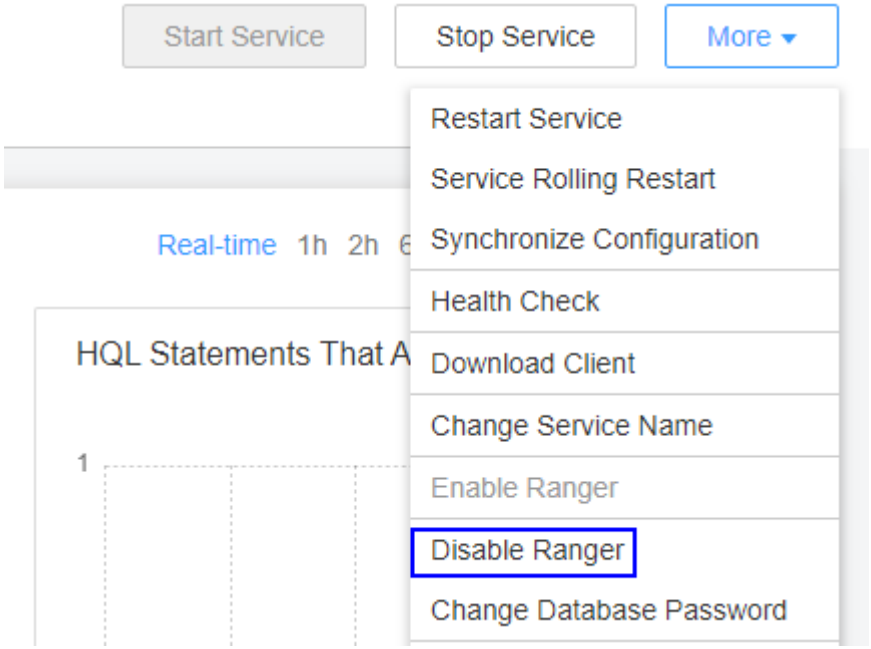

Step 3 (Optional) To use an existing authentication policy, perform this step to export the authentication policy on the Ranger web page. After the Ranger metadata is switched, you can import the existing authentication policy again. The following uses Hive as an example. After the export, a policy file in JSON format is generated in a local directory.

1. Log in to FusionInsight Manager.
2. Choose **Cluster** > **Services** > **Ranger** to go to the Ranger service overview page.

3. Click **RangerAdmin** in the **Basic Information** area to go to the Ranger web UI.

   The **admin** user in Ranger belongs to the **User** type. To view all management pages, click the username in the upper right corner and select **Log Out** to log out of the system.

4. Log in to the system as user **rangeradmin** (default password: **Rangeradmin@123**) or another user who has the Ranger administrator permissions. User Account List

5. Click the export button  in the row where the Hive component is located to export the authentication policy.

   Figure 2-29 Exporting authentication policies

   

6. Click **Export**. After the export is complete, a policy file in JSON format is generated in a local directory.

   Figure 2-30 Exporting Hive authentication policies

   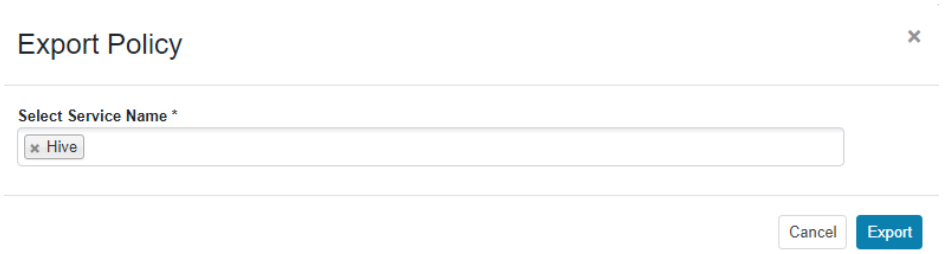

----End

Configuring a Data Connection for an MRS Cluster

Step 1 Log in to the MRS console.

Step 2 Click the name of the cluster to view its details.

Step 3 Click **Manage** on the right of **Data Connection** to go to the data connection configuration page.

Step 4 Click **Configure Data Connection** and set related parameters.

- **Component Name**: Ranger
- **Module Type**: Ranger metadata
- **Connection Type**: RDS MySQL database
- **Connection Instance**: Select a created RDS MySQL DB instance. For details about how to create a data connection, see **Creating an RDS Data Connection**.

Step 5 Select **I understand the consequences of performing the scale-in operation** and click **Test**.

Step 6 After the test is successful, click **OK** to complete the data connection configuration.

Step 7 Log in to FusionInsight Manager.

Step 8 Choose **Cluster** > **Services** > **Ranger** to go to the Ranger service overview page.

Step 9 Choose **More** > **Restart Service** or **More** > **Service Rolling Restart**.

If you choose **Restart Service**, services will be interrupted during the restart. If you select **Service Rolling Restart**, rolling restart can minimize the impact or do not affect service running.

Restarting Ranger will affect the permissions of all components controlled by Ranger and may affect service running. Restart Ranger when the cluster is idle or during off-peak hours. Before the Ranger component is restarted, the policies in the Ranger component still take effect.

Figure 2-31 Restarting a service

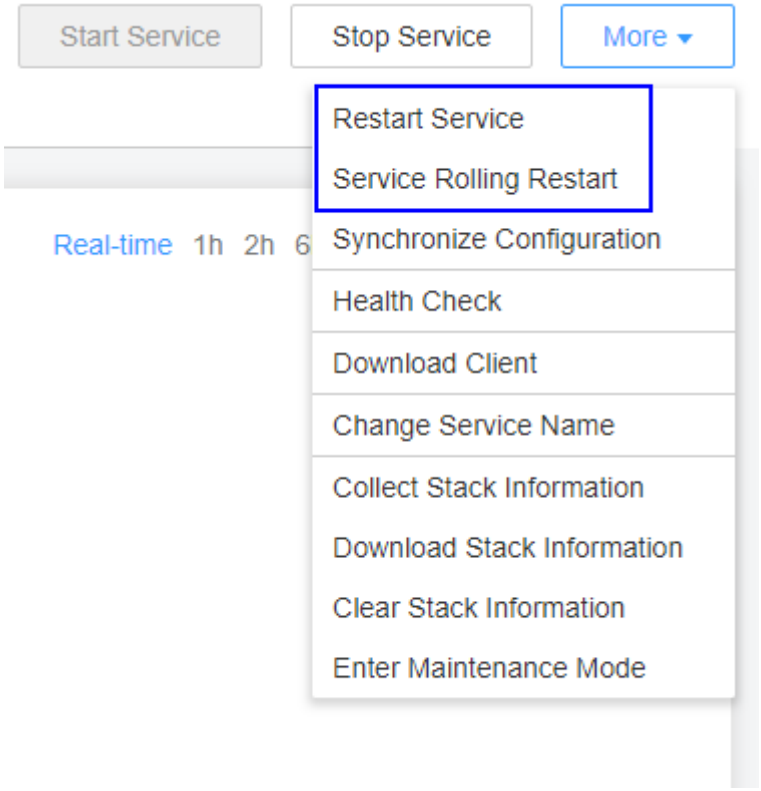

Step 10 Enable Ranger authentication for the component to be authenticated. The Hive component is used as an example.

Currently, the following components in an MRS 3.1.*x* cluster support Ranger authentication: HDFS, HBase, Hive, Spark, Impala, Storm, and Kafka.

1. Log in to FusionInsight Manager and choose **Cluster** > **Services** > *Service Name*.
2. In the upper right corner of the **Dashboard** page, click **More** and select **Enable Ranger**.

Figure 2-32 Enabling Ranger authentication

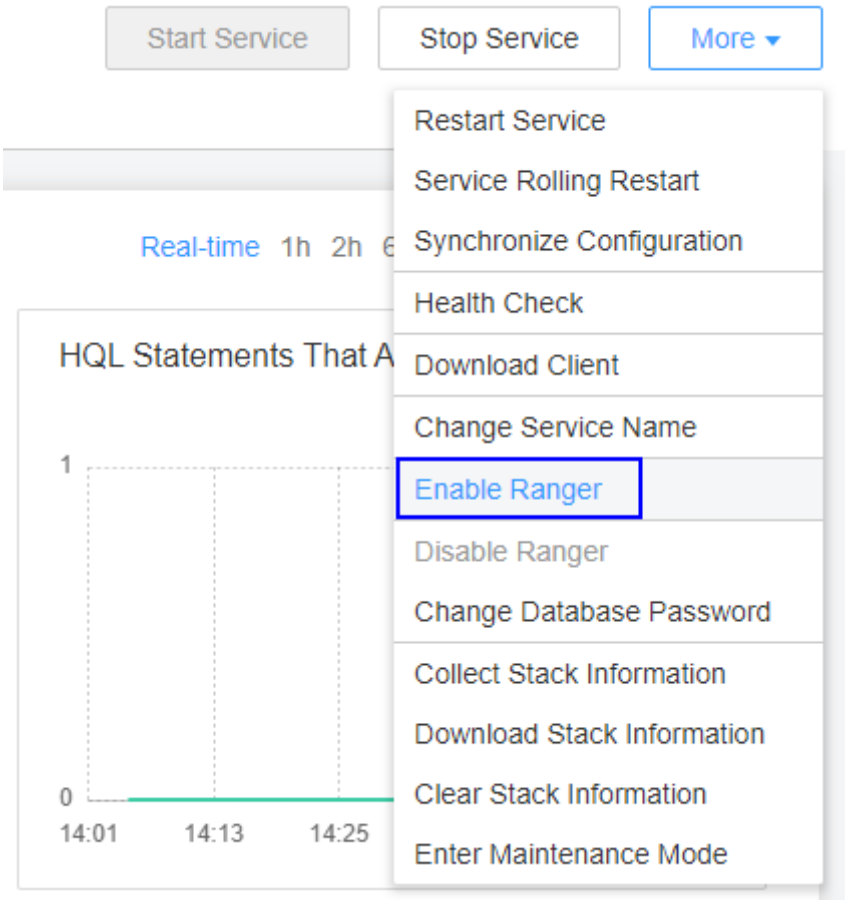

Step 11 Log in to the Ranger web UI and click the import button in the row of the Hive component.

Step 12 Import parameters.

- Click **Select file** and select the authentication policy file downloaded in **Step 3.6**.
- Select **Merge If Exist Policy**.

Figure 2-33 Importing authentication policies

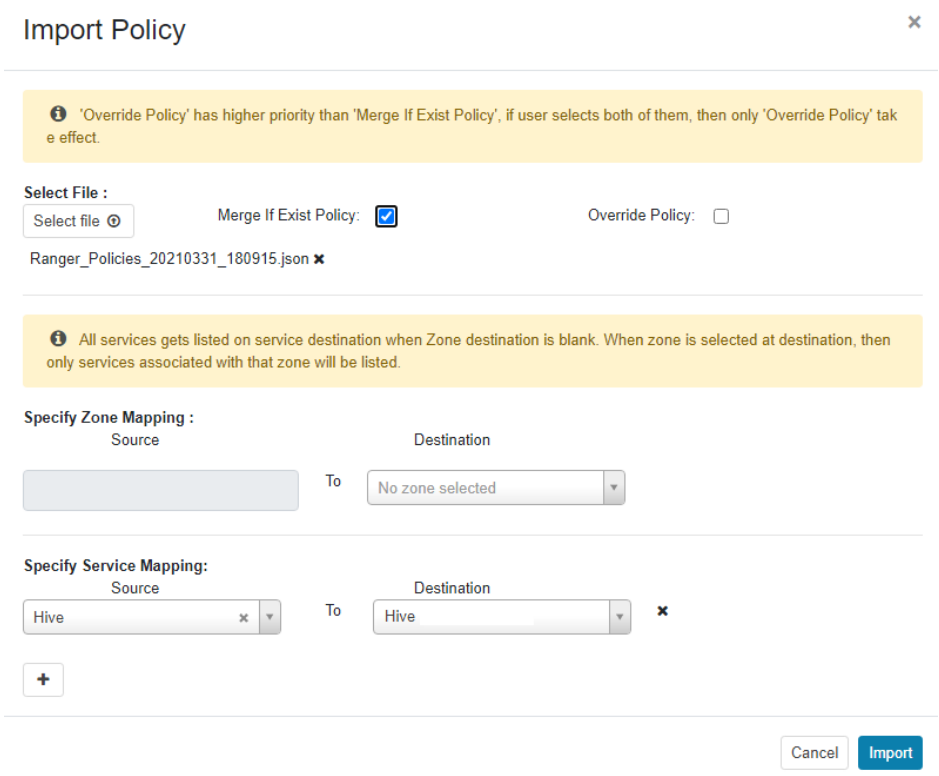

Step 13 Restart the component for which Ranger authentication is enabled.

1. Log in to FusionInsight Manager.
2. Choose **Cluster** > **Services** > **Hive** to go to the Hive service overview page.
3. Choose **More** > **Restart Service** or **More** > **Service Rolling Restart**.

Figure 2-34 Restarting a service

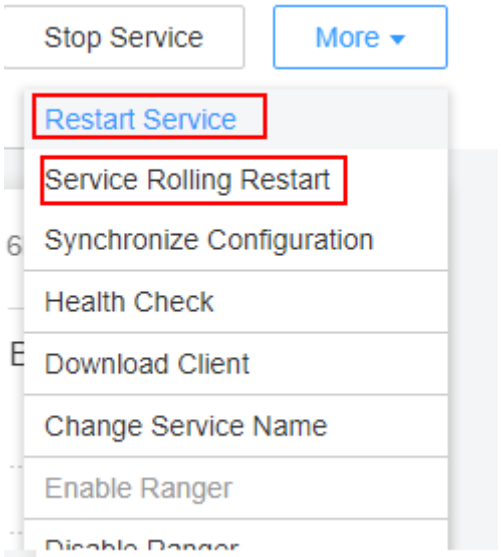

If you choose **Restart Service**, services will be interrupted during the restart. If you select **Service Rolling Restart**, rolling restart can minimize the impact or do not affect service running.

----End

2.8.2.3 Configuring a Hive Data Connection

This section describes how to switch the Hive metadata of an active cluster to the metadata stored in a local database or RDS database after you buy a cluster. This operation enables multiple MRS clusters to share the same metadata, and the metadata will not be deleted when the clusters are deleted. In this way, Hive metadata migration is not required during cluster migration.

NOTE

- When Hive metadata is switched between different clusters, MRS synchronizes only the permissions in the metadata database of the Hive component. The permission model on MRS is maintained on MRS Manager. Therefore, when Hive metadata is switched between clusters, the permissions of users or user groups cannot be automatically synchronized to MRS Manager of another cluster.
- For clusters whose version is earlier than MRS 3.*x*, if the selected data connection is **RDS MySQL database**, ensure that the database user is **root**. If the user is not **root**, create a user and grant permissions to the user by referring to **Preparations**.
- For MRS 3.*x* or later clusters, when **Type** is set to **RDS MySQL database**, **Username** must not be **root**. In this case, create a user and grant permissions to the user by referring to **Preparations**.

Configuring a Hive Data Connection

This function is not supported in MRS 3.0.5.

Step 1 Log in to the MRS console. In the navigation pane on the left, choose **Clusters** > **Active Clusters**.

Step 2 Click the name of a cluster to go to the cluster details page.

Step 3 On the **Dashboard** tab page, click **Manage** next to **Data Connection**.

Step 4 On the **Data Connection** dialog box, the data connections associated with the cluster are displayed. You can click **Edit** or **Delete** to edit or delete the data connections.

Step 5 If there is no associated data connection on the **Data Connection** dialog box, click **Configure Data Connection** to add a connection.

NOTE

Only one data connection can be configured for a module type. For example, after a data connection is configured for Hive metadata, no other data connection can be configured for it. If no module type is available, the **Configure Data Connection** button is unavailable.

Table 2-16 Configuring a Hive data connection

Parameter	Description
Component	Hive

Parameter	Description
Module Type	Hive metadata
Data Connection Type	• RDS PostgreSQL database (supported for clusters of MRS 1.9.*x*) • RDS MySQL database • Local database
Instance	This parameter is valid only when **Data Connection Type** is set to **RDS PostgreSQL database** or **RDS MySQL database**. Select the name of the connection between the MRS cluster and the RDS database. This instance must be created before being referenced here. You can click **Create Data Connection** to create a data connection. For details, see **Creating an RDS Data Connection**.

Figure 2-35 Configuring a data connection

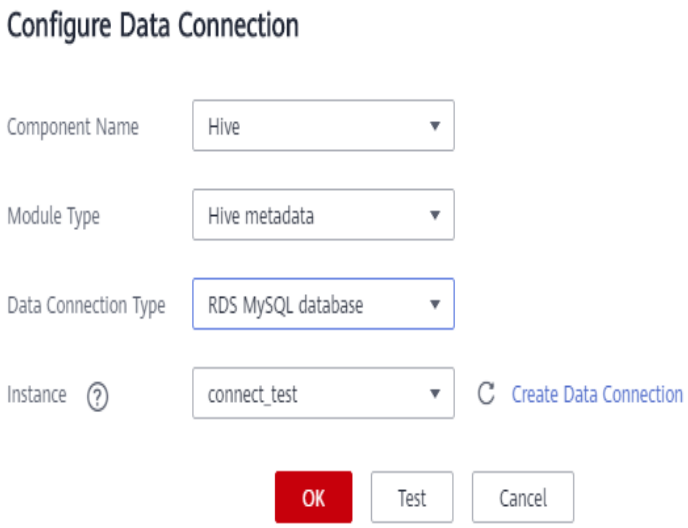

Step 6 Click **Test** to test connectivity of the data connection.

Step 7 After the data connection is successful, click **OK**.

NOTE

- After Hive metadata is configured, restart Hive. Hive will create necessary database tables in the specified database. (If tables already exist, they will not be created.)
- Before restarting the Hive service, ensure that the driver package has been installed on all nodes where Metastore instances are located.
  - Postgres: Use the open source Postgres driver package to replace the existing one of the cluster. Upload the PostgreSQL driver package **postgresql-42.2.5.jar** to the *${BIGDATA_HOME}*/**third_lib/Hive** directory on all MetaStore nodes.
  - MySQL: Go to the MySQL official website (https://www.mysql.com/). Choose **DOWNLOADS** and click **MySQL Community (GPL) Downloads**. On the displayed page, click **Connector/J** to download the driver package of the corresponding version and upload the driver package to the **/opt/Bigdata/FusionInsight_HD_*/install/FusionInsight-Hive-*/hive-*/lib/** directory on all RDSMetastore nodes.

----End

2.9 Installing Third-Party Software Using Bootstrap Actions

Prerequisites

The bootstrap action script has been prepared by referring to **Preparing the Bootstrap Action Script**.

Adding a Bootstrap Action When Creating a Cluster

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

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

Step 3 Configure the cluster software and hardware by referring to **Buying a Custom Cluster**.

Step 4 In the **Set Advanced Options** area, select **Configure** and click **Add** in the **Bootstrap Action** area.

Table 2-17 Parameters

Parameter	Description
Name	Name of a bootstrap action script The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and must not start with a space. The value can contain 1 to 64 characters. **NOTE** A name must be unique in the same cluster. You can set the same name for different clusters.

Parameter	Description
Script Path	Script path. The value can be an OBS file system path or a local VM path. • An OBS file system path must start with **obs://** and end with **.sh**, for example, **obs://mrs-samples/***xxx***.sh**. • A local VM path must start with a slash (/) and end with **.sh**. **NOTE** A path must be unique in the same cluster, but can be the same for different clusters.
Parameter	Bootstrap action script parameters
Execution Node	Select a type of the node where the bootstrap action script is executed.
Executed	Select the time when the bootstrap action script is executed. • Before initial component start • After initial component start **NOTE** You can only manually run the third-party component installation script on the node to install a running cluster component.
Action upon Failure	Whether to continue to execute subsequent scripts and create a cluster after the script fails to be executed. **NOTE** You are advised to set this parameter to **Continue** in the debugging phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.
Run as root	Whether to escalate the permission to user **root** If the bootstrap action requires root user operations, enable this function, or the bootstrap action may fail to execute. **NOTE** This parameter is available for MRS 3.1.5 clusters.

Step 5 Click **OK**.

After the bootstrap action is added, you can edit, clone, or delete it in the **Operation** column.

----End

Adding an Automation Script on the Auto Scaling Page

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**. In the cluster list, select a running cluster to access its details page.

Step 3 On the page that is displayed, click the **Auto Scaling** tab.

You can configure policies by resource pool or node group as needed. For details, see **Creating an Auto Scaling Policy for an Existing Cluster**

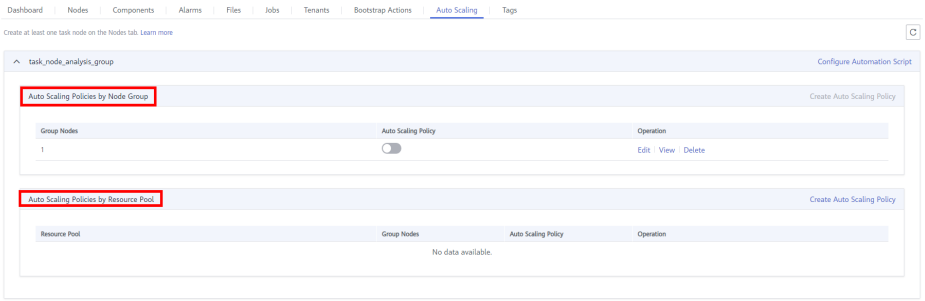

NOTE

- Auto scaling policies of different node groups are mutually exclusive. That is, you can enable auto scaling policies only for one node group.
- An auto scaling rule adjusts the number of nodes, but also affects the actual price. Exercise caution when adding an auto scaling rule.

Step 4 (Optional) Configure automation scripts.

1. Click **Configure Automation Script**.
2. Click **Add**.
3. Configure **Name**, **Script Path**, **Execution Node**, **Parameter**, **Executed**, and **Action upon Failure**. For details about the parameters, see **Table 2-14**.
4. Click **OK** to save the automation script configurations.

----End

2.10 Viewing Failed MRS Tasks

This section describes how to view and delete a failed MRS task.

Background

If a cluster fails to be created, deleted, scaled out, or scaled in, the **Manage Failed Tasks** page is displayed. Only the tasks that fail to be deleted are displayed on the **Cluster History** page. You can delete a failed task that is not required.

Procedure

Step 1 Log in to the MRS console.

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

Step 3 Click or the number on the right of **Failed Tasks**. The **Manage Failed Tasks** page is displayed.

Step 4 In the **Operation** column of the cluster that you want to start, click **Delete**.

In this step, only one job can be deleted.

Step 5 You can click **Delete All** in the upper left corner of the task list to delete all failed tasks.

----End

2.11 Viewing Information of a Historical Cluster

Choose **Clusters > Cluster History** and click the name of a target cluster. You can view the cluster configuration information, nodes, auto scaling information, component information, job information, bootstrap action, and tags.

The following table describes the parameters for the historical cluster information.

Table 2-18 Basic cluster information

Parameter	Description
Cluster Name	Name of a cluster. The cluster name is set when the cluster is created.
Cluster Status	Status of a cluster.
Billing Mode	Billing mode of a cluster. Currently, **Pay-per-use** and **Yearly/Monthly** are supported.
Cluster Version	Cluster version
Cluster Type	Type of the cluster to be created.
Obtaining a cluster ID	Unique identifier of a cluster, which is automatically assigned when a cluster is created
Created	Time when a cluster is created.
Order ID	Order ID for creating the cluster. This parameter is available only when **Billing Mode** is set to **Yearly/Monthly**.
AZ	Availability zone (AZ) in the region of a cluster, which is set when a cluster is created.
Default Subnet	Subnet selected during cluster creation. A subnet provides dedicated network resources that are isolated from other networks, improving network security.
VPC	VPC selected during cluster creation. A VPC is a secure, isolated, and logical network environment.
OBS Permission Control	Click **Manage** and modify the mapping between MRS users and OBS permissions. For details, see **Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS**.
Creating a data connection	Click **Manage** to view the data connection type associated with the cluster. For details, see **Configuring Data Connections**.

Parameter	Description
Agency	Click **Manage Agency** to bind or modify an agency for the cluster. An agency allows ECS or BMS to manage MRS resources. You can configure an agency of the ECS type to automatically obtain the AK/SK to access OBS. For details, see **Configuring a Storage-Compute Decoupled Cluster (Agency)**. The **MRS_ECS_DEFAULT_AGENCY** agency has the OBSOperateAccess permission of OBS and the CESFullAccess (for users who have enabled fine-grained policies), CES Administrator, and KMS Administrator permissions in the region where the cluster is located.
Key Pair	Name of a key pair. Set this parameter when creating a cluster. If the login mode is set to password during cluster creation, this parameter is not displayed.
Kerberos Authentication	Whether to enable Kerberos authentication when logging in to Manager. **NOTE** Kerberos authentication cannot be manually enabled or disabled after the cluster is created. Set this parameter with caution when creating a cluster. If you need to change the authentication status, you are advised to create a new cluster.
Enterprise Project	Enterprise project to which a cluster belongs. Only on the **Active Clusters** page, you can click the name of an enterprise project to go to its **Enterprise Project Management** page.
Security Group	Security group name of the cluster.
Streaming Core Node LVM	Indicates whether to enable the Logical Volume Manager (LVM) function of streaming Core nodes.
Data Disk Key Name	Name of the key used to encrypt data disks. To manage the used keys, log in to the key management console.
Data Disk Key ID	ID of the key used to encrypt data disks.
Component Version	Version of each component installed in the cluster.
Agency	Delegates ECSs or BMSs to manage some of your resources.

Go back to the historical clusters page. You can use the following buttons to perform operations. For details about the buttons, see the following table.

Table 2-19 Icon description

Icon	Description
	Click  to manually refresh the node information.
	Enter a cluster name in the search bar and click  to search for a cluster.

3 Managing Clusters

3.1 Logging In to a Cluster

3.1.1 MRS Cluster Node Overview

An MRS cluster consists of multiple ECSs. The system manages nodes in node groups based on specifications. Nodes in the same node group use same ECS specifications. Nodes in a cluster can be classified into Master nodes, Core nodes, and Task nodes based on the roles of components deployed on the nodes. For details about the node types, see **Table 3-1**.

Table 3-1 Cluster node types

Node Type	Functions
Master node	Management node of an MRS cluster. It manages and monitors the cluster. In the navigation tree of the MRS management console, choose **Clusters** > **Active Clusters**, select a running cluster, and click its name to switch to the cluster details page. On the **Nodes** tab page, view the **Name**. The node that contains **master1** in its name is the Master1 node. The node that contains **master2** in its name is the Master2 node. You can log in to a Master node either using VNC on the ECS management console or using SSH. After logging in to the Master node, you can access Core nodes without entering passwords. The system automatically deploys the Master nodes in active/standby mode and supports the high availability (HA) feature for MRS cluster management. If the active management node fails, the standby management node switches to the active state and takes over services. To determine whether the Master1 node is the active management node, see **Determining Active and Standby Management Nodes**.
Core node	Work node of an MRS cluster. It processes and analyzes data and stores process data. In the **Nodes** tab of the cluster details page, the nodes in the node group whose **Node Type** includes **Core** are core nodes.
Task node	Compute node. When the compute resources of a cluster are insufficient, you can configure elastic scaling policies to increase nodes automatically. In the **Nodes** tab of the cluster details page, the nodes in the node group whose **Node Type** is **Task** are task nodes. If only the **NodeManager (Yarn)** or **Supervisor (Storm)** role is deployed in a node group in addition to basic mandatory roles, this node group is a Task node group.

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

- GUI login: Use the remote login function provided by the ECS management console to log in to the Linux interface of the Master node in the cluster.
- SSH login: Applies to Linux ECSs only. You can use a remote login tool (such as PuTTY) to log in to an ECS. The ECS must have a bound EIP.

For details about how to apply for and bind EIP for the Master node, see **Assigning an EIP and Binding It to an ECS**.

You can log in to a Linux ECS using either a key pair or password.

NOTICE

If you need to use a key pair to access a cluster node, you need to log in to the node as user **root**. For details, see **Logging In to an ECS Using a Key Pair (SSH)**.

For details about how to access a cluster node using a password, see **Logging In to an ECS Using a Password (SSH)**.

3.1.2 Logging In to an ECS

This section describes how to remotely log in to an ECS in an MRS cluster using the remote login (VNC mode) function provided on the ECS management console or a key or password (SSH mode). Remote login (VNC mode) is mainly used for emergency O&M. In other scenarios, it is recommended that you log in to the ECS using SSH.

NOTE

To log in to a cluster node using SSH, you need to add an inbound rule to the security group of the cluster. Set **Source** to *IPv4 address of the client*/**32** or *IPv6 address of the client*/**128** and set the port number to **22**. For details, see **Adding a Security Group Rule**.

Logging In to an ECS Using VNC

Step 1 Log in to the MRS management console.

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

Step 3 On the **Nodes** tab page, click the name of a Master node in the Master node group to log in to the ECS management console.

Step 4 In the upper right corner, click **Remote Login**.

Step 5 Enter the username and password for logging in to the Master node as prompted.

1. If you select **Password** for **Login Mode**, you need to enter **root** in **Username** and the password you set during cluster creation in **Password**.

Figure 3-1 Selecting password as the login mode

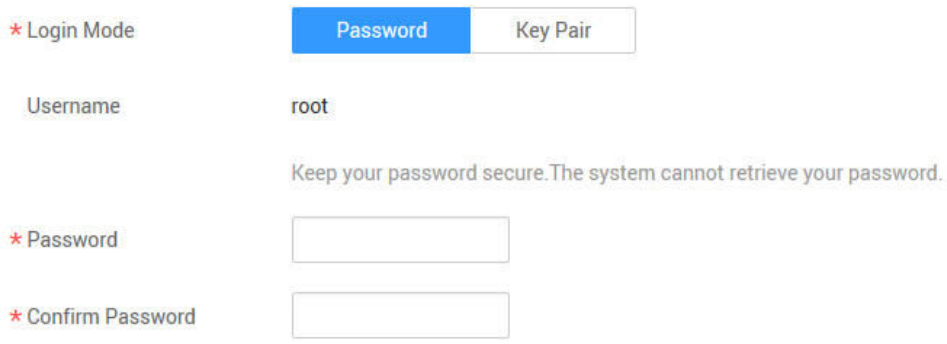

2. If you select **Key Pair** for **Login Mode** when creating a cluster, perform the following operations to log in to the cluster:
   a. After the cluster is created, assign an EIP and bind it to the Master node of the cluster. For details, see **Assigning an EIP and Binding It to an ECS**.
   b. Remotely log in to the Master node in SSH mode as user **root** using the key file.
   c. Run the **passwd root** command to set a password for user **root**.
   d. Go back to the login interface, and enter **root** and the password set in **Step 5.2.c** to log in to the node.

----End

Logging In to an ECS Using a Key Pair (SSH)

Logging In to the ECS from Local Windows

To log in to the Linux ECS from local Windows, perform the operations described in this section. The following procedure uses PuTTY as an example to log in to the ECS.

1. Log in to the MRS management console.
2. Choose **Clusters** > **Active Clusters**, select a running cluster, and click its name to switch to the cluster details page.
3. On the **Nodes** tab page, click the name of a Master node in the Master node group to log in to the ECS management console.
4. Click the **EIPs** tab, click **Bind EIP** to bind an EIP to the ECS, and record the EIP. If an EIP has been bound to the ECS, skip this step.
5. Check whether the private key file has been converted to **.ppk** format.
   - If yes, go to **10**.
   - If no, go to **6**.
6. Run PuTTY.
7. In the **Actions** area, click **Load** and import the private key file you used during ECS creation.

   Ensure that the private key file is in the format of **All files (*.*)**.
8. Click **Save private key**.

9. Save the converted private key, for example, **kp-123.ppk**, to a local directory.
10. Run PuTTY.
11. Choose **Connection** > **Data**. Enter the image username in **Auto-login username**.

    NOTE

    The image username for cluster nodes is **root**.

12. Choose **Connection** > **SSH** > **Auth**. In the last configuration item **Private key file for authentication**, click **Browse** and select the private key converted in 9.
13. Click **Session**.
    a. **Host Name (or IP address)**: Enter the EIP bound to the ECS.
    b. **Port**: Enter **22**.
    c. **Connection Type**: Select **SSH**.
    d. **Saved Sessions**: Task name, which can be clicked for remote connection when you use PuTTY next time

    Figure 3-2 Clicking **Session**

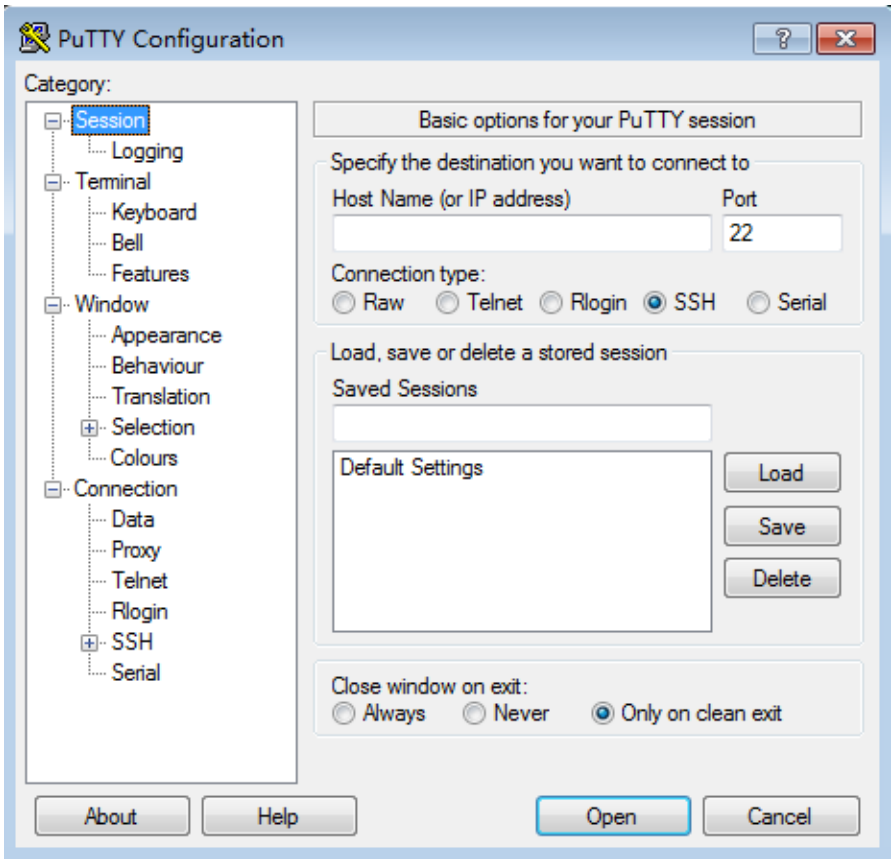

14. Click **Open** to log in to the ECS.

    If you log in to the ECS for the first time, PuTTY displays a security warning dialog box, asking you whether to accept the ECS security certificate. Click **Yes** to save the certificate to your local registry.

Logging In to the ECS from Local Linux

To log in to the Linux ECS from local Linux, perform the operations described in this section. The following procedure uses private key file **kp-123.pem** as an example to log in to the ECS. The name of your private key file may differ.

1. On the Linux CLI, run the following command to change operation permissions:

   chmod 400 /*path*/kp-123.pem

   NOTE

   In the preceding command, *path* refers to the path where the key file is saved.

2. Run the following command to log in to the ECS:

   ssh -i /*path*/kp-123.pem*Default username*@*EIP*

   For example, if the default username is **root** and the EIP is **123.123.123.123**, run the following command:

   ssh -i /*path*/kp-123.pem root@123.123.123.123

   NOTE

   - *path* indicates the path where the key file is saved.
   - *EIP* indicates the EIP bound to the ECS.
   - The image username is **root** for cluster nodes.

Logging In to an ECS Using a Password (SSH)

Logging In to the ECS from Local Windows

To log in to the Linux ECS from local Windows, perform the operations described in this section. The following procedure uses PuTTY as an example to log in to the ECS.

Step 1 Log in to the MRS management console.

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

Step 3 On the **Nodes** tab page, click the name of a Master node in the Master node group to log in to the ECS management console.

Step 4 Click the **EIPs** tab, click **Bind EIP** to bind an EIP to the ECS, and record the EIP. If an EIP has been bound to the ECS, skip this step.

Step 5 Run PuTTY.

Step 6 Click **Session**.

1. **Host Name (or IP address)**: Enter the EIP bound to the ECS.
2. **Port**: Enter **22**.
3. **Connection Type**: Select **SSH**.
4. **Saved Sessions**: Task name, which can be clicked for remote connection when you use PuTTY next time

Figure 3-3 Clicking **Session**

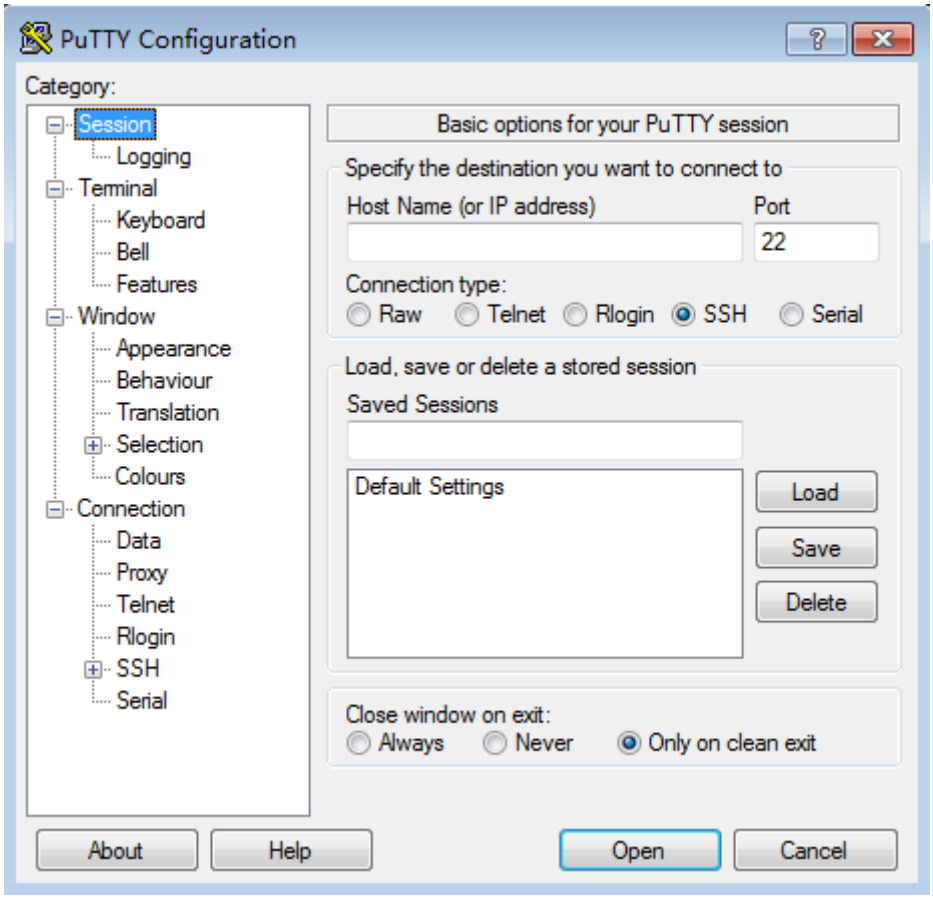

Step 7 Click **Window** and select **UTF-8** for **Remote character set:** in **Translation**.

Step 8 Click **Open** to log in to the ECS.

If you log in to the ECS for the first time, PuTTY displays a security warning dialog box, asking you whether to accept the ECS security certificate. Click **Yes** to save the certificate to your local registry.

Step 9 After the SSH connection to the ECS is set up, enter the username and password as prompted to log in to the ECS.

NOTE

The username is **root** and the password is the one you set during cluster creation.

----End

Logging In to the ECS from Local Linux

If the local host runs Linux, perform steps **Step 1** to **Step 4** to bind an EIP to the ECS, and run the following command on the CLI to log in to the ECS: **ssh *EIP bound by the ECS***

3.1.3 Determining Active and Standby Management Nodes

Scenario

Some O&M operation scripts and commands need to be run or can be run only on the active management node. You can log in to a Master node or the Manager

(MRS 3.x or later) to determine the active and standby management nodes (active and standby OMS nodes).

In active/standby mode, a switchover can be implemented between Master1 and Master2. For this reason, Master1 may not be the active management node for Manager.

Runing the Script to Determine Active and Standby Nodes

Step 1 Find the Master nodes of an MRS cluster.

1. Log in to the MRS console, choose **Clusters** > **Active Clusters** and click the name of the target cluster to access its details page.
2. On the **Nodes** tab page, view Master node names. The node that contains **master1** in its name is the Master1 node. The node that contains **master2** in its name is the Master2 node.

Step 2 Determine the active and standby management nodes of the Manager.

1. Remotely log in to the Master1 node. For details, see **Logging In to an ECS**.

   Master nodes support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation.
2. Run the following commands to switch the user:

   sudo su - root

   su - omm
3. Run the following command to identify the active and standby management nodes:

   For versions earlier than MRS 3.*x*, run the **sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh** command.

   For MRS 3.*x* or later: Run the **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** command.

   In the command output, the node whose **HAActive** is **active** is the active management node (mgtomsdat-sh-3-01-1 in the following example), and the node whose **HAActive** is **standby** is the standby management node (mgtomsdat-sh-3-01-2 in the following example).

   ```
   Ha mode
   double
   NodeName            HostName                HAVersion       StartTime           HAActive
   HAAllResOK          HARunPhase
   192-168-0-30        mgtomsdat-sh-3-01-1     V100R001C01     20xx-11-18 23:43:02
   active              normal                  Actived
   192-168-0-24        mgtomsdat-sh-3-01-2     V100R001C01     20xx-11-21 07:14:02
   standby             normal                  Deactived
   ```

   NOTE

   If the Master1 node to which you have logged in is the standby management node and you need to log in to the active management node, run the following command:

   ssh *IP address of Master2 node*

----End

Logging in to Manager to Determine Active and Standby Nodes

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

Step 1 Log in to Manager. For details, see **Accessing Manager**.

Step 2 Click **Hosts**. The **Hosts** page is displayed.

Step 3 View and record the IP addresses of the active and standby management nodes.

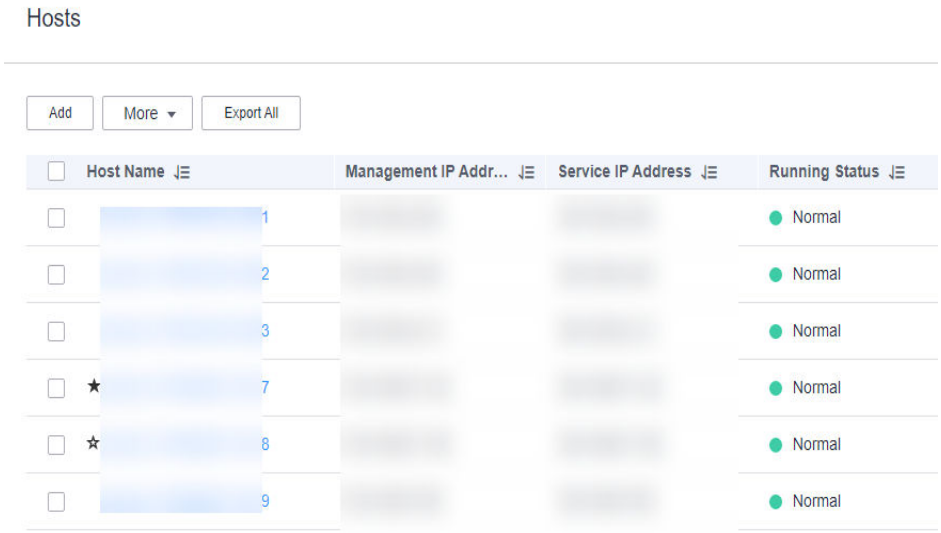

- If a host name starts with ★, it is the active management node (active OMS node). View and record the value of **Management IP Address** in the row containing the active node.
- If a host name starts with ☆, the host is a standby management node (standby OMS node). View and record the value of **Management IP Address** in the row containing the standby node.

----End

3.2 Cluster Overview

3.2.1 Cluster List

You can quickly view the status of all clusters and jobs by viewing the dashboard information, and obtain relevant MRS documents from **Help** in the left navigation pane on the MRS console.

MRS is used to manage and analyze massive data. It is easy to use. You can create a cluster and add MapReduce, Spark, and Hive jobs to the cluster to analyze and process user data. After being processed, you can transmit the data in SSL encryption mode to OBS to ensure data integrity and confidentiality.

Cluster Status

Log in to the MRS management console. You can view the status of existing clusters in the active cluster list. You search for status clusters in a specified status from the **Status** drop-down list. **Table 3-2** lists all the cluster statuses.

Table 3-2 Cluster status

Status	Description
Starting	If a cluster is being created, the cluster is in the **Starting** state.
Running	If a cluster is successfully created and is running properly, its status is **Running**.
Scaling out	If the Master, Core, or Task node in a cluster is being added, the cluster is in the **Scaling out** state. **NOTE** If the cluster scale-out fails, you can add node to the cluster again.
Scaling in	If a cluster node is being deleted, the cluster node is in the **Scaling in** state. This state shows when you scale in or elastically scale in a cluster node, decommission a node in a yearly/monthly cluster, change the OS, or reinstall the OS.
Abnormal	If some components in a cluster are abnormal, the cluster is **Abnormal**.
Deleting	If a pay-per-use cluster node is being deleted, the cluster status changes to **Deleting**. This state is displayed after you click **Delete** and confirm the deletion. **NOTE** A yearly/monthly cluster cannot be deleted.
Frozen	If the grace period of a yearly/monthly resource expires but the resource is not renewed, or if the fee of a pay-per-use resource fails to be deducted and you have not topped up your account before the grace period expires, the system freezes the resource in the **Frozen** state. **NOTE** A frozen cluster is unavailable and its all ECSs are shut down. After being unfrozen, the cluster returns to the **Running** state. If no renewal fee is paid, the cluster will be deleted after a specified period (called freeze period) and the cluster status will be changed to **Deleted**.
Restoring node...	If a faulty node in the cluster is being recovered, its status is **Restoring node...**.

Job Status

Table 3-3 describes the status of jobs that you execute after logging in to the MRS management console.

Table 3-3 Job status

Status	Description
Accepted	Initial status of a job after it is successfully submitted.

Status	Description
Running	A job is being executed.
Completed	A job has been executed and completed successfully.
Terminated	A job is stopped during execution.
Abnormal	An error occurs during job execution or job execution fails.

3.2.2 Checking the Cluster Status

The cluster list contains all clusters in MRS. You can view clusters in various states. If a large number of clusters are involved, navigate through multiple pages to view all of the clusters.

MRS, as a platform managing and analyzing massive data, provides a PB-level data processing capability. MRS allows you to buy multiple clusters. The cluster quantity is subject to that of ECSs.

Clusters are listed in chronological order by default in the cluster list, with the most recent cluster displayed at the top. **Table 3-4** describes the cluster list parameters.

- **Active Clusters**: contain all clusters except the clusters in the **Failed** and **Deleted** states.
- **Cluster History**: contains the clusters in the **Deleted** states. Only clusters deleted within the last six months are displayed. If you want to view clusters deleted six months ago, contact Huawei Cloud technical support.
- **Failed Tasks**: only contain the tasks in the **Failed** state. You can click on the **Active Cluster** page to view failed tasks.

Table 3-4 Parameters in the active cluster list

Parameter	Description
Name/ID	Cluster name, which is set when a cluster is created. Unique identifier of a cluster, which is automatically assigned when a cluster is created. • : Change the cluster name. • : Copy the cluster ID.
Cluster Version	Cluster version.
Cluster Type	The type of a cluster you want to create.
Nodes	Number of nodes that can be deployed in a cluster. This parameter is set when a cluster is created.

Parameter	Description
Status	Status and operation progress description of a cluster. The cluster creation progress includes: • Verifying cluster parameters • Applying for cluster resources • Creating VMs • Initializing VMs • Installing MRS Manager • Deploying the cluster • Cluster installation failed The cluster scale-out progress includes: • Preparing for scale-out • Creating VMs • Initializing VMs • Adding nodes to the cluster • Scale-out failed The cluster scale-in progress includes: • Preparing for scale-in • Decommissioning instance • Deleting VMs • Deleting nodes from the cluster • Scale-in failed The system will display causes of cluster installation, scale-out, and scale-in failures. For details, see **Table 2-6**.
Billing Mode	Currently, the commercial version of MRS is charged based on ECSs in a cluster. • **Yearly/Monthly**: The duration ranges from one month to one year. The minimum cluster duration is 1 month and the maximum available cluster duration is 1 year. • **Pay-per-use**: Nodes are charged by actual duration of use, with a billing cycle of one hour. The billing start time, that is, the time when the cluster node is successfully created is displayed under **Billing Mode**.
Created	The time when a cluster node is successfully created. This parameter is displayed only on the **Cluster History** page.
Deleted	Time when a cluster node stops and the cluster node begins to be deleted. This parameter is valid only for historical clusters displayed on the **Cluster History** page.
AZ	Availability zone (AZ) in the region of a cluster, which is set when a cluster is created.

Parameter	Description
Enterprise Project	Enterprise project to which a cluster belongs.
Operation	• **Delete**: If the cluster is no longer needed after the job execution is complete, you can click **Delete**. To confirm deletion, enter **DELETE** in the displayed dialog box and click **OK**. The cluster status changes from **Running** to **Deleting**. After the cluster is deleted, the cluster status changes to **Deleted** and is displayed in the historical cluster list. A cluster failed to be deployed will be automatically deleted. This parameter is displayed in **Active Clusters** only. **NOTE** Typically after data is analyzed and stored, or when the cluster encounters an exception and cannot work, you can delete a cluster. If a cluster is deleted before data processing and analysis are completed, data loss may occur. Therefore, exercise caution when deleting a cluster.

Table 3-5 Button description

Button	Description
	Select an enterprise project from the drop-down list to filter the corresponding cluster.
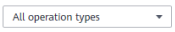	Select a status to filter clusters from the drop-down list: • All statuses • Starting • Running • Scaling out • Scaling in • Abnormal • Deleting • Restoring node... • Frozen
	Choose **Clusters > Active Clusters** and click  to go to the page for managing failed tasks.  *Num*: displays the failed tasks in the **failed** state.
	Enter a cluster name in the search bar and click  to search for a cluster.

Button	Description
Search by Tag	Click **Search by Tag**, enter the tag of the cluster to be queried, and click **Search** to search for the clusters. You can select a tag key or tag value from their drop-down lists. When the tag key or tag value is exactly matched, the system can automatically locate the target cluster. If you enter multiple tags, their intersections are used to search for the cluster.
	Click  to manually refresh the cluster list.

3.2.3 Viewing Basic Cluster Information

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

NOTE

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

On the cluster details page, click **Dashboard**. **Table 3-6**, **Table 3-7**, **Table 3-8**, and **Table 3-9** describe the parameters on the **Dashboard** tab page.

Table 3-6 Basic information

Parameter	Description
Cluster Name	The name of a cluster. Configure this parameter when creating a cluster. Click  to change the cluster name. For versions earlier than MRS 3.x, only the cluster name displayed on the MRS management console is changed, while the cluster name on MRS Manager is not changed synchronously.
Cluster Status	The cluster status. For details, see **Table 3-2**.
Cluster Version	MRS version information.

Parameter	Description
Cluster Type	There are three types of clusters: • **Analysis Cluster**: is used for offline data analysis and provides Hadoop components. • **Streaming Cluster**: is used for streaming tasks and provides stream processing components. • **Hybrid Cluster**: is used for both offline data analysis and streaming processing and provides Hadoop components and streaming processing components. • **Custom**: An MRS cluster with all custom components. MRS 3.*x* or later supports this type.
Cluster ID	Unique identifier of a cluster, which is automatically assigned when a cluster is created.
Created	Time when a cluster is created.
AZ	Availability zone (AZ) in the region of a cluster, which is set when a cluster is created.
Kerberos Authentication	Whether to enable Kerberos authentication when logging in to Manager.
Enterprise Project	The enterprise project to which a cluster belongs. Only on the **Active Clusters** page, you can click the name of an enterprise project to go to its **Enterprise Project Management** page.

Table 3-7 Network information

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

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

Table 3-8 O&M management

Parameter	Description
MRS Manager	Portal for the Manager page. • For MRS 3.*x* or later, see **Accessing FusionInsight Manager (MRS 3.x or Later)**. • For versions earlier than MRS 3.x, you need to bind an EIP and add a security group rule as prompted before accessing the MRS Manager page. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.
IAM User Sync	IAM user information (including federated users) can be synchronized to an MRS cluster for cluster management. For details, see **Synchronizing IAM Users to MRS**. **NOTE** The **Components**, **Tenants**, and **Backups & Restorations** tab pages on the cluster details page can be used only after users are synchronized. After an MRS 3.x cluster is synchronized, you can use the **Components** function. For a federated user, only information about the user that logged in the system can be synchronized.
Data Connection	Click **Manage** to view the data connection type associated with the cluster. For details, see **Configuring Data Connections**.

Parameter	Description
Agency	Click **Manage Agency** to bind or modify an agency for the cluster. An agency allows ECS or BMS to manage MRS resources. You can configure an agency of the ECS type to automatically obtain the AK/SK to access OBS. For details, see **Configuring a Storage-Compute Decoupled Cluster (Agency)**. The **MRS_ECS_DEFAULT_AGENCY** agency has the OBS OperateAccess permission of OBS and the CES FullAccess (for users who have enabled fine-grained policies), CES Administrator, and KMS Administrator permissions in the region where the cluster is located.
OBS Permission Control	Click **Manage** and modify the mapping between MRS users and OBS permissions. For details, see **Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS**.
Logging	Used to collect logs about cluster creation and scaling failures.
Secure Communications	Used to display the security authorization status. You can click  to enable or disable security authorization. Disabling security authorization brings high risks. Exercise caution when performing this operation. For details, see **Communication Security Authorization**.

Table 3-9 Billing information

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

On the cluster details page, click **Nodes**. For details about the node parameters, see **Table 3-10**.

Table 3-10 Node information

Parameter	Description
Configure Task Node	Used to add a Task node. For details, see **Adding a Task Node**. For 3.x and later versions, this operation applies only to the analysis cluster, streaming cluster, and hybrid cluster.
Add Node Group	Used to add node groups. This method applies only to customized clusters of 3.x or later. For details, see **Adding a Node Group**.
Node Group	Node group name.
Node Type	Node type: • **Master**: A Master node in an MRS cluster manages the cluster, assigns MapReduce executable files to Core nodes, traces the execution status of each job, and monitors the DataNode running status. • A task node group is a group of nodes where only data roles that do not store data are deployed. The roles include NodeManager, ThriftServer, Thrift1Server, RESTServer, Supervisor, LogViewer, HBaseIndexer, and TagSync. • If other roles are deployed in the node group in addition to the preceding roles, the node group is the Core node group. On the **Nodes** tab page, click ⌄ next to a node group name to unfold the nodes contained in the node group. Click a node name to remotely log in to the ECS using the password or key pair configured during cluster creation. For details about the parameters, see **Managing Components and Monitoring Hosts**.
Node Count	Number of nodes in a node group.
Billing Mode	Billing mode of the cluster you purchased, including **Pay-Per-Use** and **Yearly/Monthly**.

Parameter	Description
Operation	• **Scale Up Specifications**: For details, see **Scaling Up Master Node Specifications**. • **Scale Out**: For details, see . • **Scale In**: For details, see **Scaling In a Cluster**. • **Delete**: To delete a node group, ensure that there is no node in the node group. • **Decommission Node**: To decommission a node group, decommission all its nodes. This parameter is available for yearly/monthly clusters only. • **Recommission Node**: Recommission a node that has been decommissioned. This parameter is available for yearly/monthly clusters only. • **View Roles**: You can view information about roles deployed on the node group. This function applies only to custom clusters of 3.x and later.

3.2.4 Viewing Cluster Patch Information

To view patch information about cluster components, you can download the required patch if the cluster component, such as Hadoop or Spark, is faulty. On the MRS console, choose **Clusters > Active Clusters**, select a cluster, and click the cluster name. On the cluster details page that is displayed, upgrade the component and rectify the fault.

NOTE

MRS 3.x does not have patch version information. Therefore, this section is not involved.

3.2.5 Managing Components and Monitoring Hosts

You can manage the following status and metrics of all components (including role instances) and hosts on the MRS console:

- Status information: includes operation, health, configuration, and role instance status.
- Indicator information: includes key monitoring indicators for each component.
- Export monitoring metrics. (This function is not supported in MRS 3.x or later.)

NOTE

- For versions earlier than MRS 3.x, see **Managing Services and Monitoring Hosts**.
- For MRS 3.x or later, see **Procedure**.
- You can set the interval for automatically refreshing the page or click  to refresh the page immediately.
- Component management supports the following parameter values:
  - Refresh every 30 seconds
  - Refresh every 60 seconds
  - Stop refreshing

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Procedure

Manage components.

NOTE

For details about how to perform operations on MRS Manager, see **Managing Service Monitoring**.

Step 1 On the MRS cluster details page, click **Components**.

- **Table 3-11** describes the service operating status.

Table 3-11 Service operating status

Status	Description
Started	The service is started.
Stopped	The service is stopped.
Failed to start	Failed to start the role instance.
Failed to stop	Failed to stop the service.
Unknown	Indicates initial service status after the background system restarts.

- **Table 3-12** describes the service health status.

Table 3-12 Service health status

Status	Description
Good	Indicates that all role instances in the service are running properly.

Status	Description
Faulty	Indicates that the running status of at least one role instance is **Faulty** or the status of the service on which the current service depends is abnormal.
Unknown	Indicates that all role instances in the service are in the **Unknown** state.
Restoring	Indicates that the background system is restarting the service.
Partially Healthy	Indicates that the status of the service on which the service depends is abnormal, and APIs related to the abnormal service cannot be invoked by external systems.

- Table 3-13 describes the service health status.

  Table 3-13 Service configuration status

Status	Description
Synchronized	The latest configuration takes effect.
Configuration expired	The latest configuration does not take effect after the parameter modification. Related services need to be restarted.
Configuration failed	The communication is incorrect or data cannot be read or written during the parameter configuration. Use **Synchronize Configuration** to rectify the fault.
Configuring	Parameters are being configured.
Unknown	Indicates that configuration status cannot be obtained.

Step 2 Click a specified service in the list to view its status and metric information.

Step 3 Customize and view monitoring graphs.

1. In the **Charts** area, click **Customize** to customize service monitoring metrics.
2. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.

----End

Manage role instances.

NOTE

For versions earlier than MRS 3.x, see **Managing Role Instances**.

Step 1 On the MRS cluster details page, click **Components**. In the component list, click the specified service name.

Figure 3-4 Components tab page

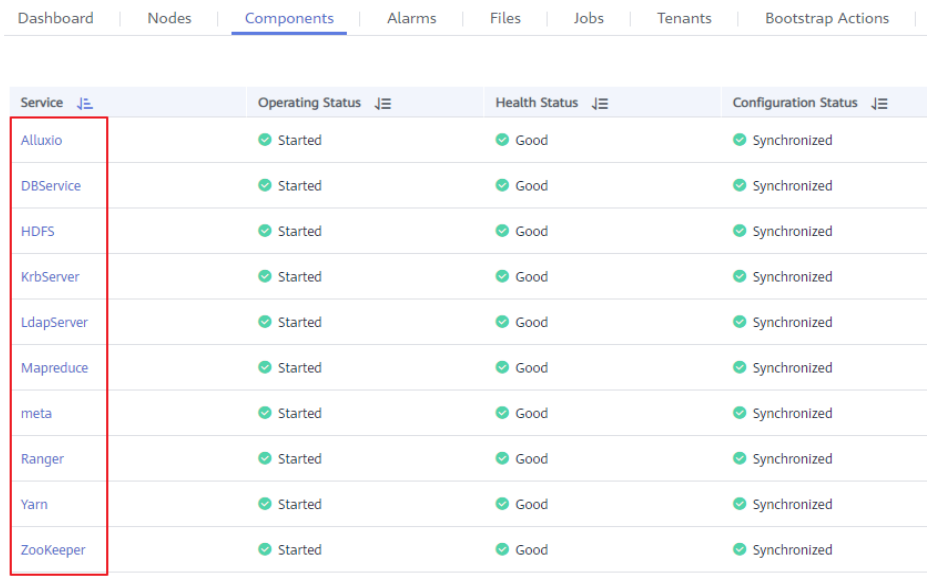

Step 2 Click **Instances** to view the role status.

Figure 3-5 Instances tab page

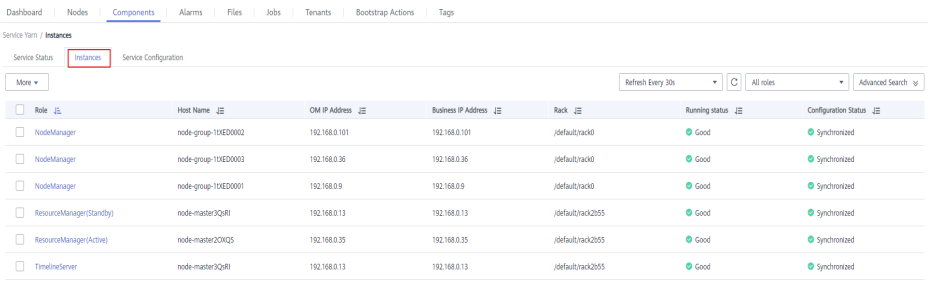

The role instance list contains the Role, Host Name, Management IP Address, Service IP Address, Rack, Running Status, and Configuration Status of each instance.

- **Table 3-14** shows the running status of a role instance.

Table 3-14 Role instance running status

Status	Description
Good	Indicates that the instance is running properly.
Bad	Indicates that the instance cannot run properly.
Decommissioned	Indicates that the instance is out of service.

Status	Description
Not started	Indicates that the instance is stopped.
Unknown	Indicates that the initial status of the instance cannot be detected.
Starting	Indicates that the instance is being started.
Stopping	Indicates that the instance is being stopped.
Restoring	Indicates that an exception may occur in the instance and the instance is being automatically rectified.
Decommissioning	Indicates that the instance is being decommissioned.
Recommissioning	Indicates that the instance is being recommissioned.
Failed to start	Indicates that the service fails to be started.
Failed to stop	Indicates that the service fails to be stopped.

- Table 3-15 shows the configuration status of a role instance.

  Table 3-15 Role instance configuration status

Status	Description
Synchronized	The latest configuration takes effect.
Configuration expired	The latest configuration does not take effect after the parameter modification. Related services need to be restarted.
Configuration failed	The communication is incorrect or data cannot be read or written during the parameter configuration. Use **Synchronize Configuration** to rectify the fault.
Configuring	Parameters are being configured.
Unknown	Current configuration status cannot be obtained.

By default, the **Role** column is sorted in ascending order. You can click the sorting icon next to **Role**, **Host Name**, **OM IP Address**, **Business IP Address**, **Rack**, **Running Status**, or **Configuration Status** to change the sorting mode.

You can filter out all instances of the same role in the **Role** column.

You can set search criteria in the role search area by clicking **Advanced Search**, and click **Search** to view specified role information. You can click **Reset** to reset the search criteria. Fuzzy search is supported.

Step 3 Click the target role instance to view its status and metric information.

Step 4 Customize and view monitoring graphs.

1. In the **Charts** area, click **Customize** to customize service monitoring metrics.
2. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.

----End

Manage hosts.

NOTE

For versions earlier than MRS 3.x, see **Managing Hosts**.

Step 1 On the MRS cluster details page, click the **Nodes** tab and expand a node group to view the host status.

The host list of a group contains the **Node Name/Resource ID**, **IP**, **Status**, **Specifications**, **Disks**, and **AZ**.

- **Table 3-16** shows the host operating status.

  Table 3-16 Host operating status

Status	Description
Normal	The host and service roles on the host are running properly.
Isolated	The host is isolated, and the service roles on the host stop running.

- **Table 3-17** describes the host health status.

  Table 3-17 Host health status

Status	Description
Good	The host can properly send heartbeats.
Bad	The host fails to send heartbeats due to timeout.
Unknown	The host initial status is unknown during the operation of adding or deleting a host.

By default, data is sorted in ascending order by node name. You can click to change the order.

Step 2 Click the target node in the list to view its status and metric information.

----End

3.3 Viewing and Customizing Cluster Monitoring Metrics

MRS cluster nodes are classified into management nodes, control nodes, and data nodes. The change trends of key host monitoring metrics on each type of node can be calculated and displayed as curve charts in reports based on the customized periods. If a host belongs to multiple node types, the metric statistics will be repeatedly collected.

This section provides overview of MRS clusters and describes how to view, customize, and export node monitoring metrics on MRS Manager.

NOTE

Cluster metrics are monitored periodically. The average historical monitoring interval is about 5 minutes.

- Scenario 1:
  - Method 1:
    i. Choose **Clusters** > **Active Clusters** and click a cluster name to access its details page.
    ii. Click the **Dashboard** tab, you can view the cluster host health status statistics on the lower part of the displayed tab page.
    iii. To view or export a report of other metrics, log in to Manager by referring to **Accessing Manager**.
    iv. On the Manager page, view, customize, and export the node monitoring metric report by referring to **Dashboard**.
  - Method 2:
    i. Log in to the MRS console.
    ii. Choose **Clusters** > **Active Clusters** and click a cluster name to access its details page.
    iii. On the **Dashboard** tab page, click **Click to synchronize** next to **IAM User Sync** to synchronize IAM users.
    iv. After the synchronization is complete, you can view the cluster monitoring metric report on the right of the page.
    v. In time range area, specify a period to view monitoring data. The options are as follows:
       - Last 1 hour
       - Last 3 hours
       - Last 12 hours
       - Last 24 hours
       - Recent 7 days
       - Recent 30 days
       - Customize: You can customize the period for viewing monitoring data.

vi. Customize a monitoring metric report.

1) Click **Customize** and select monitoring metrics to display.
2) Click **OK** to save the selected monitoring metrics for display.

NOTE

Click **Clear** to cancel all the selected monitoring metrics in a batch.

vii. Export a monitoring report.

1) Select a period.

The options are as follows:

Last 1 hour, **Last 3 hours**, **Last 12 hours**, **Last 24 hours**, **Recent 7 days**, **Recent 30 days**, or **Customize** (selection of a time range)

2) Click **Export**. MRS will generate a report about the selected monitoring metrics in a specified time range. Save the report.

- **Scenario 2 (applicable to MRS 3.*x* or later clusters)**

a. Log in to the MRS console.
b. Choose **Clusters** > **Active Clusters** and click a cluster name to access its details page.
c. On the **Dashboard** tab page, click **Click to synchronize** next to **IAM User Sync** to synchronize IAM users.
d. After the synchronization is complete, you can view the cluster monitoring metric report on the right of the page.
e. In time range area, specify a period to view monitoring data. The options are as follows:

- Last 1 hour
- Last 3 hours
- Last 12 hours
- Last 24 hours
- Recent 7 days
- Recent 30 days
- Customize: You can customize the period for viewing monitoring data.

f. Customize a monitoring metric report.

i. Click **Customize** and select monitoring metrics to be displayed.
ii. Click **OK** to save the selected monitoring metrics for display.

NOTE

Click **Clear** to cancel all the selected monitoring metrics in a batch.

3.4 Cluster O&M

3.4.1 Importing and Exporting Data

Through the **Files** tab page, you can create, delete, import, export, delete files in the analysis cluster. Currently, file creation is not supported. Streaming clusters do not support the file management function on the MRS GUI. In a cluster with Kerberos authentication enabled, to read or write the folders in the root directory, add a role that has the required permissions on the folders by referring to **Creating a Role**. Then, add the new role to the user group to which the user who submits the job belongs by referring to **Related Tasks**.

Background

Data sources processed by MRS are from OBS or HDFS. OBS is an object-based storage service that provides you with massive, secure, reliable, and cost-effective data storage capabilities. MRS can process data in OBS directly. You can view, manage, and use data by using the web page of the management control platform or OBS client. In addition, you can use REST APIs independently or integrate APIs to service applications to manage and access data.

Before creating jobs, upload the local data to OBS for MRS to compute and analyze. MRS allows exporting data from OBS to HDFS for computing and analyzing. After the data analysis and computing are completed, you can store the data in HDFS or export them to OBS. HDFS and OBS can also store the compressed data in the format of **bz2** or **gz**.

Importing Data

Currently, MRS can only import data from OBS to HDFS. The file upload rate decreases with the increase of the file size. This mode applies to scenarios where the data volume is small.

You can perform the following steps to import files and directories:

1. Log in to the MRS console.
2. Choose **Clusters > Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's information.
3. Click the **Files** tab, and go to the file management page.
4. Select **HDFS File List**.
5. Go to the data storage 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.

   The requirements for creating a folder are as follows:

   - The folder name contains a maximum of 255 characters
   - The folder name cannot be empty.
   - The folder name cannot contain the following special characters: /:*?"<>|\;&,'`!{}[]$%+
   - The value cannot start or end with a period (.).
   - The spaces at the beginning and end are ignored.
6. Click **Import Data** and configure the HDFS and OBS paths correctly. When configuring the OBS or HDFS path, click **Browse**, select a file directory, and click **Yes**.

Figure 3-6 Importing data

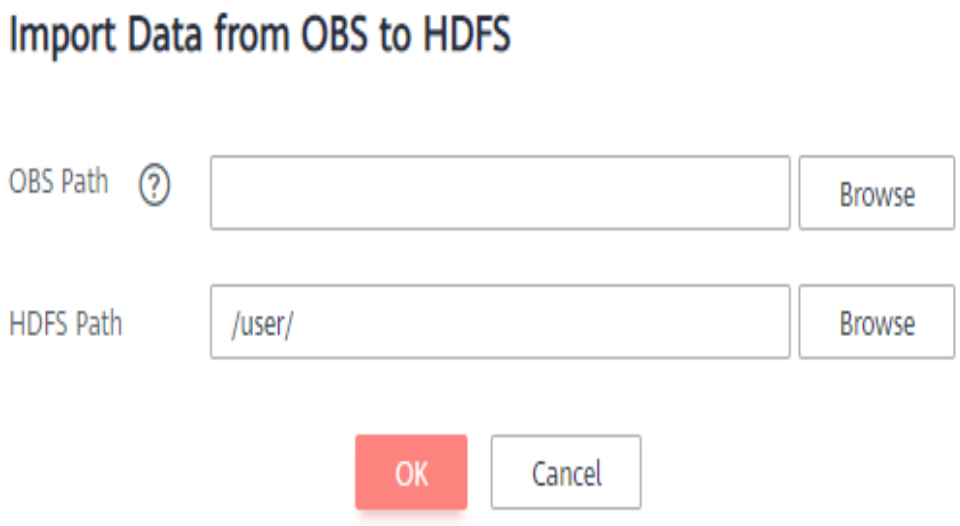

- OBS path
  - The path must start with **obs://**.
  - Files or programs encrypted by KMS cannot be imported.
  - An empty folder cannot be imported.
  - The directory and file name can contain letters, digits, hyphens (-), and underscores (_), but cannot contain the following special characters ;|&>,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

7. Click **OK**.

   You can view the file upload progress on the **File Operation Records** tab page. MRS processes the data import operation as a DistCp job. You can also check whether the DistCp job is successfully executed on the **Jobs** tab page.

Exporting Data

After the data analysis and computing are completed, you can store the data in HDFS or export them to OBS.

You can perform the following steps to export files and directories:

1. Log in to the MRS console.
2. Choose **Clusters > Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's basic information.
3. Click the **Files** tab, and the file management page is displayed.
4. Select **HDFS File List**.
5. Go to the data storage directory, for example, **bd_app1**.
6. Click **Export Data** and configure the OBS and HDFS paths. When configuring the OBS or HDFS path, click **Browse**, select a file directory, and click **Yes**.

   Figure 3-7 Exporting data

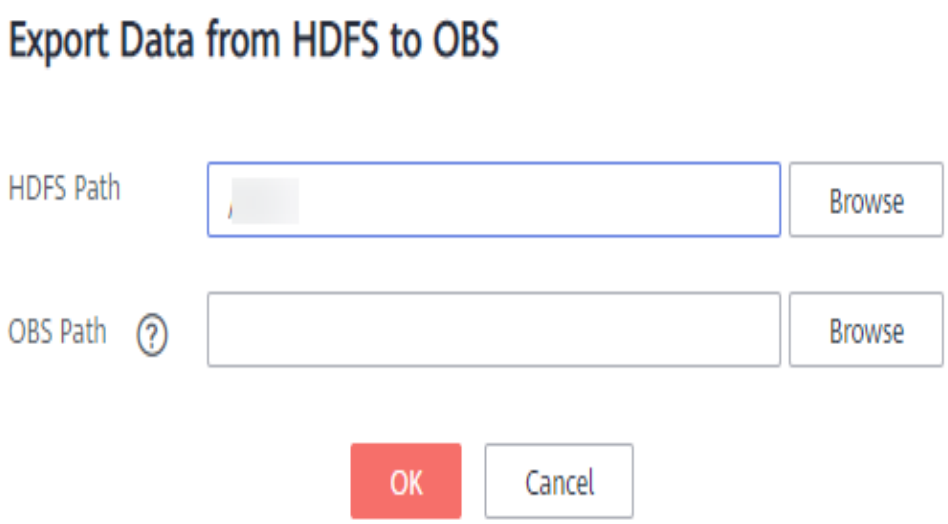

   - OBS path
     - The path must start with **obs://**.
     - The directory and file name can contain letters, digits, hyphens (-), and underscores (_), but cannot contain the following special characters: ;|&>,<'$*?\
     - The directory and file name cannot start or end with a space, but can contain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

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

NOTE

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

7. Click **OK**.

   You can view the file upload progress on the **File Operation Records** tab page. MRS processes the data export operation as a DistCp job. You can also check whether the DistCp job is successfully executed on the **Jobs** tab page.

Viewing Operation Logs

When importing and exporting data on the MRS management console, you can choose **Files > File Operation Records** to view the data import and export progress.

Table 3-18 describes the parameters of the file operation record.

Table 3-18 File operation record parameters

Parameter	Description
Created	Time when the data import or export task is created.
Source Path	Source path of data. • OBS path during data import. • HDFS path during data export.
Target Path	Target path of data. • HDFS path during data import. • OBS path during data import.
Status	Status during data import or export. • Submitted • Accepted • Running • Completed • Terminated • Abnormal

Parameter	Description
Duration (min)	Time of data import or export. The unit is minute.
Result	Result of data import or export. • Successful • Failed • Killed • Undefined
Operation	View Log: allows you to view file operation logs.

3.4.2 Changing the Subnet of a Cluster

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

For details about how to configure network ACL outbound rules, see **How Do I Configure a Network ACL Outbound Rule?**

Changing a Subnet When No Network ACL Is Associated

Step 1 Log in to the MRS console.

Step 2 Click the target cluster name to go to its details page.

Step 3 Click **Change Subnet** on the right of **Default Subnet**.

Figure 3-8 Changing a subnet

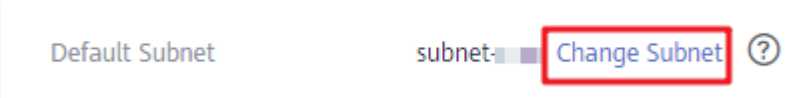

Step 4 Select the target subnet and click **OK**.

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

Figure 3-9 Selecting the target subnet

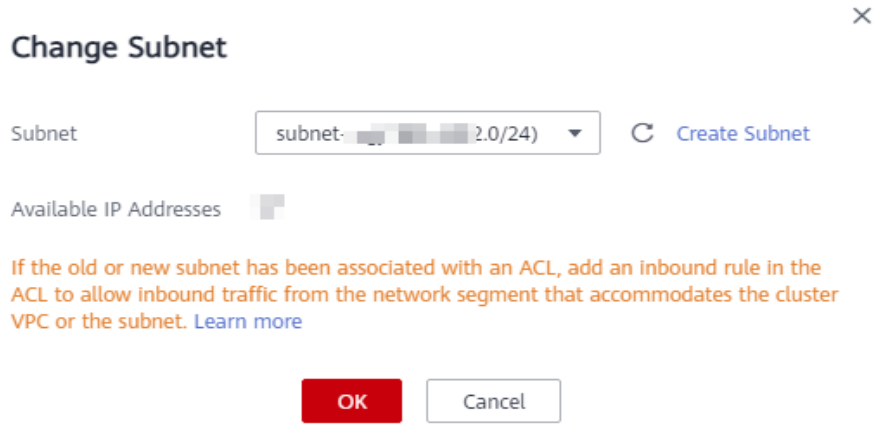

----End

Changing a Subnet When a Network ACL Is Associated

Step 1 Log in to the MRS console and click the target cluster to go to its details page.

Step 2 In the **Basic Information** area, view **VPC**.

Step 3 Log in to the VPC console. In the navigation pane on the left, choose **Virtual Private Cloud** and obtain the IPv4 CIDR block corresponding to the VPC obtained in **Step 2**.

Figure 3-10 Obtaining the IPv4 CIDR block

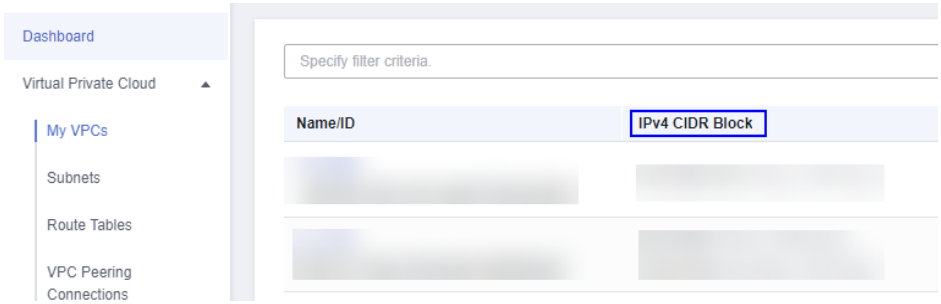

Step 4 Choose **Access Control** > **Network ACLs** and click the name of the network ACL that is associated with the default and new subnets.

NOTE

If both the default and new subnets are associated with a network ACL, add inbound rules to the network ACL by referring to **Step 5** to **Step 7**.

Figure 3-11 Network ACLs

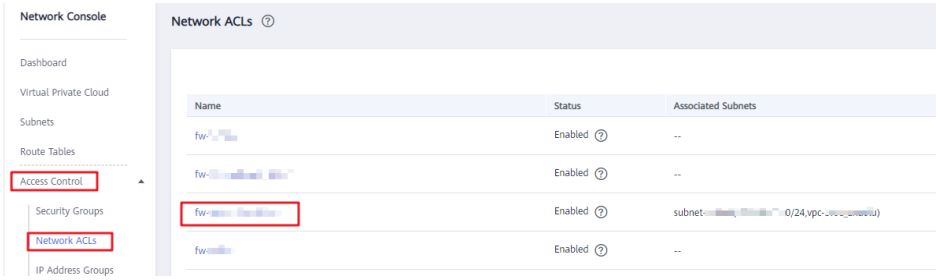

Step 5 On the **Inbound Rules** page, choose **More** > **Insert Rule Above** in the **Operation** column.

Figure 3-12 Inserting Rule Above

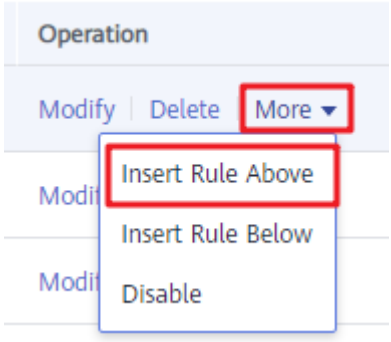

Step 6 Add a network ACL rule. Set **Action** to **Allow**, **Source** to the VPC IPv4 CIDR block obtained in **Step 3**, and retain the default values for other parameters.

Figure 3-13 Adding a network ACL rule

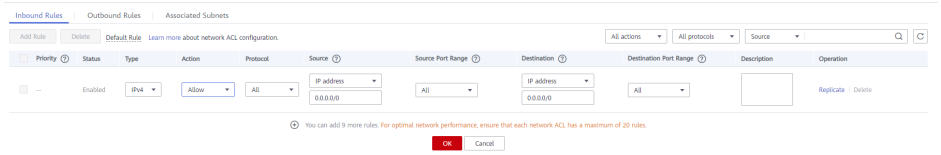

Step 7 Click **OK**.

NOTE

If you do not want to allow access from all IPv4 CIDR blocks of the VPC, add the IPv4 CIDR blocks of the default and new subnets by performing **Step 8** to **Step 12**. If the rules for VPC IPv4 CIDR blocks have been added, skip **Step 8** to **Step 12**.

Step 8 Log in to the MRS console.

Step 9 Click the target cluster to go to its details page.

Step 10 Click **Change Subnet** on the right of **Default Subnet**.

Figure 3-14 Changing a subnet

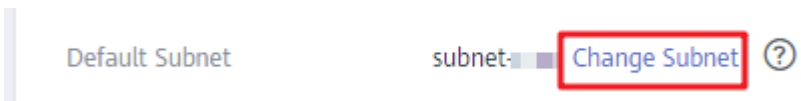

Step 11 Obtain the IPv4 CIDR blocks of the default and new subnets.

NOTICE

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

Figure 3-15 Obtaining the subnet IP address

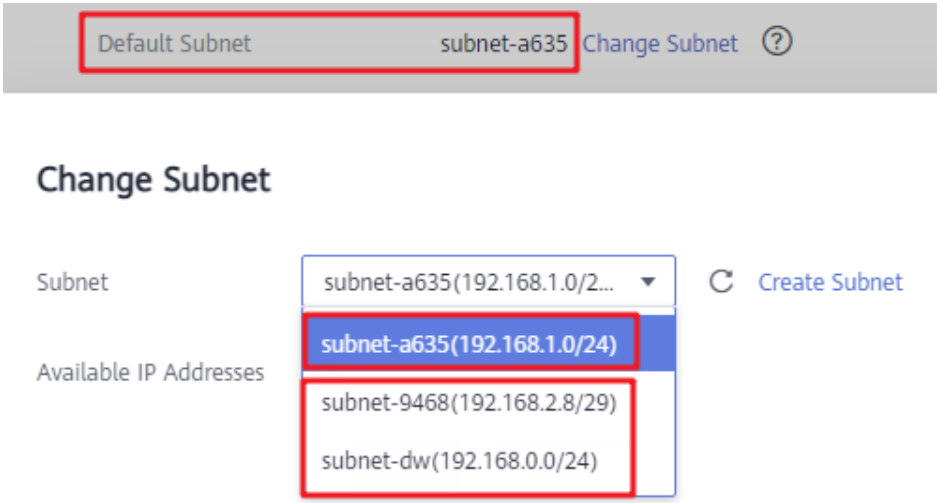

Step 12 Add the IPv4 CIDR blocks of the default and target subnets to the inbound rules of the network ACL bound to the two subnets by referring to **Step 4** to **Step 7**.

Step 13 Log in to the MRS console.

Step 14 Click the target cluster to go to its details page.

Step 15 Click **Change Subnet** on the right of **Default Subnet**.

Figure 3-16 Changing a subnet

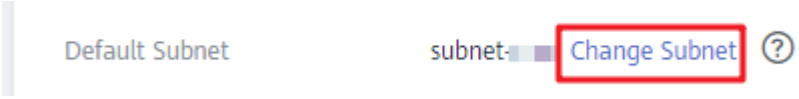

Step 16 Select the target subnet and click **OK**.

Figure 3-17 Selecting the target subnet

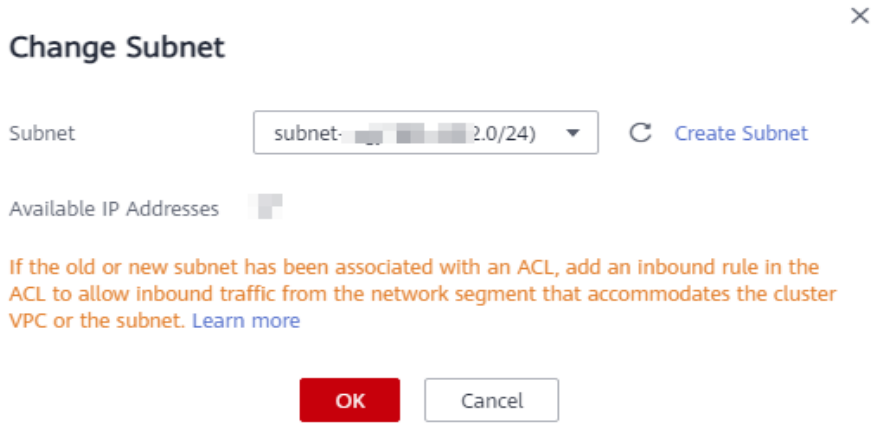

----End

How Do I Configure a Network ACL Outbound Rule?

- Method 1

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

- Method 2

  Allow the mandatory outbound rules that can ensure the successful creation of clusters. You are not advised to use this method because created clusters may not run properly due to absent outbound rules. If the preceding problem occurs, contact O&M personnel.

  Similar to the example provided in method 1, set **Action** to **Allow** and add the outbound rules whose destinations are the address with **Secure Communications** enabled, NTP server address, OBS server address, OpenStack address, and DNS server address, respectively.

3.4.3 Configuring Message Notification

MRS uses SMN to offer a publish/subscribe model to achieve one-to-multiple message subscriptions and notifications in a variety of message types (SMSs and emails).

Scenario

On the MRS management console, you can enable or disable the notification service on the **Alarms** page. The functions in the following scenarios can be implemented only after the required cluster function is enabled:

- After a user subscribes to the notification service, the MRS management plane notifies the user of success or failure of manual cluster scale-out and scale-in, cluster deletion, and auto scaling by emails or SMS messages.
- The management plane checks the alarms about the MRS cluster and sends a notification to the tenant if the alarms are critical.
- If either of the operations such as deletion, shutdown, specifications modification, restart, and OS update is performed on an ECS in a cluster, the

MRS cluster works abnormally. The management plane notifies a user when detecting that the VM of the user is in either of the preceding operations.

Creating a Topic

A topic is a specified event for message publication and notification subscription. It serves as a message sending channel, where publishers and subscribers can interact with each other.

1. Log in to the management console.
2. Click **Service List**. Under **Management & Governance**, click **Simple Message Notification**.

   The **SMN** page is displayed.
3. In the navigation pane, choose **Topic Management** > **Topics**.

   The **Topics** page is displayed.
4. Click **Create Topic**.

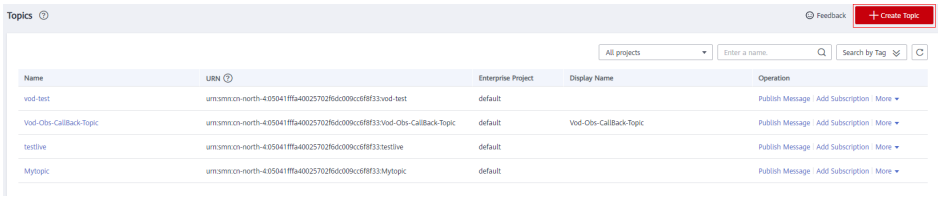

   The **Create Topic** dialog box is displayed.
5. In **Topic Name**, enter a topic name. In **Display Name**, enter a display name.
6. Select an existing project from the **Enterprise Project** drop-down list, or click **Create Enterprise Project** to create an enterprise project on the **Enterprise Project Management** page and then select it.
7. Set tag keys and tag values. Tags consist of keys and values. They identify cloud resources so that you can easily categorize and search for your resources.

Adding Subscriptions to a Topic

To deliver messages published to a topic to subscribers, you must add subscription endpoints to the topic. SMN automatically sends a confirmation message to the subscription endpoint. The confirmation message is valid only within 48 hours. The subscribers must confirm the subscription within 48 hours so that they can receive notification messages. Otherwise, the confirmation message becomes invalid, and you need to send it again.

1. Log in to the management console.
2. Under **Management & Governance**, click **Simple Message Notification**.

   The **SMN** page is displayed.
3. In the navigation pane, choose **Topic Management** > **Topics**.

   The **Topics** page is displayed.
4. Locate the topic to which you want to add a subscription, click **More** in the **Operation** column, and select **Add Subscription**.

   The **Add Subscription** box is displayed.

Protocol can be set to **SMS**, FunctionGraph (function), **HTTP**, **HTTPS**, and **Email**.

Endpoint indicates the address of the subscription endpoint. SMS and email, endpoints can be entered in batches. When adding endpoints in batches, each endpoint address occupies a line. You can enter a maximum of 10 endpoints.

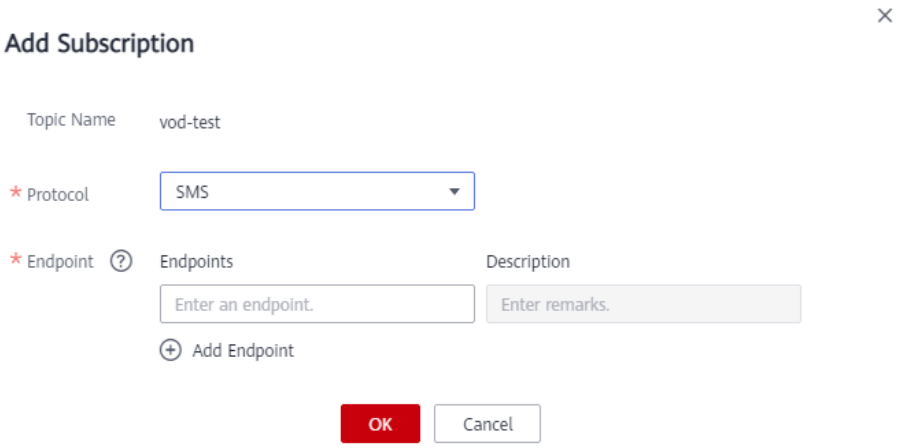

5. Click **OK**.

The subscription you added is displayed in the subscription list.

Sending Notifications to Subscribers

1. Log in to the MRS console.
2. Choose **Clusters > Active Clusters**, select a running cluster, and click its name to switch to the cluster details page.
3. Click **Alarms**.
4. Choose **Notification Rules** > **Add Notification Rule**. The **Add Notification Rule** page is displayed.

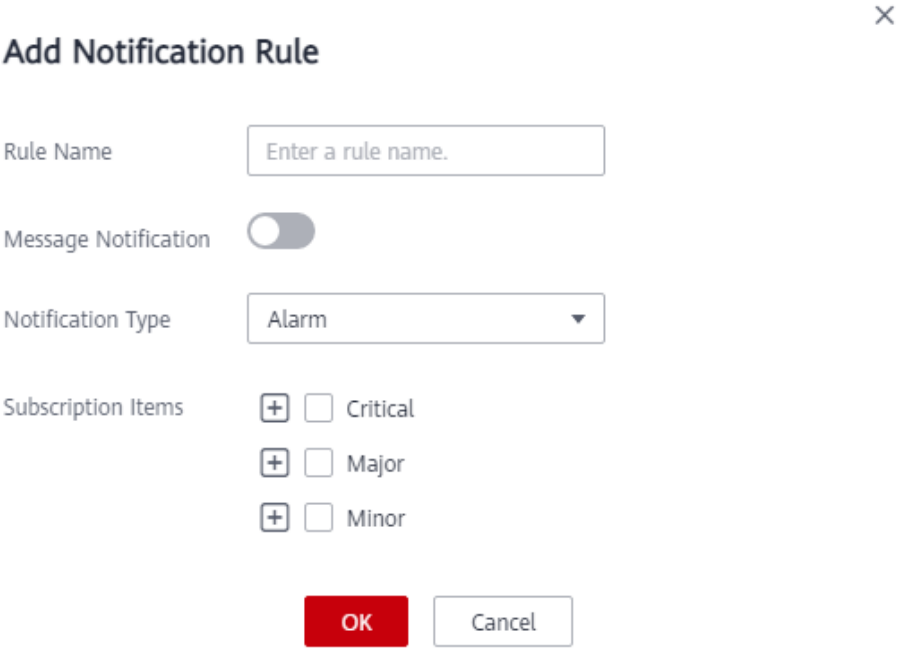

5. Set the notification rule parameters.

Table 3-19 Parameters of a notification rule

Parameter	Description
Rule Name	User-defined notification rule name. Only digits, letters, hyphens (-), and underscores (_) are allowed.
Message Notification	• If you enable this function, the system sends notifications to subscribers based on the notification rule. • If you disable this function, the rule does not take effect, that is, notifications are not sent to subscribers.
Topic Name	Select an existing topic or click **Create Topic** to create a topic.
Notification Type	Select the type of the notification to be subscribed to. • Alarm • Event
Subscription Items	Select the items to be subscribed to. You can select all or some items as required. Subscription rules in MRS 3.*x* or later: Alarm severity: critical, major, and minor Event: major, minor, and warning Subscription rules in versions earlier than MRS 3.x: • Critical • Major • Minor • Suggestion

6. Click **OK**.

NOTE

After a message subscription rule is applied, you may receive some historical alarms.

3.4.4 Checking Health Status

3.4.4.1 Before You Start

This section describes how to manage health checks on the MRS console.

Health check management operations on the MRS console apply only to clusters of **MRS 1.9.2**.

Health check management on Manager applies to all versions. For MRS 3.x and later versions, see **Viewing a Health Check Task**. For versions earlier than MRS 3.x, see **Performing a Health Check**.

3.4.4.2 Performing a Health Check

Scenario

To ensure that cluster parameters, configurations, and monitoring are correct and that the cluster can run stably for a long time, you can perform a health check during routine maintenance.

NOTE

A system health check includes MRS Manager, service-level, and host-level health checks:

- MRS Manager health checks focus on whether the unified management platform can provide management functions.
- Service-level health checks focus on whether components can provide services properly.
- Host-level health checks focus on whether host indicators are normal.

The system health check includes three types of check items: health status, related alarms, and customized monitoring indicators for each check object. The health check results are not always the same as the **Health Status** on the portal.

Procedure

- Manually perform the health check for all services.

  On the MRS details page, choose **Management Operations** > **Start Cluster Health Check**.

  NOTE

  For the operations on MRS Manager, see **Performing a Health Check**; for the operations on FusionInsight Manager of MRS 3.*x* or later, see **Overview**.

  - The cluster health check includes Manager, service, and host status checks.
  - To perform cluster health checks, you can also choose **System** > **Check Health Status** > **Start Cluster Health Check** on MRS Manager.
  - To export the health check result, click **Export Report** in the upper left corner.

- Manually perform the health check for a service.

  a. On the MRS cluster details page, click **Components**.

  b. Select the target service from the service list.

  c. Choose **More** > **Start Service Health Check** to start the health check for the service.

- Manually perform the health check for a host.

  a. On the MRS details page, click **Nodes**.

b. Expand the node group information and select the check box of the host to be checked.

c. Choose **Node** > **Start Host Health Check** to start the health check for the host.

3.4.4.3 Viewing and Exporting a Health Check Report

Scenario

You can view the health check result on MRS and export it for further analysis.

NOTE

A system health check includes MRS Manager, service-level, and host-level health checks:

- MRS Manager health checks focus on whether the unified management platform can provide management functions.
- Service-level health checks focus on whether components can provide services properly.
- Host-level health checks focus on whether host indicators are normal.

The system health check includes three types of check items: health status, related alarms, and customized monitoring indicators for each check object. The health check results are not always the same as the **Health Status** on the portal.

Prerequisites

You have performed a health check.

Procedure

Step 1 On the cluster details page, choose **Management Operations** > **View Cluster Health Check Report**.

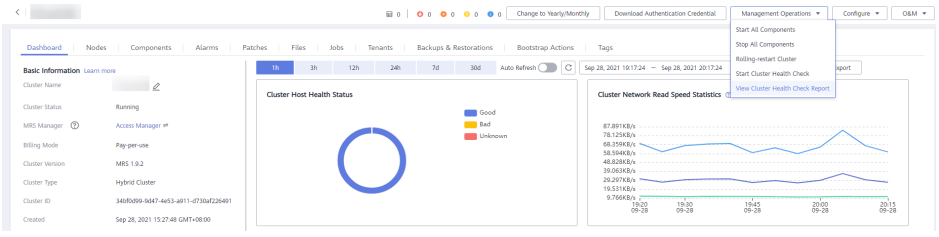

Step 2 Click **Export Report** on the health check report pane to export the report and view detailed information about check items.

----End

3.4.5 Remote O&M

3.4.5.1 Authorizing O&M

If you need HUAWEI CLOUD technical support personnel to help you with troubleshooting, you can use the O&M authorization function to authorize HUAWEI CLOUD technical support personnel to access your local host for fault location.

Procedure

Step 1 Log in to the MRS management console.

Step 2 In the navigation tree of the MRS management console, choose **Clusters** > **Active Clusters**, select a running cluster, and click its name to switch to the cluster details page.

Step 3 In the upper right corner of the page, click **O&M**, choose **Authorize for Cluster Nodes**, and select the deadline for the HUAWEI CLOUD support personnel to access the local host. Before the deadline, the support personnel have the temporary permission to access the local host.

Step 4 After the fault is rectified, click **O&M** in the upper right corner of the page and choose **Cancel Cluster Node Authorization** to retrieve the access permission granted to the support personnel.

----End

3.4.5.2 Sharing Logs

If you need HUAWEI CLOUD technical support personnel to help you with troubleshooting, you can use the log sharing function to provide logs in a specific time to HUAWEI CLOUD technical support personnel for fault location.

Procedure

Step 1 Log in to the MRS management console.

Step 2 In the navigation tree of the MRS management console, choose **Clusters** > **Active Clusters**, select a cluster, and click its name to switch to the cluster details page.

Step 3 In the upper right corner of the displayed page, choose **O&M** > **Share Log** to open the **Share Log** dialog box.

Step 4 Select the start time and end time in **Time Range**.

Figure 3-18 Sharing logs

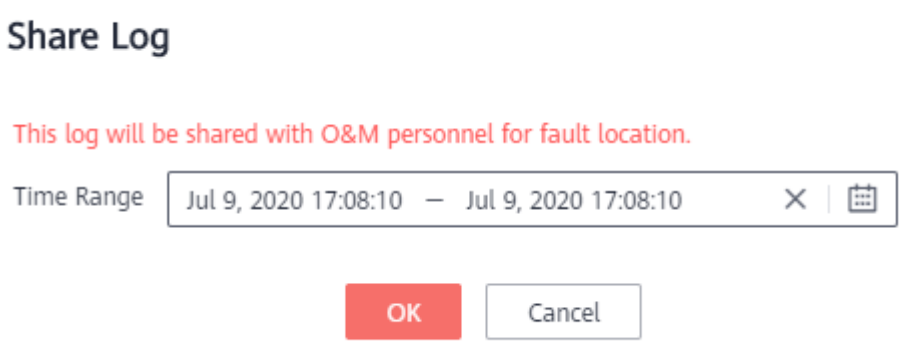

NOTE

- Select **Time Range** based on the suggestions of HUAWEI CLOUD support personnel.
- **End Date** must be later than **Start Date**. Otherwise, logs cannot be filtered by time.

----End

3.4.6 Viewing MRS Operation Logs

You can view operation logs of clusters and jobs on the **Operation Logs** page. Log information is typically used for quickly locating faults in case of cluster exceptions, helping users resolve problems.

Operation Type

Currently, the following operation logs are provided by MRS. You can filter the logs in the search box.

- Cluster operations
  - Creating, deleting, scaling out, and scaling in a cluster
  - Creating and deleting a directory, deleting a file
- Job operations: Creating, stopping, and deleting a job
- Data operations: IAM user tasks, adding user, and adding user group

Log Fields

Logs are listed in chronological order by default in the log list, with the most recent logs displayed at the top.

Table 3-20 describes various fields in a log.

Table 3-20 Log description

Parameter	Description
Operation Type	Various types of operations, including: • Cluster operations • Job operations • Data operations
Operation IP	IP address where an operation is performed. **NOTE** If an MRS cluster fails to be deployed, the cluster is automatically deleted, and the operation logs of the automatically deleted cluster do not contain the **Operation IP** of the user.
Users	The user who performs the operation
Operation	Operation details. The value can contain a maximum of 2048 characters.

Parameter	Description
Time	Operation time. For a deleted cluster, only logs generated within the last six months are displayed. To view logs generated six months ago, contact technical support.
Enterprise Project	Enterprise project to which the cluster belongs

Table 3-21 Icon description

Icon	Description
	Select an enterprise project from the drop-down list box to filter logs.
	Select an operation type from the drop-down list box to filter logs. • **All Operation Types**: Filter all logs. • **Cluster**: Filter logs for **Cluster**. • **Job**: Filter logs for **Job**. • **Data**: Filter logs for **Data**.
	Filter logs by time. 1. Click the input box. 2. Specify the date and time. 3. Click **OK**. The left-side input box indicates the start time and the right-side one indicates the end time. The start time must be earlier than or equal to the end time. Otherwise, logs cannot be filtered.
	Enter a keyword of the **Operation Details** in the search box and click  to search for logs.
	Click  to manually refresh the log list.

3.4.7 Changing Billing Mode to Yearly/Monthly

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

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

Step 1 Log in to the MRS console.

Step 2 In the navigation tree of the MRS management console, choose **Clusters** > **Active Clusters**.

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

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

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

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

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

NOTE

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

----End

3.4.8 Unsubscribing from a Cluster

If a cluster charged in **Yearly/Monthly** mode is not required after the job execution, you can unsubscribe from it. After the cluster is unsubscribed, resources and data will be deleted and cannot be restored. Ensure that the data is backed up before you unsubscribed from the cluster.

For details about the unsubscription rules, see **Conditional Unsubscription**.

Background

Typically after data is analyzed and stored, or when the cluster encounters an exception and cannot work, you can unsubscribe a cluster. When the MRS cluster fails to be deployed, the cluster is automatically unsubscribed.

Procedure

Step 1 Log in to the MRS management console.

Step 2 In the navigation tree of the MRS management console, choose **Clusters** > **Active Clusters**.

Step 3 In the **Operation** column of the cluster from which you want to unsubscribe, click **Unsubscribe**.

Step 4 On the **Unsubscribe** page, confirm the cluster information, select reasons for unsubscription, and confirm the unsubscription amount and related fees.

Step 5 Click **Confirm**.

Step 6 Confirm the unsubscription information and click **Yes** to submit the unsubscription application.

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

----End

3.4.9 Unsubscribing from a Specified Node in a Yearly/Monthly Cluster

You can reduce the number of specific nodes to scale in a cluster so that MRS delivers better storage and computing capabilities at lower O&M costs based on service requirements.

Currently, you can unsubscribe from a maximum of 20 Core nodes at a time, but there must be at least 2 Core nodes available after unsubscription.

NOTE

You can unsubscribe from a node only after the node is successfully isolated or decommissioned. Otherwise, data loss may occur.

Usage Restrictions

- If the number of Core nodes in the cluster is less than or equal to the number of HDFS copies, MRS does not support node unsubscription to ensure data reliability. The number of HDFS copies can be queried using the **dfs.replication** parameter in the HDFS parameter configuration.
- MRS does not support unsubscription from nodes where ZooKeeper, Kudu, Kafka, or ClickHouse is deployed.

How to Unsubscribe from a Specified Node in a Yearly/Monthly Cluster

Step 1 Disable the auto-renewal function of the cluster where the node to be unsubscribed from is located. For details, see **Disabling Auto-Renewal**.

Step 2 Log in to the MRS console.

Step 3 On the **Active Clusters** page, and click the name of the target cluster to be operated. The cluster details page is displayed.

Step 4 On the **Dashboard** page of the cluster, click **Synchronize** on the right of **IAM User Sync**.

Step 5 Unsubscribe from or isolate a node.

NOTE

Currently, only clusters of the following versions support unsubscription from specified nodes in yearly/monthly clusters. For clusters of other versions, contact technical support.

- MRS 2.1.0 (patch 2.1.0.5 or later)
- MRS 3.1.0 (patch 3.1.0.0.2 or later)
- MRS 3.1.5
- MRS 3.2.0-LTS.1 (patch 3.2.0-LTS.1.3 or later)

- If the cluster is earlier than MRS 2.*x*:

a. Click **Isolate Node** in the **Operation** column of the node group to be unsubscribed from.

b. Select the node to be unsubscribed from and click **OK**.

The time required for isolating a node depends on the data volume on the node. A larger data volume indicates a longer time.

After the node is isolated, the node status changes to **Isolated**. The **Unsubscribe from Node** button is displayed on the **Nodes** tab page.

- If the cluster is MRS 3.1.0, 3.1.5, or 3.2.0-LTS.1:

a. Click **Decommission Node** in the **Operation** column of the node group to be unsubscribed from.

b. Select the node to be decommissioned and click **OK**.

The time required for decommissioning a node depends on the data volume on the node. A larger data volume indicates a longer time.

After the node is decommissioned, the node status changes to **Decommissioned**.

NOTE

- For clusters of MRS 3.*x* (except MRS 3.1.0, 3.1.5, and 3.2.0-LTS.1), perform the following operations:
  1. If the DataNode of HDFS, NodeManager of Yarn, or RegionServer of HBase is deployed on the node to be unsubscribed from, log in to FusionInsight Manager and decommission the related instances. For details, see **Decommissioning and Recommissioning an Instance**.
  2. Log in to the MRS management console, on the **Nodes** tab, select the nodes to be unsubscribed from, and choose **Node Operation** > **Isolate Host**. After the host is isolated, contact the technical engineer to unsubscribe from the node.
- Only one node can be isolated or decommissioned at a time. You can unsubscribe from a node only after the node is successfully isolated or decommissioned.
- If the node fails to isolate or decommission, log in to Manager. Click , search for the name of the task that fails to isolate or decommission the host in the task list, click the name, and rectify the fault as prompted.

Step 6 On the cluster details page, choose **Nodes** > **Unsubscribe from Node**.

Step 7 Select the node to be unsubscribed from and click **OK**.

Currently, you can unsubscribe from a maximum of 20 Core nodes at a time, but there must be at least 2 Core nodes available after unsubscription.

Figure 3-19 Selecting a node to be unsubscribed from

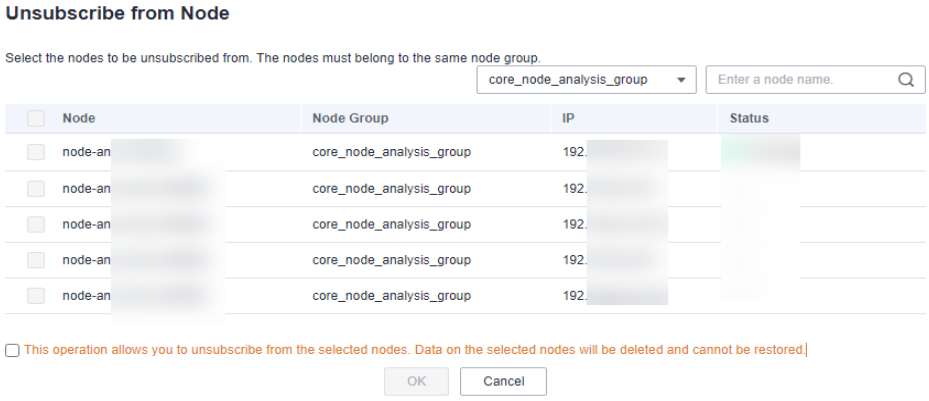

Step 8 On the **Refund Details** page, select "I understand a handling fee will be charged for this unsubscription" as prompted if needed, and click **Confirm**. If this check box does not appear, skip this step.

The cluster status changes to **Scaling in**. After the scale-in is complete, the cluster status changes to **Running**, and the specified node is deleted.

Figure 3-20 Refund details

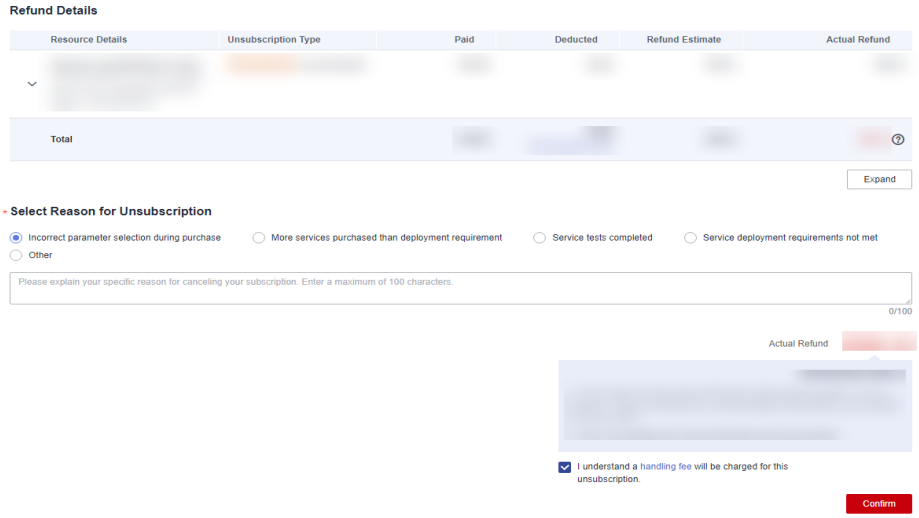

Step 9 (Optional) To enable auto-renewal for a cluster, see **Enabling Auto-Renewal**.

----End

3.4.10 Deleting a Cluster

You can delete an MRS cluster after job execution is complete. The deleted or unsubscribed cluster is no longer billed.

Background

You can manually delete a cluster after data analysis is complete or when the cluster encounters an exception. A cluster failed to be deployed will be automatically deleted.

A yearly/monthly billed cluster cannot not be deleted but can be unsubscribed from. For details about how to unsubscribe from a cluster, see **Unsubscribing from a Cluster**.

Procedure

Step 1 Log in to the MRS management console.

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

Step 3 In the cluster list, locate the row containing the cluster to be deleted, and click **Delete** in the **Operation** column. In the **Delete Cluster** dialog box, enter **DELETE** in the confirmation text box and click **OK**.

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

----End

3.5 Managing Nodes

3.5.1 Scaling Out a Cluster

The storage and computing capabilities of MRS can be improved by simply adding Core nodes or Task nodes instead of modifying system architecture, reducing O&M costs. Core nodes can process and store data. You can add Core nodes to expand the node quantities and handle peak loads. Task nodes are used to process data instead of storing persistent data.

Background

Only Master, Core, and Task nodes can be added.

When you log in to a node added for scale-out as user **root**, the password set during cluster creation is required.

Constraints

- When you expand a node group where HBase is installed:

  If automatic DNS registration is not enabled for a node in the cluster, do not start HBase when you expand the node group. Then, update the HBase client configuration by referring to **Updating a Client** and start the HBase instances on the node to be expanded.

  In the following versions, automatic DNS registration is enabled by default:

  MRS 1.9.3, MRS 3.1.0, MRS 3.1.2-LTS, MRS 3.1.5, and MRS 3.2.0-LTS

  You can query the metadata of a version to check whether the **features** field in the response body contains **register_dns_server**.
- After a scale-out, the clients installed on nodes in the cluster do not need to be updated. For details about how to update the client installed on nodes outside the cluster, see **Updating a Client**.

- If you need to balance HDFS data after scale-out, see **Balancing DataNode Capacity**. For details about how to balance Kafka data, see **Kafka Balancing Tool Instructions**.

Scaling Out a Cluster Billed in Pay-per-Use Mode

Step 1 Log in to the MRS console.

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

Step 3 Click the **Nodes** tab. In the **Operation** column of the node group, click **Scale Out**. The **Scale Out** page is displayed.

The scale-out operation can only be performed on the running clusters.

Step 4 Set the type of **System Disk** and **Data Disk**, **Scale-Out Nodes**, **Enable Components** and **Run Bootstrap Action**, and click **OK**. The **Enable Components** and **Run Bootstrap Action** parameters may not be supported by clusters of some versions. The operations are subject to the UI.

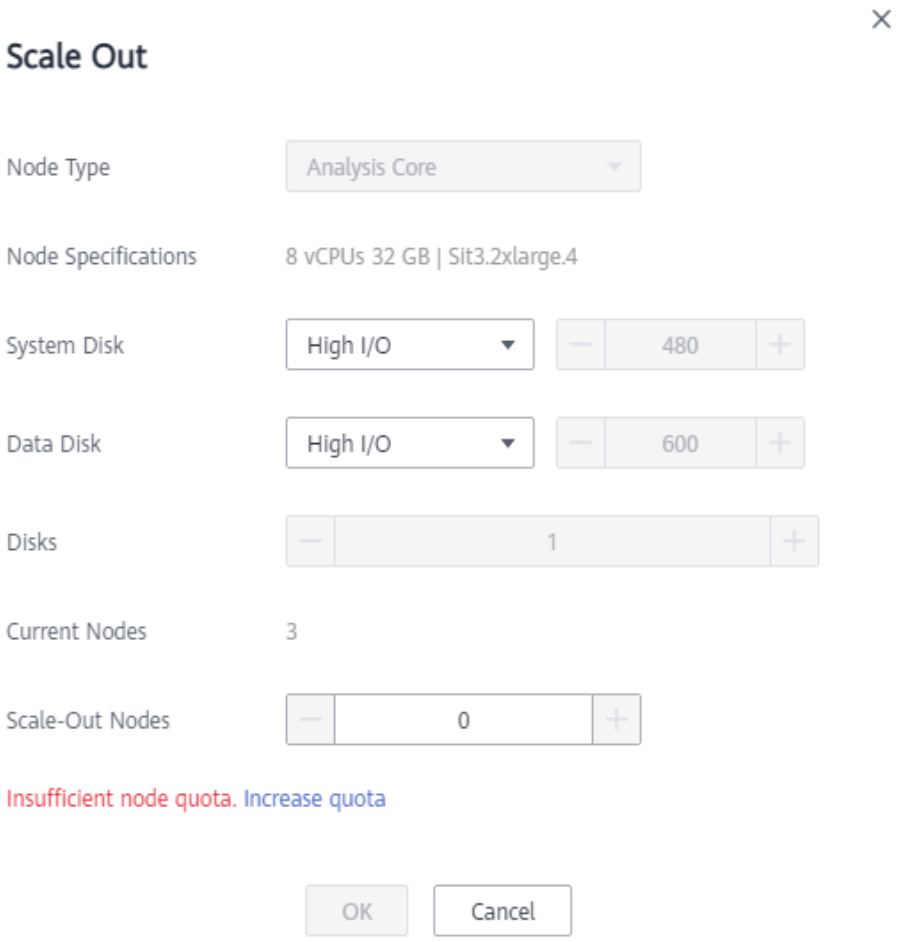

NOTE

- If the Task node group does not exist in the cluster, configure the Task node by referring to **Related Tasks**.
- If a bootstrap action is added during cluster creation, the **Run Bootstrap Action** parameter is valid. If this function is enabled, the bootstrap actions added during cluster creation will be run on all the scaled out nodes.
- If the **New Specifications** parameter is available, the specifications that are the same as those of the original nodes have been sold out or discontinued. Nodes with new specifications will be added.
- Before scaling out the cluster, check whether its security group configuration is correct. Ensure that an inbound security group rule contains a rule in which **Protocol & Port** is set to **All**, and **Source** is set to a trusted accessible IP address range.

Step 5 In the **Add Node** dialog box, click **OK**.

Step 6 A dialog box is displayed in the upper right corner of the page, indicating that the scale-out task is submitted successfully.

The following parameters explain the cluster scale-out process:

- During scale-out: If a cluster is being scaled out, its status is **Scaling out**. The submitted jobs will be executed and you can submit new jobs. You are not allowed to continue to scale out or delete the cluster. You are advised not to restart the cluster or modify the cluster configuration.
- Successful scale-out: The cluster status is **Running**. The resources used in the old nodes and expanded nodes are charged.
- Failed scale-out: The cluster status is **Running**. You can execute jobs and scale out the cluster again.

After the cluster is scaled out, you can view the node information of the cluster on the **Nodes** page.

----End

Scaling Out a Cluster Billed in Yearly/Monthly Mode

Step 1 Log in to the MRS console.

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

Step 3 Click the **Nodes** tab. In the **Operation** column of the node group, click **Scale Out**. The **Scale Out** page is displayed.

The scale-out operation can only be performed on the running clusters.

Step 4 Set the type of the **System Disk** and **Data Disk**, **Scale-Out Nodes**, **Enable Components** and **Run Bootstrap Action**. The system displays the expiration time of the cluster and the fee required for adding nodes. **Enable Components** and **Run Bootstrap Action** may not be supported by clusters of some versions. The operations are subject to the UI.

NOTE

- If a bootstrap action is added during cluster creation, the **Run Bootstrap Action** parameter is valid. If this function is enabled, the bootstrap actions added during cluster creation will be run on all the scaled out nodes.
- If the **New Specifications** parameter is available, the specifications that are the same as those of the original nodes have been sold out or discontinued. Nodes with new specifications will be added.
- Before scaling out the cluster, check whether its security group configuration is correct. Ensure that an inbound security group rule contains a rule in which **Protocol & Port** is set to **All**, and **Source** is set to a trusted accessible IP address range.

- Click **Submit Order**.

On the **Purchase MapReduce Service** page, click **Pay**.

- Click **Confirm order, not pay**.

On the cluster information page, choose **Fee > My Order** and click **Pay**.

Step 5 After the payment is successful, return to the MRS management console to view the cluster status.

The following parameters explain the cluster scale-out process:

- During scale-out: If a cluster is being scaled out, its status is **Scaling out**. The submitted jobs will be executed and you can submit new jobs. You are not allowed to continue to scale out or delete the cluster. You are advised not to restart the cluster or modify the cluster configuration.
- Successful scale-out: The cluster status is **Running**. The resources used in the old nodes and expanded nodes are charged.
- Failed scale-out: The cluster status is **Running**. You can execute jobs and scale out the cluster again.

After the cluster is scaled out, you can view the node information of the cluster on the **Nodes** page.

----End

Adding a Task Node

You can scale out an MRS cluster by manually adding task nodes.

To add a task node to a custom cluster, perform the following steps:

1. On the cluster details page, click the **Nodes** tab and click **Add Node Group**. The **Add Node Group** page is displayed.
2. Select **Task** for**Node Type**. Retain the default value **NM** for **Deploy Roles**. To deploy the NodeManager role, the node type must be **Task**. Set other parameters as required.

Figure 3-21 Adding a task node group

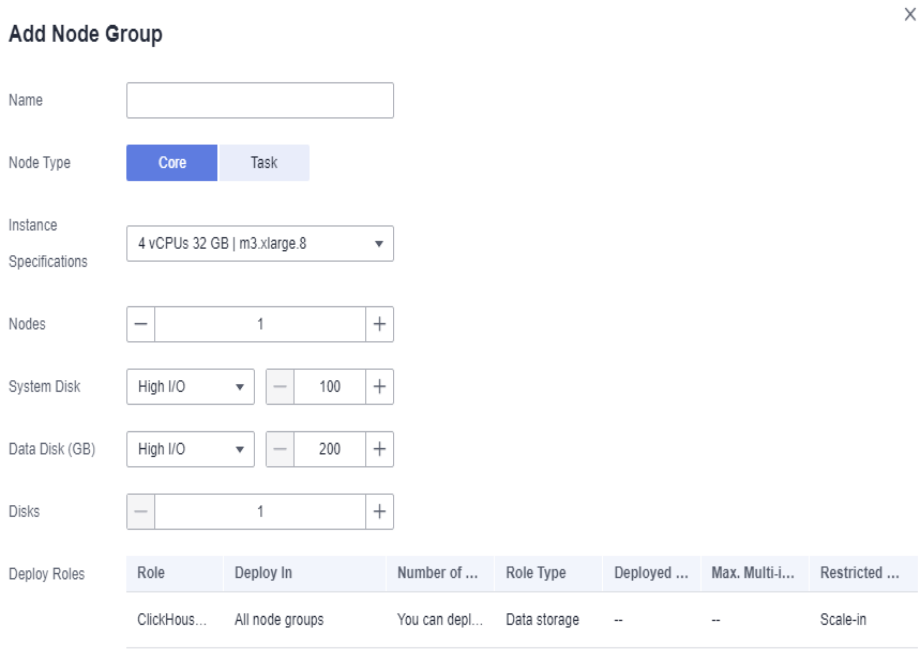

To add a task node to a non-custom cluster, perform the following steps:

1. On the cluster details page, click the **Nodes** tab and click **Configure Task Node**. The **Configure Task Node** page is displayed.
2. On the **Configure Task Node** page, set **Node Type**, **Instance Specifications**, **Nodes**, **System Disk**. In addition, if **Add Data Disk** is enabled, configure the storage type, size, and number of data disks.

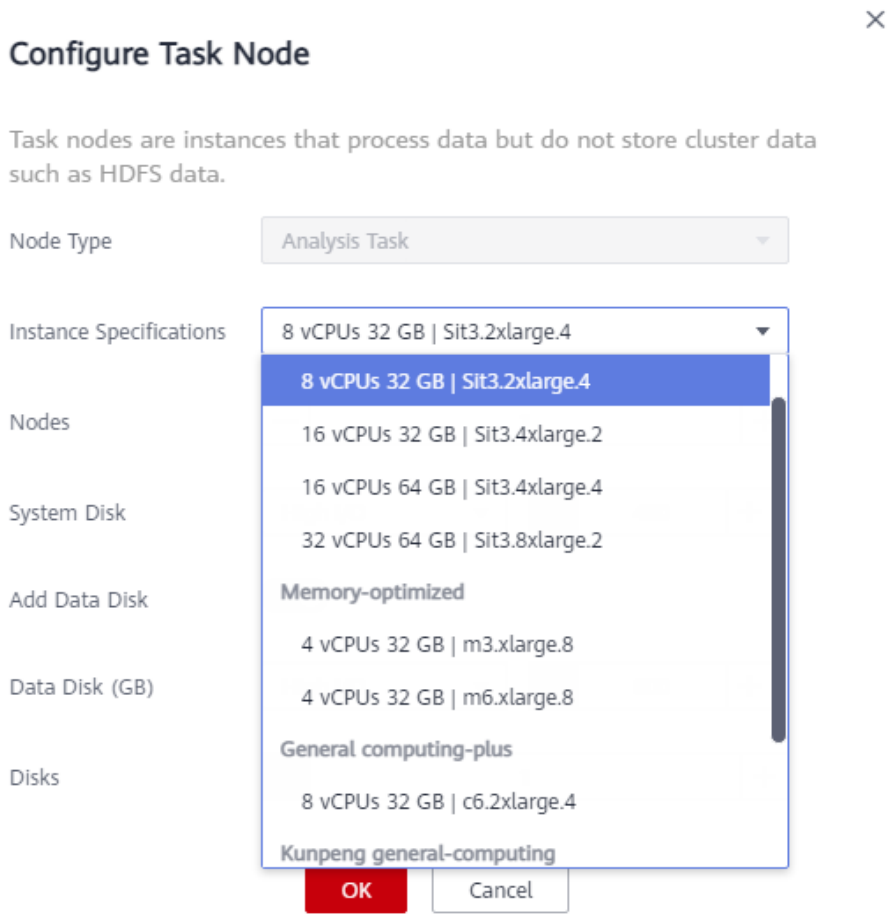

3. Click **OK**.

Adding a Node Group

NOTE

Used to add node groups and applies to customized clusters of MRS 3.*x*.

1. On the cluster details page, click the **Nodes** tab and click **Add Node Group**. The **Add Node Group** page is displayed.

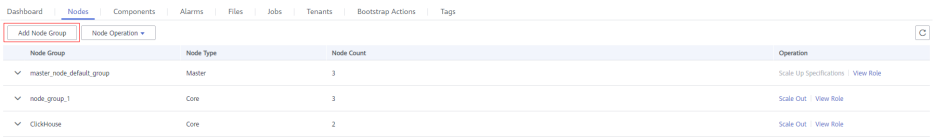

2. Set the parameters as needed.
3. Click **OK**.

3.5.2 Scaling In a Cluster

You can reduce the number of core or task nodes to scale in a cluster based on service requirements so that MRS delivers better storage and computing capabilities at lower O&M costs.

The scale-in operation is not allowed for a cluster that is performing active/standby synchronization.

NOTE

Only pay-per-use clusters can be scaled in. For details about how to scale in a yearly/monthly node, see **Unsubscribing from a Specified Node in a Yearly/Monthly Cluster**.

Background

A cluster can have three types of nodes, master, core, and task nodes. Currently, only core and task nodes can be removed. To scale in a cluster, you only need to adjust the number of nodes on the MRS console. MRS then automatically selects the nodes to be removed.

The policies for MRS to automatically select nodes are as follows:

- MRS does not select the nodes with basic components installed, such as ZooKeeper, DBService, KrbServer, and LdapServer, because these basic components are the basis for the cluster to run.
- Core nodes store cluster service data. When scaling in a cluster, ensure that all data on the core nodes to be removed has been migrated to other nodes. You can perform follow-up scale-in operations only after all component services are decommissioned, for example, removing nodes from Manager and deleting ECSs. When selecting core nodes, MRS preferentially selects the nodes with a small amount of data and healthy instances to be decommissioned to prevent decommissioning failures. For example, if DataNodes are installed on core nodes in an analysis cluster, MRS preferentially selects the nodes with small data volume and good health status during scale-in.

  When core nodes are removed, their data is migrated to other nodes. If the user business has cached the data storage path, the client will automatically update the path, which may increase the service processing latency temporarily. Cluster scale-in may slow the response of the first access to some HBase on HDFS data. You can restart HBase or disable or enable related tables to resolve this issue.
- Task nodes are computing nodes and do not store cluster data. Data migration is not involved in removing task nodes. Therefore, when selecting task nodes, MRS preferentially selects nodes whose health status is faulty, unknown, or subhealthy. On the **Components** tab of the MRS console, click a service and then the **Instances** tab to view the health status of the node instances.

Scale-In Verification Policy

To prevent component decommissioning failures, components provide different decommissioning constraints. Scale-in is allowed only when the constraints of all installed components are met. **Table 3-22** describes the scale-in verification policies.

Table 3-22 Decommissioning constraints

Component	Constraint
HDFS/DataNode	The number of available nodes after the scale-in is greater than or equal to the number of HDFS copies and the total HDFS data volume does not exceed 80% of the total HDFS cluster capacity. This ensures that the remaining space is sufficient for storing existing data after the scale-in and reserves some space for future use. **NOTE** To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total.
HBase/ RegionServer	The total available memory of RegionServers on all nodes except the nodes to be removed is greater than 1.2 times of the memory which is currently used by RegionServers on these nodes. This ensures that the node to which the region on a decommissioned node is migrated has sufficient memory to bear the region of the decommissioned node.
Storm/ Supervisor	After the scale-in, ensure that the number of slots in the cluster is sufficient for running the submitted tasks. This prevents no sufficient resources being available for running the stream processing tasks after the scale-in.
Flume/ FlumeServer	If FlumeServer is installed on a node and Flume tasks have been configured for the node, the node cannot be deleted. This prevents the deployed service program from being deleted by mistake.
ClickHouse/ ClickHouseServer	For details, see **Constraints on ClickHouseServer Scale-in**. This ensures that data on the decommissioned nodes is migrated to in-use nodes.

Scaling In a Cluster by Specifying the Node Quantity

Step 1 Log in to the MRS console.

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

Step 3 Click the **Nodes** tab. In the **Operation** column of the node group, click **Scale In** to go to the **Scale In** page.

This operation can be performed only when the cluster and all nodes in it are running.

Step 4 Set **Scale-In Type** to **Node quantity**.

Step 5 Set **Scale-In Nodes** and click **OK**.

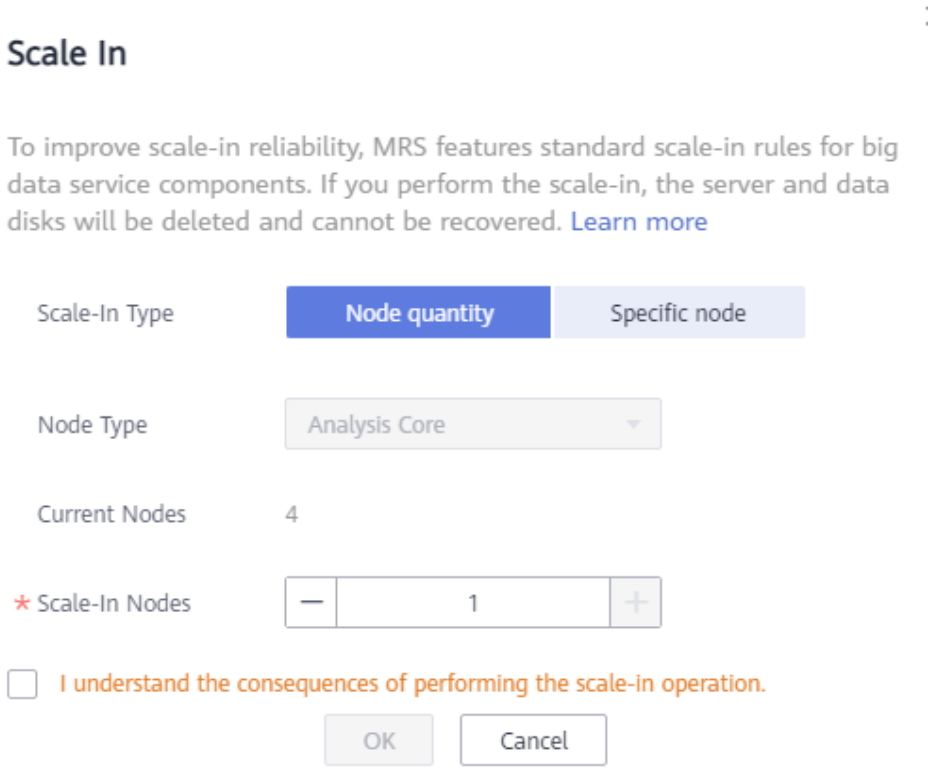

NOTE

- Before scaling in the cluster, check whether its security group configuration is correct. Ensure that an inbound security group rule contains a rule in which **Protocol & Port** is set to **All**, and **Source** is set to a trusted accessible IP address range.
- If damaged data blocks exist in HDFS, the cluster may fail to be scaled in. Contact Huawei Cloud technical support.

Step 6 A dialog box displayed in the upper right corner of the page indicates that the task of removing the node is submitted successfully.

The cluster scale-in process is explained as follows:

- During scale-in: The cluster status is **Scaling In**. The submitted jobs will be executed, and you can submit new jobs. You are not allowed to continue to scale in or delete the cluster. You are advised not to restart the cluster or modify the cluster configuration.
- Successful scale-in: The cluster status is **Running**. The resources used after the cluster scale-in are billed.
- Failed scale-in: The cluster status is **Running**. You can execute jobs or scale-in the cluster again.

After the cluster is scaled in, you can view the node information of the cluster on the **Nodes** page.

----End

Scaling In a Cluster by Removing Nodes that Are No Longer Needed

If a faulty node is no longer needed, you can use this function to remove it. When the node is removed, the instance of the component role will not be

decommissioned. Before deleting the node, ensure that the data on the node has been backed up. For details about how to remove ClickHouseServer nodes, see **Removing ClickHouseServer Instance Nodes**.

Step 1 Log in to MRS Manager and choose **Hosts**.

Step 2 Select the host to be removed, choose **More**, and select **Isolate** to isolate the host.

For versions earlier than MRS 3.x, isolate the node to be removed by referring to **Isolating a Host**.

The time required for isolating a host depends on the data volume on the host. A larger data volume requires a longer time.

After the node is isolated, the node status changes to **Isolated**.

NOTE

- If the host isolation fails, log in to MRS Manager, click  to search for the task that fails to isolate the host in the task list, and rectify the fault as prompted.
- Isolating a host helps you decommission a node. If data on the node has been backed up, you can skip the operation of isolating a host, directly stop the host on the ECS console, and scale in the host.
- If a host is faulty, forcibly remove the node.

Step 3 Log in to the MRS console.

Step 4 Click the name of the cluster to go to its details page.

Step 5 Click the **Nodes** tab.

Step 6 Locate the row that contains the target node group and click **Scale In** in the **Operation** column to go to the **Scale In** page.

Step 7 Set **Scale-In Type** to **Specific node** and select the node to be removed.

Nodes in the **Stopped**, **Lost**, **Unknown**, **Isolated**, or **Faulty** status can be specified for scale-in. If the node cannot be selected, click **Stop ECS** to go to the ECS console to stop the node. On the cluster details page of the MRS console, click the **Alarms** tab and check whether any service fault alarms are generated after the node is stopped. If no such an alarm is generated, go back to the **Scale In** page and select the corresponding node for scale-in. If such an alarm is generated, clear the alarm before the scale-in.

Step 8 Select **I understand the consequences of performing the scale-in operation**, and click **OK**.

Step 9 Click the **Components** tab and check whether each component is normal. If any component is abnormal, wait for 5 to 10 minutes and check the component status again. If the fault persists, contact technical support.

Step 10 Click the **Alarms** tab and check whether there are exception alarms. If there are exception alarms, clear them before performing other operations.

----End

3.5.3 Removing ClickHouseServer Instance Nodes

3.5.3.1 Constraints on ClickHouseServer Scale-in

Cluster Scale

- If a cluster has only one shard, the instance nodes cannot be removed from the cluster.
- Multiple instance nodes in the same shard **must be decommissioned or recommissioned at the same time**.

  To query cluster shard information, perform the following steps:

  a. Log in to the node where the client is installed as the client installation user.

  cd *Client installation directory*

  source bigdata_env

  In security mode, run the following commands:

  kinit *ClickHouse service user*

  clickhouse client --host *IP address of the ClickHouse instance* **--port 9440 --secure**

  In normal mode, run the following command:

  clickhouse client --host *IP address of the ClickHouse instance* **--user** *Username* **--password --port** *9000*

  Enter the password.

  b. Run the following command to query the cluster shard information:

  select cluster,shard_num,replica_num,host_name from system.clusters;

Cluster Storage

Ensure that the disk space of nodes that will not be decommissioned is sufficient for storing data of all decommissioned nodes. There must be approximately 10% redundant storage space after decommissioning to ensure that the remaining instances can run properly. The procedure is as follows:

1. Run the following command to check the disk usage on each node:

   select * from system.disks;

   free_space indicates the free disk space, and **total_space** indicates the total disk space. The used space is calculated by subtracting the value of **free_space** from that of **total_space**, and its unit is byte.

2. Run the preceding command on a node you want to decommission and calculate the data volume on the node using the preceding formula.
3. Run the preceding command on a node that will not be decommissioned, and then use the following formula: (Value of **free_space** – Data volume of the node to be decommissioned)/Value of **total_space**. If the result is greater than 10%, the node can be decommissioned.

Cluster Status

If there is any faulty ClickHouseServer instance node in the cluster, all instance nodes in the cluster cannot be decommissioned. Log in to Manager, choose **Cluster** > **Services** > **ClickHouse**, click **Instance**, and view the running status of each node in the cluster.

Database

If a database is deployed only on an instance node you want to decommission, the instance node cannot be decommissioned. To remove the instance node, you need to create the database on all ClickHouseServer instance nodes in the cluster. The procedure is as follows:

1. Run the **select * from system.databases;** command to collect the database list of each node.

   name indicates the database name. **engine** indicates the database engine, and the default value is **Atomic**. If the default engine is used, you do not need to specify the engine when creating a table.

2. For the database deployed only on the instance node to be decommissioned, run the following command to create the database:

   create database xxx engine=xxx on cluster xxx;

Local Non-replicated Table

If a local non-replicated table is deployed only on an instance node you want to decommission, the instance node cannot be decommissioned. To decommission the node, create a local non-replicated table with the same name on any node that will not be decommissioned.

For example, the current cluster has two shards, shard 1 has two nodes A and B, and shard 2 has two nodes C and D. The non-replicated table **test** was created without the **ON CLUSTER** keyword, so the table is created only on node A.

In this case, to decommission nodes A and B in shard 1, you need to create the table **test** on node C or D in shard 2.

Run the following command to list the data tables of each node:

select database,name,engine,create_table_query from system.tables where database != 'system';

Perform the following operations according to the result:

- Check the **engine** column. The table that does not contain the **Replicated** field is a local non-replicated table.
- If there are no replicated tables on any nodes that will not be decommissioned, create one based the table created by **create_table_query**. The following creation statement is an example:

  CREATE TABLE {database}.{table} ('column name' type...) ENGINE = MergeTree;

Replicated Table

If a replicated table exists only on some nodes in the cluster, the nodes where the replicated table is deployed cannot be decommissioned. You need to manually

create the replicated table on all instance nodes where no replicated table is deployed in the cluster before decommissioning.

For example, the current cluster has two shards, shard 1 has two nodes A and B, and shard 2 has two nodes C and D. The replicated table **test** was created without the **ON CLUSTER** keyword, so the table is created only on nodes A and B.

To decommission nodes A and B in shard 1, you need to create the table **test** on nodes C and D in shard 2.

Run the following command to list the data tables of each node:

select database,name,engine,create_table_query from system.tables where database != 'system';

Perform the following operations according to the result:

- Check the **engine** column. The table that contains the **Replicated** field is a replicated table.
- If there are no replicated tables on any nodes that will not be decommissioned, create one based the table created by **create_table_query**.

Distributed Table

Distributed tables will not be migrated automatically for decommissioning. Create distributed tables on the nodes that will not be decommissioned.

Run the following command to list data tables of each node and check the **engine** column. These tables are distributed tables if this column contains field **Distributed**.

select database,name,engine from system.tables where database != 'system';

NOTE

Creating distributed tables on these nodes will not affect the decommissioning, but may affect subsequent service operations.

View

Views will not be automatically migrated for decommissioning, and views do not store data. Run the following command to list data tables of each node and check the **engine** column. These tables are views if this column contains field **View**.

select database,name,engine from system.tables where database != 'system';

Run the following command to delete the views one by one:

drop view {database_name}.{table_name};

Materialized Views

Materialized views will not be automatically migrated for Decommissioning. Create materialized views on the nodes that will not be decommissioned. If the materialized view of a node to be decommissioned does not display the specified aggregation table but uses an embedded table, the node cannot be decommissioned.

Run the following command to list data tables of each node and check the **engine** column. These tables are materialized views if this column contains field **MaterializedView**.

select database,name,engine, create_table_query from system.tables where database != 'system';

The table whose **create_table_query** column contains the **POPULATE** field is an embedded table. Views are initialized when they are created, and newly inserted data is ignored during the initialization. A table that does not contain the **POPULATE** field is an aggregation table. Newly inserted data is directly inserted into the view charts and support tables, and the original data is manually loaded into the views and support tables. The table creation operations of the aggregation table and embedded table are different.

Perform the following operations to process the materialized views of the node to be decommissioned:

1. Record the materialized views and delete them.

   drop view {database_name}.{table_name};

2. After the node decommissioning is complete, delete and recreate the corresponding materialized views on in-use nodes to update the materialized views.
3. To create an aggregation table, specify **WHERE** to search for historical data and manually import the historical data to the materialized views. Otherwise, historical data cannot be imported to the materialized views based on unified conditions. As a result, data is imported repeatedly. For example, an update point can be specified to ensure that data before the update point is manually loaded in **INSERT** mode.
   - Add **WHERE** *{ Time field (for example, date)}***>= toDate (***{ Current time (for example, '2022-12-01 00:00:00')}***)** to the table creation statement.
   - **insert into** *{table}* **select** *{Table field}* **from** *{Source table}* **where** *{Time field}***< toDate** *({Current time})* is used to load original data.
4. Embedded tables will lose data generated during table creation. You can specify **WHERE** to filter out all historical data. In this case, an empty table is created, and you only need to manually insert all data in the historical data source table.

Tables of Third-Party Engines

Currently, tables of third-party engines cannot be automatically migrated for decommissioning.

Run the following command to list data tables of each node and check the **engine** column. These tables are tables of third-party engines if this column does not contain any of the following fields: **MergeTree**, **View**, **MaterializedView**, **Distributed**, and **Log**. (The **engine** column of a third-party engine table may contain field **Memory**, **HDFS**, or **MySQL**.)

select database,name,engine from system.tables where database != 'system';

Create third-party engine tables on the nodes that will not be decommissioned and delete those from the nodes that will be decommissioned.

Detached Data

If the table on a node to be decommissioned has been detached and data still exists in the **detached** directory, the node cannot be decommissioned. You need to attach the data in the **detached** directory to other directories before decommissioning.

1. Run the following command to view the **system.detached_parts** system catalog of the node to be decommissioned:

   select * from system.detached_parts;

2. If **detached part** data exists and these partitions are no longer used, run the following command to delete the **detached part** data:

   ALTER TABLE {table_name} DROP DETACHED PARTITION {partition_expr} SETTINGS allow_drop_detached = 1;

3. Run the following command to check whether there is any **detached part** data in the **system.detached_parts** system catalog:

   select * from system.detached_parts;

   If the command output is empty, there is no **detached part** data in this system catalog.

3.5.3.2 Scaling In ClickHouseServer Nodes

Before removing ClickHouseServer instance nodes, you need to decommission them. Multiple node replicas of the same shard **must be decommissioned at the same time**. If there is a faulty ClickHouseServer instance node in the cluster, all instance nodes of the cluster cannot be decommissioned. For more constraints, see **Constraints on ClickHouseServer Scale-in**.

NOTE

- Perform the decommissioning in idle hours because the operation will occupy certain bandwidth resources.
- The decommissioning operation can be performed only to ClickHouseServer. ClickHouseBalancer cannot be decommissioned.
- **This operation is only supported for MRS 3.1.2 and later.**

Step 1 Use PuTTY to log in to the node where ClickHouseServer is installed as user **root** and run the following command:

echo 'select * from system.clusters' | curl -k 'https://*IP address of the node where the ClickHouseServer instance is located***:***Port number***/' -u** *ck_user*:*Password* **--data-binary @-**

Record the nodes of the same shard. In the following command output, the nodes with the same number in bold belong to the same shard.

```
[root@kwephispra44948 ~]# echo 'select * from  system.clusters' | curl -k 'https://10.112.17.189:21422/' -u
ck_user:Bigdata_2013  --data-binary @-
default_cluster 1       1       1       kwephispra44947 10.112.17.150  21427   0                       0       0
default_cluster 1       1       2       kwephispra44948 10.112.17.189  21427   0                       0       0
```


NOTE

- To view the port number of ClickHouseServer instance nodes, log in to FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**, click **Configuration** > **All Configurations**, and choose **ClickHouseServer (Role)** on the left.

  In security mode (Kerberos authentication enabled), check the value of **https_port**, which is the port of a ClickHouseServer instance node.

  In common mode (Kerberos authentication disabled), check the value of **http_port**, which is the port of a ClickHouseServer instance node.
- **ck_user** indicates the created ClickHouse user, which must be bound to a role with the ClickHouse administrator permission. For details about how to create a user and a role, see **Creating a User** and **Managing Roles**, respectively.

Step 2 Log in to the MRS console and click the cluster name to go to the cluster details page.

Step 3 Click the **Components** tab and click **ClickHouse**. Then switch to **Instances**, select the **ClickHouseServer** instances to be removed, click **More**, and select **Decommission**.

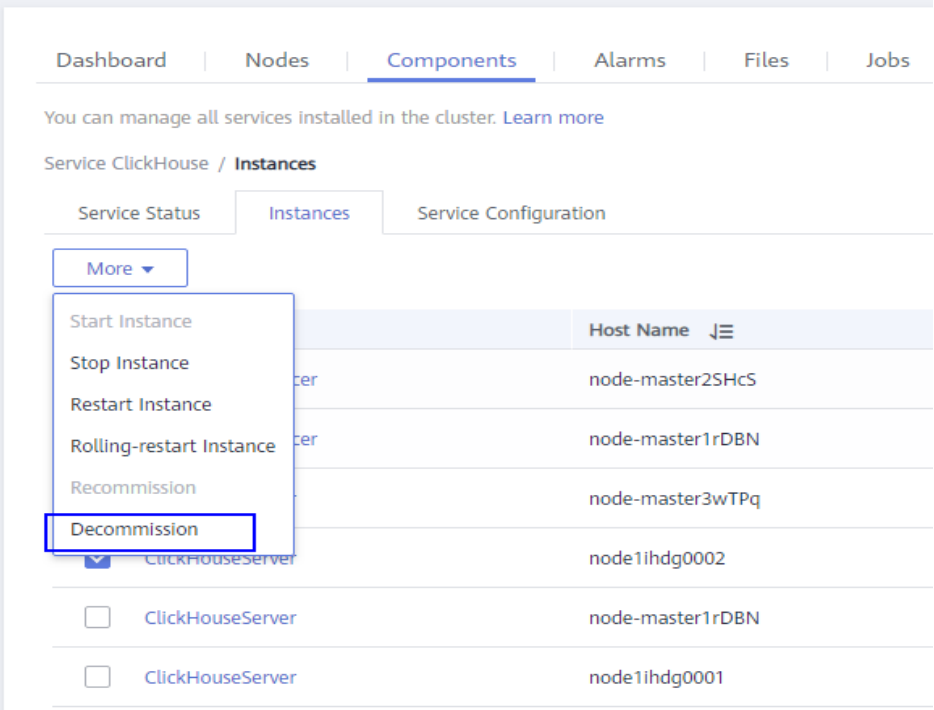

Step 4 Click the **Components** tab and click **ClickHouse**. Then click **More**, and select **Synchronize Configuration**.

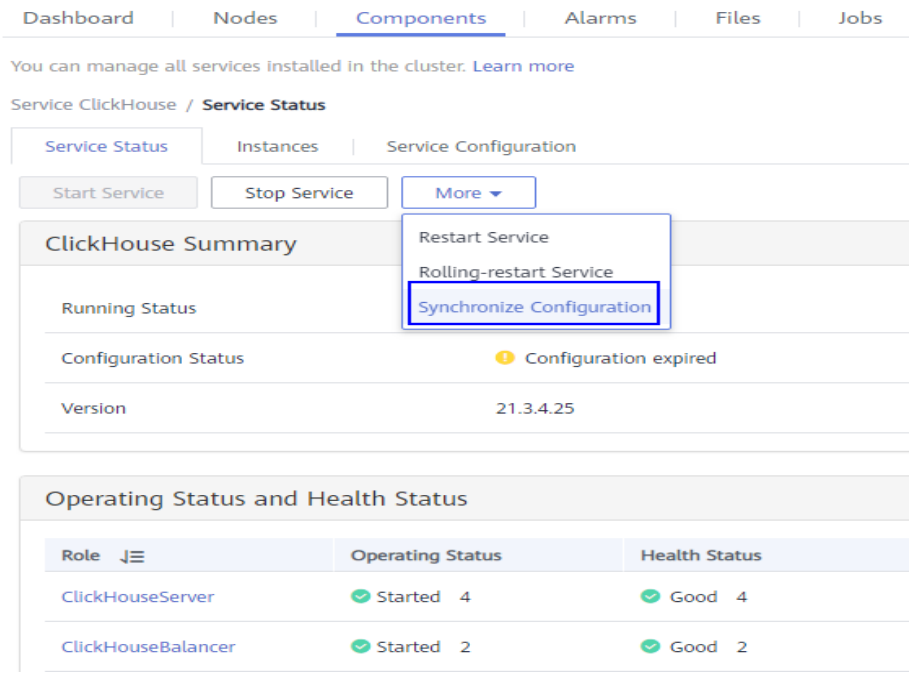

Step 5 Click the **Nodes** tab and click the ClickHouseServer instance node that has been decommissioned.

Step 6 On the ECS page, click **Stop**. In the displayed dialog box, select **Forcibly stop the preceding ECSs** and click **Yes**.

Step 7 Go back to the MRS console, click the **Nodes** tab, locate the row that contains the target node group, and click **Scale In** in the **Operation** column to go to the **Scale In** page.

Step 8 Set **Scale-In Type** to **Specific node** and select the node to be removed.

Step 9 Select **I understand the consequences of performing the scale-in operation**. Click **OK**.

Step 10 Click the **Components** tab and check whether each component is normal. If any component is abnormal, wait for 5 to 10 minutes and check the component status again. If the fault persists, contact technical support.

Step 11 Click the **Alarms** tab and check whether there are exception alarms. If there are exception alarms, clear them before performing other operations.

----End

3.5.4 Managing a Host (Node)

Scenario

To check an abnormal or faulty host (node), you need to stop all host roles on MRS. To recover host services after the host fault is rectified, restart all roles.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-22 Synchronizing IAM users (using MRS 1.9.2 as an example)

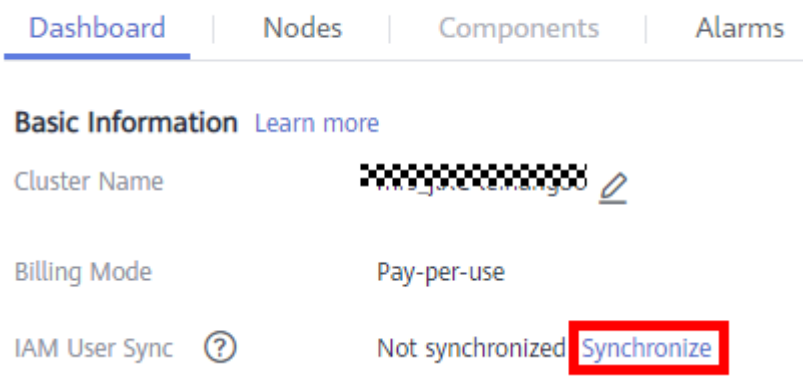

Procedure

Step 1 On the MRS details page, click **Nodes**.

Step 2 Unfold the node group information and select the check box of the target node.

Step 3 Choose **Node Operation** > **Start All Roles** or **Stop All Roles** to perform the required operation.

----End

3.5.5 Isolating a Host

Scenario

If a host is found to be abnormal or faulty, affecting cluster performance or preventing services from being provided, you can temporarily exclude that host from the available nodes in the cluster. In this way, the client can access other available nodes. In scenarios where patches are to be installed in a cluster, you can also exclude a specified node from patch installation.

You can isolate a host manually on MRS based on the actual service requirements or O&M plan. Only non-management nodes can be isolated.

Impact on the System

- After a host is isolated, all role instances on the host will be stopped. You cannot start, stop, or configure the host and any instances on the host.
- After a host is isolated, statistics of the monitoring status and indicator data of the host hardware and instances cannot be collected or displayed.

Prerequisites

You have synchronized IAM users. (To synchronize IAM users, on the **Dashboard** tab page, click **Synchronize** next to **IAM User Sync**.)

Figure 3-23 Synchronizing IAM users (using MRS 1.9.2 as an example)

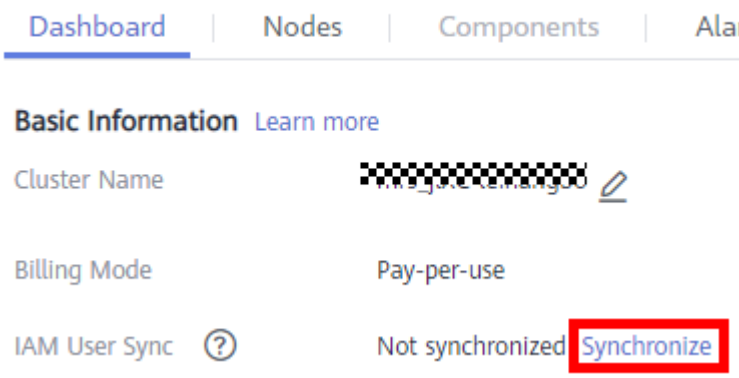

Procedure

Step 1 On the MRS details page, click **Nodes**.

Step 2 Unfold the node group information and select the check box of the target host.

Step 3 Choose **Node Operation** > **Isolate Host**.

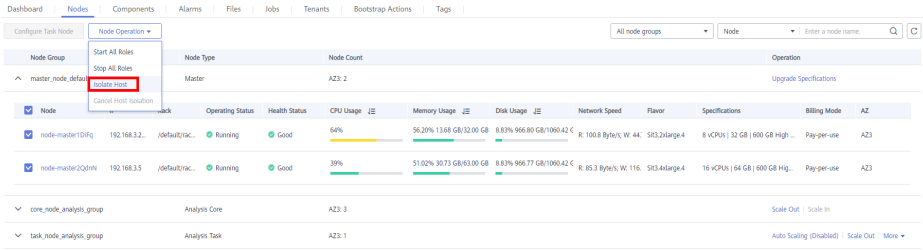

Step 4 Confirm the information about the host to be isolated and click **OK**.

When **Operation successful** is displayed, click **Finish**. The host is isolated successfully, and the value of **Operating Status** becomes **Isolated**.

NOTE

For isolated hosts, you can cancel the isolation and add them to the cluster again. For details, see **Canceling Host Isolation**.

----End

3.5.6 Canceling Host Isolation

Scenario

After the exception or fault of a host is handled, you must cancel the isolation of the host for proper usage.

You can cancel the isolation of a host on MRS.

Prerequisites

- The host is in the **Isolated** state.
- The exception or fault of the host has been rectified.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

  Figure 3-24 Synchronizing IAM users (using MRS 1.9.2 as an example)

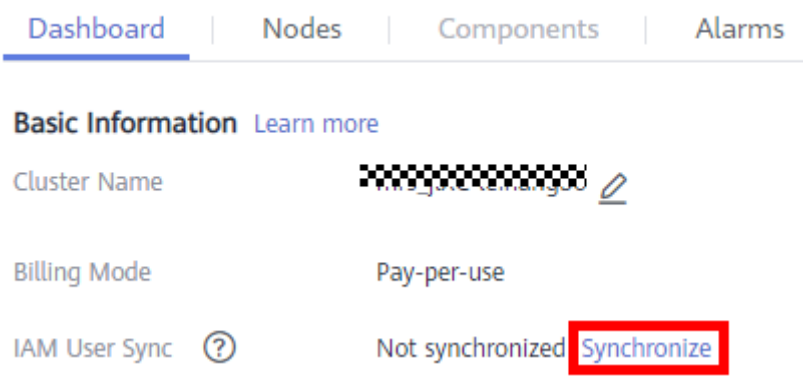

Procedure

Step 1 On the MRS details page, click **Nodes**.

Step 2 Unfold the node group information and select the check box of the target host that you want to cancel its isolation.

Step 3 Choose **Node Operation** > **Cancel Host Isolation**.

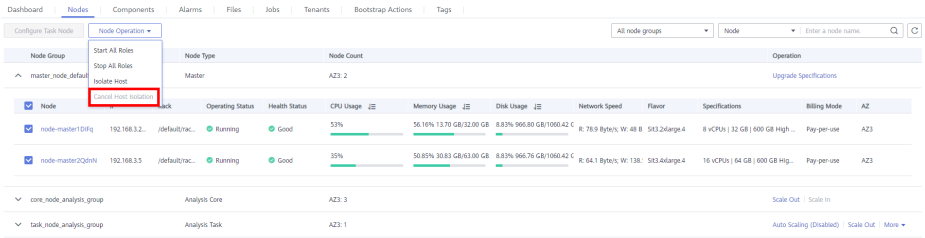

Step 4 Confirm the information about the host for which the isolation is to be cancelled and click **OK**.

When **Operation successful** is displayed, click **Finish**. The host is de-isolated successfully, and the value of **Operating Status** becomes **Normal**.

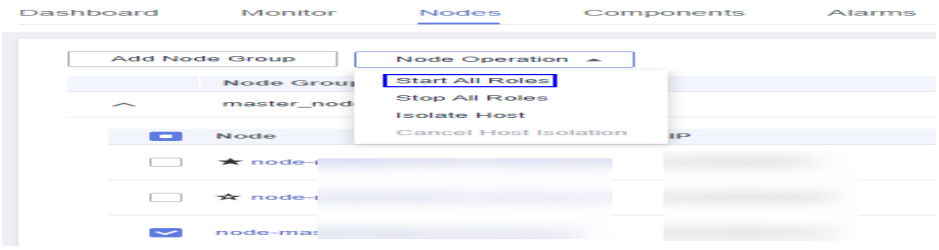

----End

3.5.7 Scaling Up Master Node Specifications

As users' increasing services lead to Core node scale-out and high CPU usage, Master node specifications cannot meet user requirements and need to be scaled up. This section describes how to scale up Master node specifications.

Prerequisites

- You have checked whether the Host Security Service (HSS) is enabled. If HSS is enabled, disable the HSS monitoring on the MRS cluster before you scale up master node specifications.
- Ensure that sufficient specification resources are available throughout the steps in **Scaling Up Master Node Specifications (Step-by-Step Upgrade)**.

Use Restrictions

- Master nodes can be scaled up for clusters with two or more master nodes.
- The specifications of the Master node in a BMS cluster cannot be upgraded.
- For MRS 1.8.2 or later to a version earlier than MRS 3.*x* or MRS 3.1.0 or later, see **Scaling Up Master Node Specifications (One-Click Upgrade)**.
- For MRS 3.0.5 and a version earlier than MRS 1.8.2, see **Scaling Up Master Node Specifications (Step-by-Step Upgrade)**.
- Do not perform other operations on the cluster during the scale-up.

Scaling Up Master Node Specifications (One-Click Upgrade)

Step 1 Log in to the MRS console.

Step 2 In the left navigation pane, choose **Clusters** > **Active Clusters**, select the cluster for which you want to scale up Master node specifications, and click its name to switch to the cluster details page.

Step 3 On the **Nodes** tab page, select **Scale Up Specifications** in the **Operation** column of the Master node group. The **Scale Up Master Node Specifications** page is displayed.

Step 4 Select the target specifications and click **Submit Order**. The order has been submitted successfully.

The node specification scale-up takes some time. After the scale-up is successful, the cluster status changes to **Running**.

NOTE

- The VM to be scaled up is automatically stopped during the scale-up and started after the scale-up is complete.
- The scale-up does not automatically upgrade the memory of components due to different component usage requirements. You can adjust the memory of components as needed.

----End

Scaling Up Master Node Specifications (Step-by-Step Upgrade)

Preparing for Scaling Up Master Node Specifications

Step 1 Log in to the MRS console.

Step 2 In the left navigation pane, choose **Clusters** > **Active Clusters**, select the cluster for which you want to scale up Master node specifications, and click its name to switch to the cluster details page.

Step 3 Ensure that the cluster status is **Running**.

Step 4 On the **Nodes** tab page, ensure that all nodes in the cluster are in the **Running** state.

Step 5 Log in to Manager and go to the cluster management page. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 6 Choose **Cluster** > **Services** > **ZooKeeper** > **Dashboard** and ensure that **Running Status** of the ZooKeeper service is **Normal**.

Figure 3-25 ZooKeeper service status

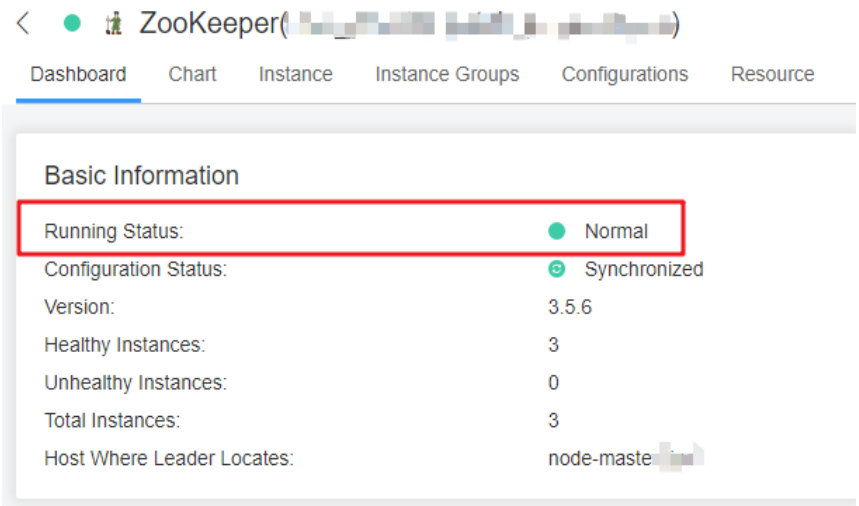

Step 7 Update service parameter settings as required. For details, see **Configuring Service Parameters**.

NOTE

You need to perform this step only once before scaling up the standby Master node.

Step 8 Choose **Cluster** > **Services** > **HDFS** > **Instance**.

Step 9 Record the business IP address of **NameNode (Standby)**. When upgrading the specifications of the active Master node, record the business IP address of **NameNode (Active)**. **Figure 3-26** shows the location of the business IP address.

Figure 3-26 Business IP address of the NameNode

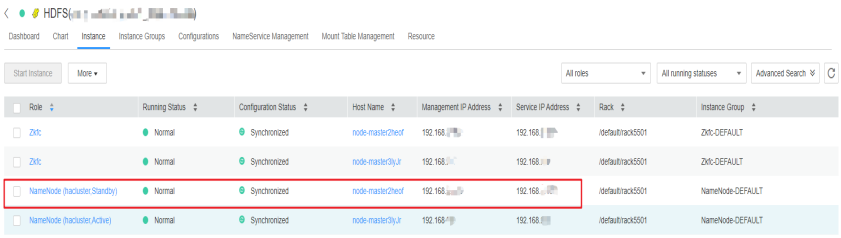

NOTE

Only when the cluster is an analysis cluster, you can perform **Step 8** to **Step 9** to record the IP addresses of the active and standby nodes.

Step 10 On the upper right of the Manager page, check the number next to the 

icon. If the number is 0, there is no running tasks in the cluster.

Step 11 Click **Hosts**. If the cluster is an analysis cluster, select the checkbox of the host corresponding to the business IP address of the **NameNode** recorded in **Step 9**. If the cluster is a streaming cluster, you do not need to distinguish the active and standby nodes. You only need to choose hosts for the scale-up.

Step 12 Choose **More** > **Stop All Instances** and wait until all instances are stopped.

NOTE

- When the node where Manager resides is scaled up, Manager may not be accessed due to an active/standby switchover. It is a normal phenomenon. Try to log in to Manager later. If the login fails for a long time, contact O&M personnel.
- After all roles are stopped, the following alarms may be generated. After the scale-up of Master node specifications is complete and all roles are started, the alarms are automatically cleared.
  - **ALM-12006 Node Fault**
  - **ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes**
  - **ALM-12039 Active/Standby OMS Databases Not Synchronized**
  - **ALM-14000 HDFS Service Unavailable**
  - **ALM-14010 NameService Service Is Abnormal**
  - **ALM-14012 JournalNode Is Out of Synchronization**
  - **ALM-16004 Hive Service Unavailable**
  - **ALM-18000 Yarn Service Unavailable**
  - **ALM-19000 HBase Service Unavailable**
  - **ALM-20002 Hue Service Unavailable**
  - **ALM-27001 DBService Service Unavailable**
  - **ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes**
  - **ALM-27004 Data Inconsistency Between Active and Standby DBServices**
  - **ALM-43001 Spark2x Service Unavailable**

----End

Scaling Up Master Node Specifications

Step 1 Log in to the MRS console.

Step 2 In the left navigation pane, choose **Clusters** > **Active Clusters**, select the cluster for which you want to scale up Master node specifications, and click its name to switch to the cluster details page.

Step 3 On the **Nodes** tab page, select **Scale Up Specifications** in the **Operation** column of the Master node group.

Step 4 Select the target specifications and click **Next**.

NOTE

Ensure that target specification resources are sufficient. Otherwise, the active node cannot be scaled up.

Step 5 On the **Confirm** page that is displayed, confirm the target node specifications and fees and click **OK**.

Step 6 Ensure that all services on the standby Master node have been stopped. For operation details, refer to **Step 1** to **Step 12** in the **Preparing for Scaling Up Master Node Specifications** part. On the **Scale Up Master Node Specifications** page, select **Are you sure you have stopped all services on the standby Master node?** and **If not all services are stopped before the scale-up, data saving failure or data damage may occur.** and click **Submit Order**.

Step 7 On the **Warning** page that is displayed, confirm again that all services on the standby Master node are stopped and click **OK** to start scaling up the specifications of the standby Master node.

The node specification scale-up takes some time. Please wait. After the scale-up is successful, the cluster status changes to **Scaled-up-first**. Otherwise, contact O&M personnel.

Step 8 After the standby Master node has been scaled up successfully, start all services and set parameters on the standby Master node by referring to **Step 1** to **Step 11** in the **Operations After the Master Node Specifications Scale-up** part.

Step 9 After the services on the standby Master node are started, perform an active/standby NameNode switchover. Perform this step only for an analysis cluster and skip this step for a streaming cluster.

1. Access the NameNode web UI of the active and standby nodes separately. For details about how to access the NameNode web UI, see **Step 11**.
2. In the navigation bar on the NameNode web UI, choose **Overview** and record the NameNode IDs of the active and standby nodes. Do not close the page after recording.

Figure 3-27 NameNode ID of the active node

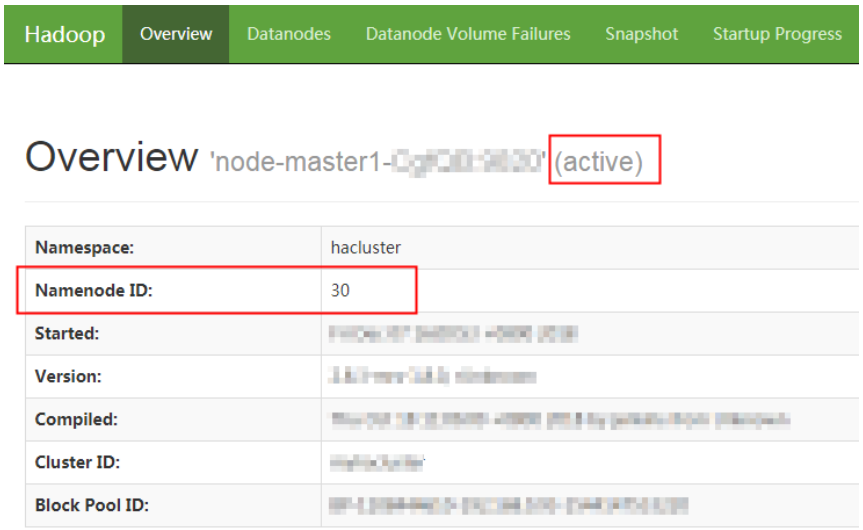

3. Log in to the ECS of any Master node and run the following command to configure environment variables:
   ```
   source /opt/Bigdata/client/bigdata_env
   ```
4. If the Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If the Kerberos authentication is disabled for the current cluster, skip this step.
   ```
   kinit MRS cluster user
   ```
   For example, **kinit admin**.
5. Run the following command to perform an active/standby NameNode switchover:
   ```
   hdfs haadmin -failover <NameNode ID of the active node> <NameNode ID of the standby node>
   ```
6. Go to the NameNode web UI page that is not closed in **Step 9.2** and refresh the page. You can view that the active/standby NameNode switchover is complete.

Figure 3-28 NameNode

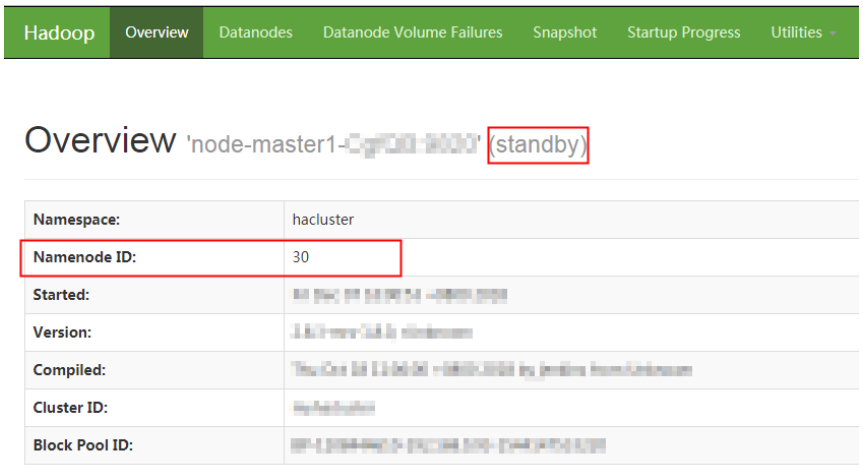

Step 10 Stop all services on the active Master node by referring to **Step 1** to **Step 12** in the **Preparing for Scaling Up Master Node Specifications** part.

Step 11 On the **Scale Up Master Node Specifications** page, select **I confirm that all services on the standby Master node have been started.** and **I confirm that all services on the active Master node have been stopped**, and click **Submit**.

Step 12 On the **Confirm** page that is displayed, confirm again that all services on the active Master node are stopped and click **OK** to start scaling up the specifications of the active Master node.

The node specification scale-up takes some time. Please wait. After the scale-up is successful, the cluster status changes to **Scaled-up-success**. Otherwise, contact O&M personnel.

Step 13 Start all services and set parameters on the active Master node by referring to **Step 1** to **Step 11** in the **Operations After the Master Node Specifications Scale-up** part.

Step 14 On the **Scale Up Master Node Specifications** page, select **Are you sure you have started all services on the active Master node?** and click **OK** to complete the scale-up.

----End

Operations After the Master Node Specifications Scale-up

Step 1 Log in to Manager and go to the cluster management page. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 Click **Hosts**. Check basic information about the host corresponding to the business IP address of the NameNode recorded in **Step 9** in the **Preparing for Scaling Up Master Node Specifications** part. If the **Running Status** is **Good** and **Disk**, **Memory**, and **CPU** have values, perform **Step 9**. If any of the preceding conditions is not met, go to the next step.

Step 3 Log in to the standby Master node remotely. For details, see **Logging In to an ECS**.

Step 4 Run the following command to switch to user **omm**:

```
su - omm
```


Step 5 Run the following command to start the Agent:

```
sh /opt/Bigdata/nodeagent/bin/start-agent.sh
```

Step 6 Run the following command to check whether the Agent is started successfully:

```
jps | grep NodeAgent
```

Step 7 Log in to Manager and go to the cluster management page. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 8 Click **Hosts**. Check basic information about the host corresponding to the business IP address of the NameNode recorded in **Step 9** in the **Preparing for Scaling Up Master Node Specifications** part to ensure that **Running Status** is **Good** and **Disk**, **Memory**, and **CPU** have values.

NOTE

It may take 3 minutes until the host status is normal after the Agent is started successfully. Please wait. If the host status is abnormal for a long time, contact O&M personnel.

Step 9 On the Manager page, click **Hosts** and select the checkbox of the host corresponding to the business IP address of the NameNode recorded in **Step 9** in the **Preparing for Scaling Up Master Node Specifications** part.

Step 10 Choose **More** > **Start All Instances** and wait until all instances are started.

Step 11 Access the NameNode web UI and check the NameNode startup status.

1. On FusionInsight Manager, choose **Cluster** > **Services** > **HDFS** > **Dashboard**.
2. In the **HDFS Summary** column, click **NameNode** of the active or standby node that has been scaled up on the right of **NameNode Web UI**.
3. Go to the **NameNode Web UI** page, choose **Startup Progress** in the navigation bar. After ensuring that **Percent Complete** is displayed as 100%, go to the next step, as shown in **Figure 3-29**.

Figure 3-29 NameNode startup status

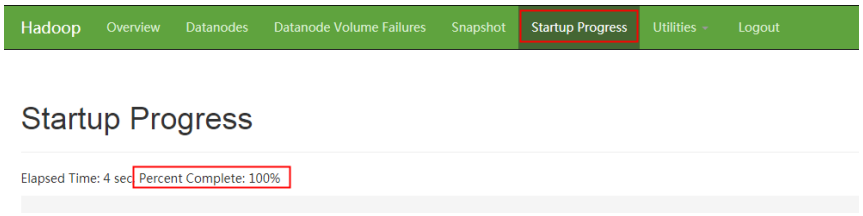

NOTE

Perform **Step 11** for an analysis cluster and skip this step for a streaming cluster.

----End

3.6 Job Management

3.6.1 Introduction to MRS Jobs

An MRS job is the program execution platform of MRS. It is used to process and analyze user data. After a job is created, all job information is displayed on the

Jobs tab page. You can view a list of all jobs and create and manage jobs. If the **Jobs** tab is not displayed on the cluster details page, submit a job in the background.

Data sources processed by MRS are from OBS or HDFS. OBS is an object-based storage service that provides you with massive, secure, reliable, and cost-effective data storage capabilities. MRS can process data in OBS directly. You can view, manage, and use data by using the web page of the management control platform or OBS client. In addition, you can use REST APIs independently or integrate APIs to service applications to manage and access data.

Before creating jobs, upload the local data to OBS for MRS to compute and analyze. MRS allows exporting data from OBS to HDFS for computing and analyzing. After the analyzing and computing are complete, you can store the data in HDFS or export them to OBS. HDFS and OBS can also store the compressed data in the format of **bz2** or **gz**.

Category

An MRS cluster allows creating and managing the following jobs: If a cluster in the **Running** state fails to create a job, check the health status of related components on the cluster management page. For details, see **Viewing and Customizing Cluster Monitoring Metrics**.

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

Job List

Tasks are listed in chronological order by default in the task list, with the most recent jobs displayed at the top. **Table 3-23** describes the parameters in the job list.

Table 3-23 Job list parameters

Parameter	Description
Name/ID	Job name, which is set when a job is created. ID is the unique identifier of a job. After a job is added, the system automatically assigns a value to ID.
Username	Name of the user who submits a job.
Type	The following data types are supported: • DistCp: importing and exporting data • MapReduce • Spark • SparkSubmit • SparkScript • Spark SQL • Hive SQL • HiveScript • Flink • Flink SQL • HadoopStreaming **NOTE** • After importing and exporting files on the **Files** tab page, you can view the DistCp job on the **Jobs** tab page. • Spark, Hive, and Flink jobs can be added only when the Spark, Hive, and Flink components are selected during cluster creation and the cluster is running.
Status	Job status. • Submitted • Accepted • Running • Completed • Terminated • Abnormal
Result	Execution result of a job. • Undefined: indicates that the job is being executed. • **Successful**: indicates that the job has been successfully executed. • **Killed**: indicates that the job is manually terminated during execution. • **Failed**: indicates that the job fails to be executed. **NOTE** Once a job has succeeded or failed, you cannot execute it again. However, you can add a job, and set job parameters to submit a job again.

Parameter	Description
Queue Name	Name of the queue bound to the user who submits the job
Submitted	Time when a job is submitted.
Ended	Time when a job is completed or manually stopped.
Operation	Operations can be performed on the job. Click **More** for more operations in the drop-down list. • Viewing Log: Click **View Log** to view the real-time logs of running jobs. For details, see **Viewing Job Configuration and Logs**. • View Details: Click **View Details** to view the detailed configuration information about jobs. For details, see **Viewing Job Configuration and Logs**. • Stop: You can click **Stop** to stop a running job. For details, see **Stopping a Job**. • View Result: Click **View Result** to view the execution results of SparkSQL and SparkScript jobs whose status is **Completed** and result is **Successful**. • **Delete**: Delete the job. For details, see **Deleting a Job**. **NOTE** • A deleted job cannot be restored. • If you choose to save job logs to OBS or HDFS, the system compresses and saves the logs to the corresponding path after the job execution is completed. Therefore, after a job execution of this type is completed, the job status is still **Running**. After the log is successfully stored, the job status changes to **Completed**. The log storage duration depends on the log size and takes several minutes.

Table 3-24 Icon description

Icon	Description
	Select a time range for job submission to filter jobs submitted in the time range.
	Select a certain job execution result from the drop-down list to display jobs of the status. • All statuses: Filter all jobs. • Successful: Filter jobs that are successfully executed. • Undefined: Filter jobs that are being executed. • Killed: Filter jobs that are manually stopped. • Failed: Filter jobs that fail to be executed.

Icon	Description
	Select a certain job type from the drop-down list to display jobs of the type. • All types • MapReduce • HiveScript • Distcp • SparkScript • Spark SQL • Hive SQL • SparkSubmit • Flink • Flink SQL • HadoopStreaming
	In the search box, search for a job by setting the corresponding search condition and click  . • Job name. • Job ID. • Username. • Queue name.
	Click  to manually refresh the job list.

Job Execution Permission Description

For a security cluster with Kerberos authentication enabled, a user needs to synchronize an IAM user before submitting a job on the MRS web UI. After the synchronization is completed, the MRS system generates a user with the same IAM username. Whether a user has the permission to submit jobs depends on the IAM policy bound to the user during IAM synchronization. For details about the job submission policy, see **Table 1-3** in **Synchronizing IAM Users to MRS**.

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

Step 1 Log in to Manager as user **admin**.

Step 2 Add the role of the component whose permission is required by the user. For details, see **Creating a Role**.

Step 3 Change the user group to which the user who submits the job belongs and add the new component role to the user group. For details, see **Related Tasks**.

NOTE

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

----End

3.6.2 Running a MapReduce Job

You can submit programs developed by yourself to MRS to execute them, and obtain the results. This section describes how to submit a MapReduce job on the MRS management console. MapReduce jobs are used to submit JAR programs to quickly process massive amounts of data in parallel and create a distributed data processing and execution environment.

If the job and file management functions are not supported on the cluster details page, submit the jobs in the background.

Prerequisites

You have uploaded the program packages and data files required for running jobs to OBS or HDFS.

Before you upload the program packages and data files to OBS, you need to create an OBS agency and bind it to the MRS cluster. For details, see **Configuring a Storage-Compute Decoupled Cluster (Agency)**.

Submitting a Job on the GUI

Step 1 Log in to the MRS console.

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

Step 3 If Kerberos authentication is enabled for the cluster, perform the following steps. If Kerberos authentication is not enabled for the cluster, skip this step.

In the **Basic Information** area on the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users. For details, see **Synchronizing IAM Users to MRS**.

NOTE

- When the policy of the user group to which the IAM user belongs changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from MRS CommonOperations, MRS FullAccess, or MRS Administrator to MRS ReadOnlyAccess, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.

Step 4 Click the **Jobs** tab.

Step 5 Click **Create**. The **Create Job** page is displayed.

NOTE

If the IAM username contains spaces (for example, **admin 01**), a job cannot be created.

Step 6 In **Type**, select **MapReduce**. Configure other job information.

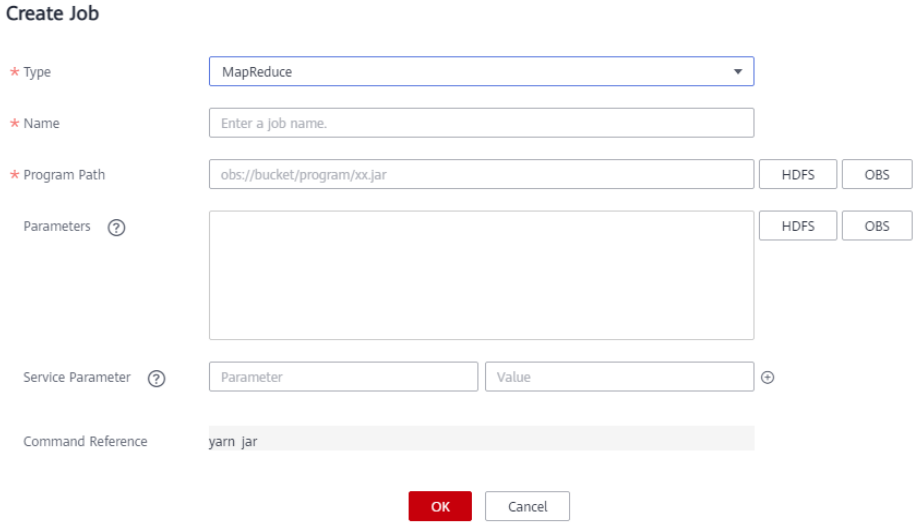

Table 3-25 Job configuration information

Parameter	Description
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. **NOTE** You are advised to set different names for different jobs.
Program Path	Path of the program package to be executed. The following requirements must be met: • Contains a maximum of 1,023 characters, excluding special characters such as ;|&><'$. The parameter value cannot be empty or full of spaces. • The path of the program to be executed can be stored in HDFS or OBS. The path varies depending on the file system. – OBS: The path must start with **obs://**. Example: **obs://wordcount/program/*xxx*.jar** – HDFS: The path must start with **/user**. For details about how to import data to HDFS, see **Importing Data**. • For SparkScript and HiveScript, the path must end with **.sql**. For MapReduce, the path must end with **.jar**. For Flink and SparkSubmit, the path must end with **.jar** or **.py**. The **.sql**, **.jar**, and **.py** are case-insensitive.

Parameter	Description
Parameters	(Optional) It is the key parameter for program execution. Multiple parameters are separated by space. Configuration method: *Program class name Data input path Data output path* • Program class name: It is specified by a function in your program. MRS is responsible for transferring parameters only. • Data input path: Click **HDFS** or **OBS** to select a path or manually enter a correct path. • Data output path: Enter a directory that does not exist. The parameter contains a maximum of 150,000 characters. It cannot contain special characters ;|&><'$, but can be left blank. **CAUTION** If you enter a parameter with sensitive information (such as the login password), the parameter may be exposed in the job details display and log printing. Exercise caution when performing this operation.
Service Parameter	(Optional) It is used to modify service parameters for the job. The parameter modification applies only to the current job. To make the modification take effect permanently for the cluster, follow instructions in **Configuring Service Parameters**. To add multiple parameters, click ⊕ on the right. To delete a parameter, click **Delete** on the right. **Table 3-26** lists the common service configuration parameters.
Command Reference	Command submitted to the background for execution when a job is submitted.

Table 3-26 Service Parameter parameters

Parameter	Description	Example Value
fs.obs.access.key	Key ID for accessing OBS.	-
fs.obs.secret.key	Key corresponding to the key ID for accessing OBS.	-

Step 7 Confirm job configuration information and click **OK**.

After the job is created, you can manage it.

----End

Submitting a Job in the Background

The default client installation path for MRS 3.x or later is /opt/Bigdata/client, and for versions earlier than MRS 3.x is /opt/client. Configure the path based on site requirements.

Step 1 Log in to the MRS console.

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

Step 3 On the **Nodes** tab page, click the name of a Master node to go to the ECS management console.

Step 4 Click **Remote Login** in the upper right corner of the page.

Step 5 Enter the username and password of the Master node as prompted. The username is **root** and the password is the one set during cluster creation.

Step 6 Run the following command to initialize environment variables:

source /opt/Bigdata/client/bigdata_env

Step 7 If the Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If the Kerberos authentication is disabled for the current cluster, skip this step.

kinit *MRS cluster user*

Example: **kinit admin**

Step 8 Run the following command to copy the program in the OBS file system to the Master node in the cluster:

hadoop fs -Dfs.obs.access.key=AK -Dfs.obs.secret.key=SK -copyToLocal source_path.jar target_path.jar

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

NOTE

- Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.
- You can log in to the OBS Console using AK/SK. To obtain AK/SK information, click the username in the upper right corner of the management console and choose **My Credentials** > **Access Keys**.

Step 9 Run the following command to submit a wordcount job. If data needs to be read from OBS or outputted to OBS, the AK/SK parameters need to be added.

source /opt/Bigdata/client/bigdata_env;hadoop jar execute_jar wordcount input_path output_path

Example: **source /opt/Bigdata/client/bigdata_env;hadoop jar /home/omm/hadoop-mapreduce-examples-XXX.jar wordcount -Dfs.obs.access.key=XXXX -Dfs.obs.secret.key=XXXX "obs://mrs-word/input/*" "obs://mrs-word/output/"**

In the preceding command, **input_path** indicates a path for storing job input files on OBS. **output_path** indicates a path for storing job output files on OBS and needs to be set to a directory that does not exist

----End

3.6.3 Running a SparkSubmit or Spark Job

You can submit programs developed by yourself to MRS to execute them, and obtain the results. This section describes how to submit a Spark job on the MRS console.

Prerequisites

You have uploaded the program packages and data files required for running jobs to OBS or HDFS.

NOTE

- The example JAR file provided by the system is *{Client installation directory}***/Spark2x/spark/examples/jars/spark-examples_*.jar**.
- Log in to the client node and run the following command to upload the JAR package to HDFS (for example, **/tmp**):

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

Submitting a Job on the GUI

Step 1 Log in to the MRS console.

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

Step 3 If Kerberos authentication is enabled for the cluster, perform the following steps. If Kerberos authentication is not enabled for the cluster, skip this step.

In the **Basic Information** area on the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users. For details, see **Synchronizing IAM Users to MRS**.

NOTE

- When the policy of the user group to which the IAM user belongs changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from MRS CommonOperations, MRS FullAccess, or MRS Administrator to MRS ReadOnlyAccess, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.

Step 4 Click the **Jobs** tab.

Step 5 Click **Create**. The **Create Job** page is displayed.

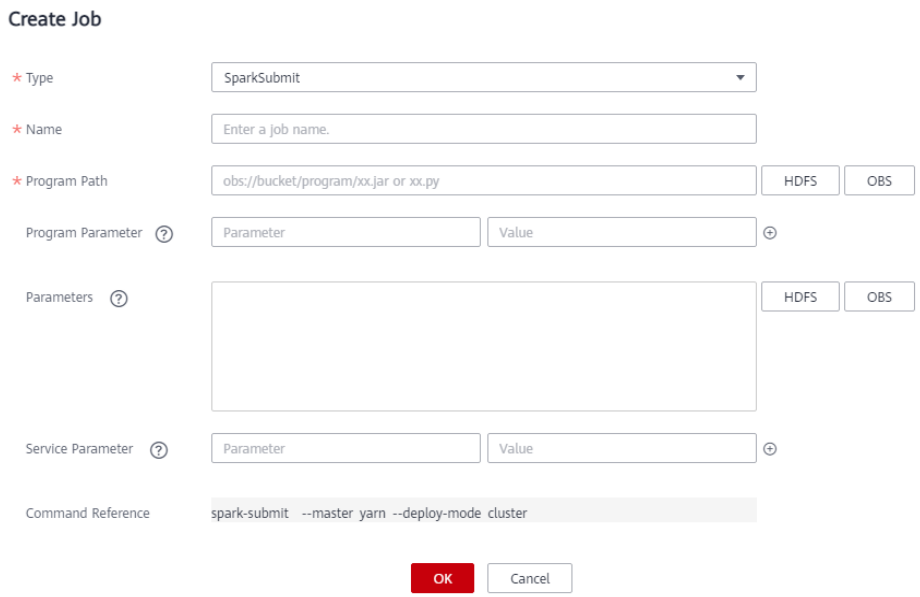

Step 6 Configure job information.

Table 3-27 Job configuration information

Parameter	Description
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. **NOTE** You are advised to set different names for different jobs.
Program Path	Path of the program package to be executed. The following requirements must be met: • Contains a maximum of 1,023 characters, excluding special characters such as ;|&><'$. The parameter value cannot be empty or full of spaces. • The path of the program to be executed can be stored in HDFS or OBS. The path varies depending on the file system. – OBS: The path must start with **obs://**. Example: **obs://wordcount/program/*xxx*.jar** (supported in MRS 1.9.2 or later) – HDFS: The path must start with **/user**. For details about how to import data to HDFS, see **Importing Data**. • For SparkScript and HiveScript, the path must end with **.sql**. For MapReduce, the path must end with **.jar**. For Flink and SparkSubmit, the path must end with **.jar** or **.py**. The **.sql**, **.jar**, and **.py** are case-insensitive.
Program Parameter	(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance. **Table 3-28** describes the common parameters of a running program.

Parameter	Description
Parameters	(Optional) Key parameter for program execution. The parameter is specified by the function of the user's program. MRS is only responsible for loading the parameter. Multiple parameters are separated by space. The parameter contains a maximum of 150,000 characters. It cannot contain special characters ;|&><'$, but can be left blank. **CAUTION** If you enter a parameter with sensitive information (such as the login password), the parameter may be exposed in the job details display and log printing. Exercise caution when performing this operation.
Service Parameter	(Optional) It is used to modify service parameters for the job. The parameter modification applies only to the current job. To make the modification take effect permanently for the cluster, follow instructions in **Configuring Service Parameters**. To add multiple parameters, click ⊕ on the right. To delete a parameter, click **Delete** on the right. **Table 3-29** lists the common service configuration parameters. **NOTE** If you need to run a long-term job, such as SparkStreaming, and access OBS, you need to use **Service Parameter** to import the AK/SK for accessing OBS.
Command Reference	Command submitted to the background for execution when a job is submitted.

Table 3-28 Program parameters

Parameter	Description	Example Value
--conf	Add the task configuration items.	spark.executor.memory=2G
--driver-memory	Set the running memory of driver.	2G
--num-executors	Set the number of executors to be started.	5
--executor-cores	Set the number of executor cores.	2
--class	Set the main class of a task.	org.apache.spark.examples.SparkPi
--files	Upload files to a task. The files can be custom configuration files or some data files from OBS or HDFS.	-
--jars	Upload additional dependency packages of a task to add the external dependency packages to the task.	-

Parameter	Description	Example Value
--executor-memory	Set executor memory.	2G
--conf spark-yarn.maxAppAttempts	Control the number of AM retries.	If this parameter is set to **0**, retry is not allowed. If this parameter is set to **1**, one retry is allowed.

Table 3-29 Service Parameter parameters

Parameter	Description	Example Value
fs.obs.access.key	Key ID for accessing OBS.	-
fs.obs.secret.key	Key corresponding to the key ID for accessing OBS.	-

Step 7 Confirm job configuration information and click **OK**.

After the job is created, you can manage it.

----End

Submitting a Job in the Background

The default client installation path for MRS 3.x or later is /opt/Bigdata/client, and for versions earlier than MRS 3.x is /opt/client. Configure the path based on site requirements.

Step 1 Create a user for submitting jobs. For details, see **Creating a User**.

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

Step 2 Download the authentication credential.

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

Step 3 Upload JAR files related to the job to the cluster. In this example, the sample JAR file built in Spark is used. It is stored in **$SPARK_HOME/examples/jars**.

Step 4 Upload the authentication credential of the user created in **Step 2** to the **/opt** directory of the cluster and run the following command to decompress the credential:

tar –xvf MRSTest _*xxxxxx*_keytab.tar

You will obtain two files: **user.keytab** and **krb5.conf**.

Step 5 Before performing operations on the cluster, run the following commands:

source /opt/Bigdata/client/bigdata_env

cd $SPARK_HOME

Step 6 Run the following command to submit the Spark job:

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

Parameter description:

1. Computing capability of Yarn, which specifies that the job is submitted in client mode.
2. Configuration item of the Spark job. The authentication file and username are transferred here.
3. **spark.yarn.principal**: user created in step 1
4. **spark.yarn.keytab**: keytab file used for authentication
5. *xx*.**jar**: JAR file used by the job

----End

3.6.4 Running a HiveSQL Job

You can submit programs developed by yourself to MRS to execute them, and obtain the results. This section describes how to submit a HiveSQL job on the MRS management console. HiveSQL jobs are used to submit SQL statements and script files for data query and analysis. Both SQL statements and scripts are supported. If SQL statements contain sensitive information, use Script to submit them.

Prerequisites

You have uploaded the program packages and data files required for running jobs to OBS or HDFS.

Submitting a Job on the GUI

Step 1 Log in to the MRS console.

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

Step 3 If Kerberos authentication is enabled for the cluster, perform the following steps. If Kerberos authentication is not enabled for the cluster, skip this step.

In the **Basic Information** area on the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users. For details, see **Synchronizing IAM Users to MRS**.

NOTE

- When the policy of the user group to which the IAM user belongs changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from MRS CommonOperations, MRS FullAccess, or MRS Administrator to MRS ReadOnlyAccess, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.

Step 4 Click the **Jobs** tab.

Step 5 Click **Create**. The **Create Job** page is displayed.

Step 6 Configure job information. Set **Type** to **HiveSql** and configure HiveSQL job information be referring to **Table 3-30**.

Table 3-30 Job configuration information

Parameter	Description
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. **NOTE** You are advised to set different names for different jobs.
SQL Type	Submission type of the SQL statement • SQL • Script
SQL Statement	This parameter is valid only when **SQL Type** is set to **SQL**. Enter the SQL statement to be executed, and then click **Check** to check whether the SQL statement is correct. If you want to submit and execute multiple statements at the same time, use semicolons (;) to separate them.

Parameter	Description
SQL File	This parameter is valid only when **SQL Type** is set to **Script**. The path of the SQL file to be executed must meet the following requirements: • Contains a maximum of 1,023 characters, excluding special characters such as ;|&><'$. The parameter value cannot be empty or full of spaces. • The path of the program to be executed can be stored in HDFS or OBS. The path varies depending on the file system. – OBS: The path must start with **obs://**. Example: **obs://wordcount/program/*xxx*.jar** – HDFS: The path must start with **/user**. For details about how to import data to HDFS, see **Importing Data**. • For SparkScript and HiveScript, the path must end with **.sql**. For MapReduce, the path must end with **.jar**. For Flink and SparkSubmit, the path must end with **.jar** or **.py**. The **.sql**, **.jar**, and **.py** are case-insensitive. **NOTE** A file path on OBS can start with **obs://**. To submit jobs in this format, you need to configure permissions for accessing OBS. • If the OBS permission control function is enabled during cluster creation, you can use the **obs://** directory without extra configuration. • If the OBS permission control function is not enabled or is not supported when you create a cluster, configure the function by following instructions in **Accessing OBS**.
Program Parameter	(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance. **Table 3-31** describes the common parameters of a running program.
Service Parameter	(Optional) It is used to modify service parameters for the job. The parameter modification applies only to the current job. To make the modification take effect permanently for the cluster, follow instructions in **Configuring Service Parameters**. To add multiple parameters, click ⊕ on the right. To delete a parameter, click **Delete** on the right. **Table 3-32** lists the common service configuration parameters.
Command Reference	Command submitted to the background for execution when a job is submitted.

Table 3-31 Program parameters

Parameter	Description	Example Value
--hiveconf	Hive service configuration, for example, set the execution engine to MapReduce.	Setting the execution engine to MR: **--hiveconf "hive.execution.engine=mr"**
--hivevar	Custom variable, for example, variable ID.	Setting the variable ID: **--hivevar id="123" select * from test where id = ${hivevar:id}**

Table 3-32 Service parameters

Parameter	Description	Example Value
fs.obs.access.key	Key ID for accessing OBS.	-
fs.obs.secret.key	Key corresponding to the key ID for accessing OBS.	-
hive.execution.engine	Engine for running a job.	• mr • tez

Step 7 Confirm job configuration information and click **OK**.

After the job is created, you can manage it.

----End

Submitting a Job in the Background

Step 1 Log in to the MRS console.

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

Step 3 On the **Nodes** tab page, click the name of a Master node to go to the ECS management console.

Step 4 Click **Remote Login** in the upper right corner of the page.

Step 5 Enter the username and password of the Master node as prompted. The username is **root** and the password is the one set during cluster creation.

Step 6 Run the following command to initialize environment variables:

source /opt/BigData/client/bigdata_env

NOTE

- The default client installation path for MRS 3.x or later is /opt/Bigdata/client, and for versions earlier than MRS 3.x is /opt/client. Configure the path based on site requirements.

Step 7 If the Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If the Kerberos authentication is disabled for the current cluster, that is, the cluster is in normal mode, skip this step.

kinit *MRS cluster user* (The user must be in the **hive** user group.)

Step 8 Run the **beeline** command to connect to HiveServer and run tasks.

beeline

For clusters in normal mode, run the following commands. If no component service user is specified, the current OS user is used to log in to the HiveServer.

beeline -n *Component service user*

beeline -f *SQL files* (SQLs in the execution files)

----End

3.6.5 Running a SparkSql Job

You can submit programs developed by yourself to MRS to execute them, and obtain the results. This section describes how to submit a SparkSQL job on the MRS console. SparkSQL jobs are used for data query and analysis. Both SQL statements and scripts are supported. If SQL statements contain sensitive information, use Spark Script to submit them.

Prerequisites

You have uploaded the program packages and data files required for running jobs to OBS or HDFS.

Submitting a Job on the GUI

Step 1 Log in to the MRS console.

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

Step 3 If Kerberos authentication is enabled for the cluster, perform the following steps. If Kerberos authentication is not enabled for the cluster, skip this step.

In the **Basic Information** area on the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users. For details, see **Synchronizing IAM Users to MRS**.

NOTE

- When the policy of the user group to which the IAM user belongs changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from MRS CommonOperations, MRS FullAccess, or MRS Administrator to MRS ReadOnlyAccess, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.

Step 4 Click the **Jobs** tab.

Step 5 Click **Create**. On the displayed **Create Job** page, set **Type** to **SparkSql** and configure SparkSql job information by referring to **Table 3-33**.

Table 3-33 Job configuration information

Parameter	Description
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. **NOTE** You are advised to set different names for different jobs.
SQL Type	Submission type of the SQL statement • SQL • Script
SQL Statement	This parameter is valid only when **SQL Type** is set to **SQL**. Enter the SQL statement to be executed, and then click **Check** to check whether the SQL statement is correct. If you want to submit and execute multiple statements at the same time, use semicolons (;) to separate them.

Parameter	Description
SQL File	This parameter is valid only when **SQL Type** is set to **Script**. The path of the SQL file to be executed must meet the following requirements: • Contains a maximum of 1,023 characters, excluding special characters such as ;|&><'$. The parameter value cannot be empty or full of spaces. • The path of the program to be executed can be stored in HDFS or OBS. The path varies depending on the file system. – OBS: The path must start with **obs://**. Example: **obs://wordcount/program/*xxx*.jar** – HDFS: The path must start with **/user**. For details about how to import data to HDFS, see **Importing Data**. • For SparkScript and HiveScript, the path must end with **.sql**. For MapReduce, the path must end with **.jar**. For Flink and SparkSubmit, the path must end with **.jar** or **.py**. The **.sql**, **.jar**, and **.py** are case-insensitive. **NOTE** A file path on OBS can start with **obs://**. To submit jobs in this format, you need to configure permissions for accessing OBS. • If the OBS permission control function is enabled during cluster creation, you can use the **obs://** directory without extra configuration. • If the OBS permission control function is not enabled or is not supported when you create a cluster, configure the function by following instructions in **Accessing OBS**.
Program Parameter	(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance. **Table 3-34** describes the common parameters of a running program.
Service Parameter	(Optional) It is used to modify service parameters for the job. The parameter modification applies only to the current job. To make the modification take effect permanently for the cluster, follow instructions in **Configuring Service Parameters**. To add multiple parameters, click ⊕ on the right. To delete a parameter, click **Delete** on the right. **Table 3-35** lists the common service configuration parameters.
Command Reference	Command submitted to the background for execution when a job is submitted.

Table 3-34 Program parameters

Parameter	Description	Example Value
--conf	Task configuration items to be added.	spark.executor.memory=2G
--driver-memory	Running memory of a driver.	2G
--num-executors	Number of executors to be started.	5
--executor-cores	Number of executor cores.	2
--jars	Additional dependency packages of a task, which is used to add the external dependency packages to the task.	-
--executor-memory	Executor memory.	2G

Table 3-35 Service parameters

Parameter	Description	Example Value
fs.obs.access.key	Key ID for accessing OBS.	-
fs.obs.secret.key	Key corresponding to the key ID for accessing OBS.	-

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

After the job is created, you can manage it.

----End

Submitting a Job in the Background

The default client installation path for MRS 3.x or later is /opt/Bigdata/client, and for versions earlier than MRS 3.x is /opt/client. Configure the path based on site requirements.

Step 1 Create a user for submitting jobs. For details, see **Creating a User**.

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

Step 2 Download the authentication credential.

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

Step 3 Log in to the node where the Spark client is located, upload the user authentication credential created in 2 to the **/opt** directory of the cluster, and run the following command to decompress the package:

tar -xvf MRSTest *_xxxxxx_***keytab.tar**

After the decompression, you obtain the **user.keytab** and **krb5.conf** files.

Step 4 Before performing operations on the cluster, run the following commands:

source /opt/Bigdata/client/bigdata_env

cd $SPARK_HOME

Step 5 Open the **spark-sql** CLI and run the following SQL statement:

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

To execute the SQL file, you need to upload the SQL file (for example, to the **/opt/** directory). After the file is uploaded, run the following command:

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

----End

3.6.6 Running a Flink Job

You can submit programs developed by yourself to MRS to execute them, and obtain the results. This section describes how to submit a Flink job on the MRS management console. Flink jobs are used to submit JAR programs to process streaming data.

Prerequisites

You have uploaded the program packages and data files required for running jobs to OBS or HDFS.

Submitting a Job on the GUI

Step 1 Log in to the MRS console.

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

Step 3 If Kerberos authentication is enabled for the cluster, perform the following steps. If Kerberos authentication is not enabled for the cluster, skip this step.

In the **Basic Information** area on the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users. For details, see **Synchronizing IAM Users to MRS**.

NOTE

- When the policy of the user group to which the IAM user belongs changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from MRS CommonOperations, MRS FullAccess, or MRS Administrator to MRS ReadOnlyAccess, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.

Step 4 Click the **Jobs** tab.

Step 5 Click **Create**. The **Create Job** page is displayed.

Step 6 Set **Type** to **Flink**. Configure Flink job information by referring to **Table 3-36**.

Table 3-36 Job configuration information

Parameter	Description
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. **NOTE** You are advised to set different names for different jobs.

Parameter	Description
Program Path	Path of the program package to be executed. The following requirements must be met: • Contains a maximum of 1,023 characters, excluding special characters such as ;|&><'$. The parameter value cannot be empty or full of spaces. • The path of the program to be executed can be stored in HDFS or OBS. The path varies depending on the file system. – OBS: The path must start with **obs://**. Example: **obs://wordcount/program/*xxx*.jar** (Supported in MRS 3.x or later) – HDFS: The path must start with **/user**. For details about how to import data to HDFS, see **Importing Data**. **NOTE** If you use an OBS path starting with **obs://**, configure permission for accessing OBS as follows: • If the OBS permission control function is enabled during cluster creation, you can use the **obs://** directory without extra configuration. • If the OBS permission control function is not enabled or not available during cluster creation, perform the following steps: 1. On the MRS cluster details page, click the **Nodes** tab and expand a node group. 2. Click a node name to go to the cloud server console. 3. Click on the right of **Agency**, select **MRS_ECS_DEFAULT_AGENCY** and add it. 4. Repeat the preceding steps to add agencies for all nodes in the cluster.
Program Parameter	(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance. **Table 3-37** describes the common parameters of a running program.
Parameters	(Optional) Key parameter for program execution. The parameter is specified by the function of the user's program. MRS is only responsible for loading the parameter. Multiple parameters are separated by space. The parameter contains a maximum of 150,000 characters. It cannot contain special characters ;|&><'$, but can be left blank. **CAUTION** If you enter a parameter with sensitive information (such as the login password), the parameter may be exposed in the job details display and log printing. Exercise caution when performing this operation.

Parameter	Description
Service Parameter	(Optional) It is used to modify service parameters for the job. The parameter modification applies only to the current job. To make the modification take effect permanently for the cluster, follow instructions in **Configuring Service Parameters**. To add multiple parameters, click ⊕ on the right. To delete a parameter, click **Delete** on the right. **Table 3-38** describes the common parameters of a service.
Command Reference	Command submitted to the background for execution when a job is submitted.

Table 3-37 Program parameters

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

NOTE

For MRS 3.x or later, the **-yn** parameter is not supported.

Table 3-38 Service parameters

Parameter	Description	Example Value
fs.obs.access.key	Key ID for accessing OBS.	-
fs.obs.secret.key	Key corresponding to the key ID for accessing OBS.	-

Step 7 Confirm job configuration information and click **OK**.

After the job is created, you can manage it.

----End

Submitting a Job in the Background

The default client installation path for MRS 3.x or later is /opt/Bigdata/client, and for versions earlier than MRS 3.x is /opt/client. Configure the path based on site requirements.

Step 1 Log in to the MRS client.

Step 2 Run the following command to initialize environment variables:

source /opt/Bigdata/client/bigdata_env

Step 3 If Kerberos authentication is enabled for the cluster, perform the following steps. If Kerberos authentication is not enabled for the cluster, skip this step.

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

   Example:

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

   sh generate_keystore.sh

   This script automatically replaces the SSL value in the **/opt/Bigdata/client/Flink/flink/conf/flink-conf.yaml** file. For MRS 3.x or earlier, external SSL is disabled by default in security clusters. To enable external SSL, run this script

again after configuration. The configuration parameters do not exist in the default Flink configuration of MRS, if you enable SSL for external connections, you need to add the parameters listed in **Table 3-39**.

Table 3-39 Parameter description

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

NOTE

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

  curl -k -i -u <*user name*>:<*password*> -X POST -HContent-type:application/json -d '{"plainText":"<*password*>"}' 'https://*x.x.x.x*:28443/web/api/v2/tools/encrypt'; In the preceding command, <*password*> must be the same as the password used for issuing the certificate, and *x.x.x.x* indicates the floating IP address of Manager in the cluster.

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

7. Configure paths for the client to access the **flink.keystore** and **flink.truststore** files.
   - Absolute path: After the script is executed, the file path of **flink.keystore** and **flink.truststore** is automatically set to the absolute path **opt/Bigdata/client/Flink/flink/conf/** in the **flink-conf.yaml** file. In this case,

you need to move the **flink.keystore** and **flink.truststore** files from the **conf** directory to this absolute path on the Flink client and Yarn nodes.

- Relative path: Perform the following steps to set the file path of **flink.keystore** and **flink.truststore** to the relative path and ensure that the directory where the Flink client command is executed can directly access the relative paths.
  i. In the **/opt/Bigdata/client/Flink/flink/conf/**directory, create a new directory, for example, **ssl**.
  ii. Move the **flink.keystore** and **flink.truststore** file to the /**opt/Bigdata/client/Flink/flink/conf/ssl/** directory.
  iii. For MRS 3.*x* or later: Change the values of the following parameters in the **flink-conf.yaml** file to relative paths:
  ```
  security.ssl.keystore: ssl/flink.keystore
  security.ssl.truststore: ssl/flink.truststore
  ```
  iv. For MRS 3.x or earlier: Change the values of the following parameters in the **flink-conf.yaml** file to relative paths:
  ```
  security.ssl.internal.keystore: ssl/flink.keystore
  security.ssl.internal.truststore: ssl/flink.truststore
  ```

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

Step 4 Run a wordcount job.

- Normal cluster (Kerberos authentication disabled)
  - Run the following commands to start a session and submit a job in the session:
    ```
    yarn-session.sh -nm "session-name" -d
    flink run /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
    ```
  - Run the following command to submit a single job on Yarn:
    ```
    flink run -m yarn-cluster /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
    ```
- Security cluster (Kerberos authentication enabled)
  - If the **flink.keystore** and **flink.truststore** file are stored in the absolute path:
    - Run the following commands to start a session and submit a job in the session:
      ```
      yarn-session.sh -nm "session-name" -d
      flink run /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
      ```
    - Run the following command to submit a single job on Yarn:
      ```
      flink run -m yarn-cluster /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
      ```
  - If the **flink.keystore** and **flink.truststore** file are stored in the relative path:
    - In the same directory of SSL, run the following command to start a session and submit jobs in the session. The SSL directory is a relative path. For example, if the SSL directory is **opt/Bigdata/client/Flink/flink/conf/**, then run the following command in this directory:
      ```
      yarn-session.sh -t ssl/ -nm "session-name" -d
      flink run /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar
      ```


- Run the following command to submit a single job on Yarn:
  flink run -m yarn-cluster -yt ssl/ /opt/Bigdata/client/Flink/flink/examples/streaming/WordCount.jar

----End

3.6.7 Running a HadoopStreaming Job

You can submit programs developed by yourself to run them on MRS, and obtain the results. This topic describes how to submit a HadoopStreaming job on the MRS management console.

Submitting a Job on the UI

Step 1 Log in to the MRS console.

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

Step 3 If Kerberos authentication is enabled for the cluster, perform the following steps. If Kerberos authentication is not enabled for the cluster, skip this step.

In the **Basic Information** area on the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users. For details, see **Synchronizing IAM Users to MRS**.

NOTE

- When the policy of the user group to which the IAM user belongs changes from MRS ReadOnlyAccess to MRS CommonOperations, MRS FullAccess, or MRS Administrator, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** (System Security Services Daemon) cache of cluster nodes needs time to be updated. Then, submit a job. Otherwise, the job may fail to be submitted.
- When the policy of the user group to which the IAM user belongs changes from MRS CommonOperations, MRS FullAccess, or MRS Administrator to MRS ReadOnlyAccess, wait for 5 minutes until the new policy takes effect after the synchronization is complete because the **SSSD** cache of cluster nodes needs time to be updated.

Step 4 Click the **Jobs** tab.

Step 5 Click **Create**. The **Create Job** page is displayed.

Step 6 Set **Type** to **HadoopStreaming**. Configure job information by referring to **Table 3-40**.

Table 3-40 Job parameters

Parameter	Description
Name	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. **NOTE** You are advised to set different names for different jobs.

Parameter	Description
Program Parameter	(Optional) Used to configure optimization parameters such as threads, memory, and vCPUs for the job to optimize resource usage and improve job execution performance. **Table 3-41** describes the common parameters of a running program.
Parameters	(Optional) Key parameter for program execution. The parameter is specified by the function of the custom program. MRS is only responsible for loading the parameters. Use spaces to separate parameters. To prevent parameters from being saved as plaintext, add an at sign (@) before parameters. The value can contain a maximum of 150,000 characters. It cannot contain special characters ;|&><'$, but can be left blank. **CAUTION** If you enter a parameter with sensitive information (such as the login password), the parameter may be exposed in the job details and logs. Exercise caution when performing this operation.
Service Parameter	(Optional) Service parameters for the job. The parameters apply only to this job. To apply modifications to the cluster, follow instructions in **Configuring Service Parameters**. To add more parameters, click ⊕ on the right. To delete a parameter, click **Delete** on the right. **Table 3-42** describes the typical service parameters.
Command Reference	Commands submitted to the background when the job is submitted.

Table 3-41 Program parameters

Parameter	Description	Example Value
-ytm	Memory size of each TaskManager container. (Optional unit. The unit is MB by default.)	1024
-yjm	Memory size of JobManager container. (Optional unit. The unit is MB by default.)	1024
-yn	Number of Yarn containers allocated to applications. The value is the same as the number of TaskManagers.	2
-ys	Number of TaskManager cores	2
-ynm	Custom name of an application on Yarn	test

Parameter	Description	Example Value
-c	Class of the program entry method (for example, the **main** or **getPlan()** method). This parameter is required only when the JAR file does not specify the class of its manifest.	com.bigdata.mrs.test

NOTE

For MRS 3.x or later, the **-yn** parameter is not supported.

Table 3-42 Service Parameter

Parameter	Description	Example Value
fs.obs.access.key	Key ID for accessing OBS	-
fs.obs.secret.key	Key (corresponding to the key ID) for accessing OBS	-

Step 7 Confirm job configuration information and click **OK**.

After the job is created, you can manage it.

----End

3.6.8 Viewing Job Configuration and Logs

This section describes how to view job configuration and logs.

Background

- You can view configuration information of all jobs.
- You can only view logs of running jobs.

  Because logs of Spark SQL and DistCp jobs are not in the background, you cannot view logs of running Spark SQL and DistCp jobs.

Procedure

Step 1 Log in to the MRS management console.

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

Step 3 Click **Jobs**.

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

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

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

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

Each tenant can submit and view 10 jobs concurrently.

----End

3.6.9 Stopping a Job

This section describes how to stop running MRS jobs.

Background

You cannot stop Spark SQL jobs. After a job is stopped, its status changes to **Terminated** and the job cannot be executed again.

Procedure

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**, select a running cluster, and click its name.

The cluster details page is displayed.

Step 3 Click **Jobs**.

Step 4 Select a running job, and choose **More > Stop** in the **Operation** column.

The job status changes from **Running** to **Terminated**.

----End

3.6.10 Deleting a Job

This section describes how to delete an MRS job. After a job is executed, you can delete it if you do not need to view its information.

Background

Jobs can be deleted one after another or in a batch. A deleted job cannot be restored. Therefore, exercise caution when deleting a job.

Procedure

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters**, select a running cluster, and click its name.

The cluster details page is displayed.

Step 3 Click **Jobs**.

Step 4 Choose **More > Delete** from the **Operation** in the row of the target job to be deleted. Enter **DELETE** in the **Delete Job** dialog box and click **OK**.

In this step, you can only delete one job only.

Step 5 If you select multiple jobs and click **Delete** on the upper left of the job list. Enter **DELETE** in the **Delete Job** dialog box and click **OK**.

You can delete one, multiple, or all jobs.

----End

3.6.11 Using Encrypted OBS Data for Job Running

In MRS 1.9.x encrypted data in OBS file systems can be used to run jobs, and the encrypted job running results can be stored in OBS file systems. Currently, data can be accessed only through an OBS protocol.

OBS supports data encryption and decryption using KMS keys. All encryption and decryption operations are performed on OBS, and keys are managed by DEW.

To use the OBS encryption function in MRS, you must have the KMS Administrator permissions and configure the following settings for the corresponding component:

NOTE

If the **OBS permission control** function is enabled in a cluster, the default agency **MRS_ECS_DEFAULT_AGENCY** configured on the ECS or the AK/SK of the custom agency is used for accessing OBS. OBS uses the received AK/SK to access DEW to obtain the KMS key status. Therefore, you need to bind the KMS Administrator policy to the used agency. Otherwise, OBS returns the "403 Forbidden" error when processing encrypted data. Currently, the KMS Administrator policy is bound to the agency **MRS_ECS_DEFAULT_AGENCY** by default. If you use a custom agency, you need to manually bind the policy to your custom agency.

Prerequisites

You have configured the function of accessing OBS from MRS first to use the OBS encryption function. For details, see **Configuring a Storage-Compute Decoupled Cluster (Agency)**.

Hive Configuration

Step 1 Log in to the MRS management console. In the navigation tree on the left, choose **Clusters** > **Active Clusters** and click the cluster name.

Step 2 Choose **Components** > **Hive** > **Service Configuration**.

Step 3 Switch **Basic** to **All**, and search for and set the following parameters:

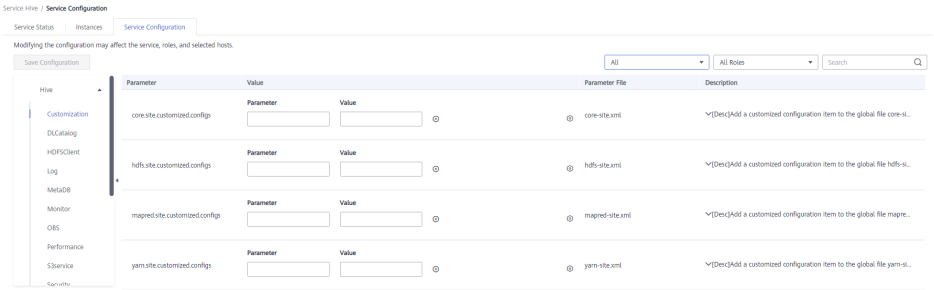

Table 3-43 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	• **SSE-KMS**: KMS keys are used for encryption and decryption • **NONE**: The encryption function is disabled.
fs.obs.server-side-encryption-key	-	(Optional) This parameter indicates an ID of the KMS key used for encryption. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. • **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. • **false**: The secure connection is disabled.

Step 4 Click **Save Configuration** and save the modified parameters as prompted.

----End

Hadoop Configuration

Method 1: Configuration on the GUI

Step 1 Log in to the MRS management console. In the navigation tree on the left, choose **Clusters** > **Active Clusters** and click the cluster name.

Step 2 Choose **Components** > **HDFS** > **Service Configuration**.

Step 3 Switch **Basic** to **All**, and search for and set the following parameters:

Table 3-44 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	• **SSE-KMS**: KMS keys are used for encryption and decryption • **NONE**: The encryption function is disabled.

Parameter	Value	Description
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. • **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. • **false**: The secure connection is disabled.

Step 4 Click **Save Configuration** and operate as prompted.

Step 5 Log in to the Master node as user **root**. The password is the password of user **root** you set when you create the cluster. If the cluster has multiple Master nodes, log in to each Master node and repeat **Step 5** to **Step 7**.

Step 6 Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

cd /opt/Bigdata/client

Step 7 Run the following command to update client configurations, and enter the username and password. The username is **admin**, and the password is the password of user **admin** you set when you create the cluster.

./ autoRefreshConfig.sh

----End

Method 2: Configuration Through the Client Configuration File

Add the following parameter settings to the client configuration file, for example, **/opt/Bigdata/client/HDFS/hadoop/etc/hadoop/core-site.xml**, on the Master node. If the cluster has multiple Master nodes, log in to each Master node and perform this operation.

Table 3-45 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	• **SSE-KMS**: KMS keys are used for encryption and decryption • **NONE**: The encryption function is disabled.

Parameter	Value	Description
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. ● **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. ● **false**: The secure connection is disabled.

HBase Configuration

Method 1: Configuration on the GUI

Step 1 Log in to the MRS management console. In the navigation tree on the left, choose **Clusters** > **Active Clusters** and click the cluster name.

Step 2 Choose **Components** > **HBase** > **Service Configuration**.

Step 3 Switch **Basic** to **All**, and search for and set the following parameters:

Table 3-46 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	● **SSE-KMS**: KMS keys are used for encryption and decryption ● **NONE**: The encryption function is disabled.
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.

Parameter	Value	Description
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. • **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. • **false**: The secure connection is disabled.

Step 4 Click **Save Configuration** and operate as prompted.

Step 5 Log in to the Master node as user **root**. The password is the password of user **root** you set when you create the cluster. If the cluster has multiple Master nodes, log in to each Master node and repeat **Step 5** to **Step 7**.

Step 6 Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

cd /opt/Bigdata/client

Step 7 Run the following command to update client configurations, and enter the username and password. The username is **admin**, and the password is the password of user **admin** you set when you create the cluster.

./ autoRefreshConfig.sh

----End

Method 2: Configuration Through the Client Configuration File

Add the following parameter settings to the client configuration file, for example, **/opt/Bigdata/client/HBase/hbase/conf/core-site.xml**, on the Master node. If the cluster has multiple Master nodes, log in to each Master node and perform this operation.

Table 3-47 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	• **SSE-KMS**: KMS keys are used for encryption and decryption • **NONE**: The encryption function is disabled.
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.

Parameter	Value	Description
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. • **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. • **false**: The secure connection is disabled.

Spark Configuration

Method 1: Configuration on the GUI

Step 1 Log in to the MRS management console. In the navigation tree on the left, choose **Clusters** > **Active Clusters** and click the cluster name.

Step 2 Choose **Components** > **Spark** > **Service Configuration**.

Step 3 Switch **Basic** to **All**, and search for and set the following parameters:

Table 3-48 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	• **SSE-KMS**: KMS keys are used for encryption and decryption • **NONE**: The encryption function is disabled.
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. • **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. • **false**: The secure connection is disabled.

Step 4 Click **Save Configuration** and operate as prompted.

Step 5 Log in to the Master node as user **root**. The password is the password of user **root** you set when you create the cluster. If the cluster has multiple Master nodes, log in to each Master node and repeat **Step 5** to **Step 7**.

Step 6 Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

cd /opt/Bigdata/client

Step 7 Run the following command to update client configurations, and enter the username and password. The username is **admin**, and the password is the password of user **admin** you set when you create the cluster.

./autoRefreshConfig.sh

----End

Method 2: Configuration Through the Client Configuration File

Add the following parameter settings to the client configuration file, for example, **/opt/Bigdata/client/Spark/spark/conf/core-site.xml**, on the Master node. If the cluster has multiple Master nodes, log in to each Master node and perform this operation.

Table 3-49 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	• **SSE-KMS**: KMS keys are used for encryption and decryption • **NONE**: The encryption function is disabled.
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. • **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. • **false**: The secure connection is disabled.

Presto Configuration

Step 1 Log in to the MRS management console. In the navigation tree on the left, choose **Clusters** > **Active Clusters** and click the cluster name.

Step 2 Choose **Components** > **Presto** > **Service Configuration**.

Step 3 Switch **Basic** to **All**, and search for and set the following parameters:

Table 3-50 Data encryption parameters

Parameter	Value	Description
fs.obs.server-side-encryption-type	SSE-KMS	• **SSE-KMS**: KMS keys are used for encryption and decryption • **NONE**: The encryption function is disabled.
fs.obs.server-side-encryption-key	-	ID of the KMS key used for encryption. This parameter is optional. If **fs.obs.server-side-encryption-type** is set to **SSE-KMS** and this parameter is not set, OBS uses the default KMS key for encryption.
fs.obs.connection.ssl.enabled	true	Whether to establish a secure connection with OBS. • **true**: The secure connection is enabled. To use OBS encryption and decryption, this parameter must be set to **true**. • **false**: The secure connection is disabled.

Step 4 Click **Save Configuration** and operate as prompted.

----End

3.6.12 Configuring Job Notification Rules

MRS uses SMN to offer a publish/subscribe model to achieve one-to-multiple message subscriptions and notifications in a variety of message types (SMSs and emails). You can configure job notification rules to receive notifications immediately upon a job execution success or failure.

Procedure

Step 1 Log in to the management console.

Step 2 Click **Service List**. Under **Management & Governance**, click **Simple Message Notification**.

Step 3 Create a topic and add subscriptions to the topic. For details, see **Configuring Message Notification**.

Step 4 Go to the MRS management console, and click the cluster name to go to the cluster details page.

Step 5 Click the **Alarms** tab, and choose **Notification Rules** > **Add Notification Rule**.

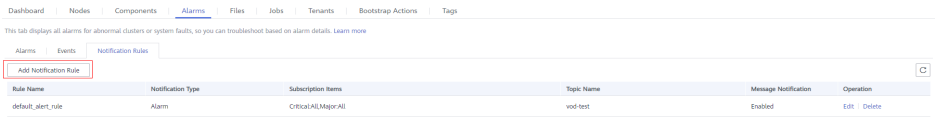

Step 6 Configure a notification rule for sending job execution results to subscribers.

Table 3-51 Parameters of adding a notification rule

Parameter	Description
Rule Name	User-defined notification rule name. Only digits, letters, hyphens (-), and underscores (_) are allowed.
Message Notification	If you enable this function, subscription messages will be sent to subscribers.
Topic Name	Select an existing topic or click **Create Topic** to create a topic.
Notification Type	Select **Event**.
Subscription Items	1. Click ⊕ next to **Suggestion**. 2. Click ⊕ next to **Manager**. 3. Select **Job Running Succeeded** and **Job Running Failed**.

----End

3.7 Component Management

3.7.1 Object Management

MRS contains different types of basic objects. **Table 3-52** describes these objects.

Table 3-52 MRS basic object overview

Object	Description	Example
Service	Function set that can complete specific business.	KrbServer service and LdapServer service
Service instance	Specific instance of a service, usually called service.	KrbServer service

Object	Description	Example
Service role	Function entity that forms a complete service, usually called role.	KrbServer is composed of the KerberosAdmin role and KerberosServer role.
Role instance	Specific instance of a service role running on a host.	KerberosAdmin that is running on Host2 and KerberosServer that is running on Host3
Host	An ECS running Linux OS.	Host1 to Host5
Rack	Physical entity that contains multiple hosts connecting to the same switch.	Rack1 contains Host1 to Host5.
Cluster	Logical entity that consists of multiple hosts and provides various services.	Cluster1 cluster consists of five hosts (Host1 to Host5) and provides services such as KrbServer and LdapServer.

3.7.2 Viewing Configuration

On MRS, you can view the configuration of services (including roles) and role instances.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-30 Synchronizing IAM users

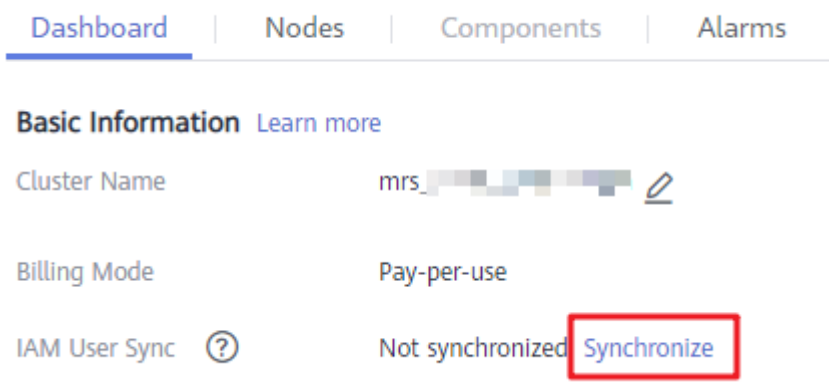

Procedure

- Query service configuration.
  a. On the cluster details page, click the **Components** tab.

Figure 3-31 Components tab page (using MRS 1.9.2 as an example)

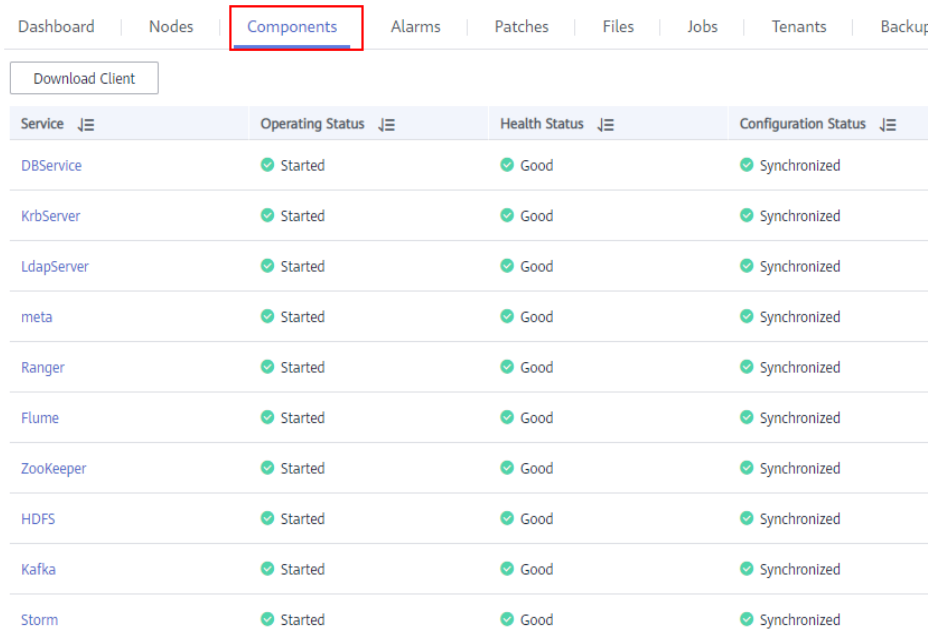

b. Select the target service from the service list.

c. Click **Service Configuration**.

d. Switch **Basic** to **All**. All configuration parameters of the service are displayed in the navigation tree. The service name and role names are displayed from upper to lower in the navigation tree.

Figure 3-32 All configurations

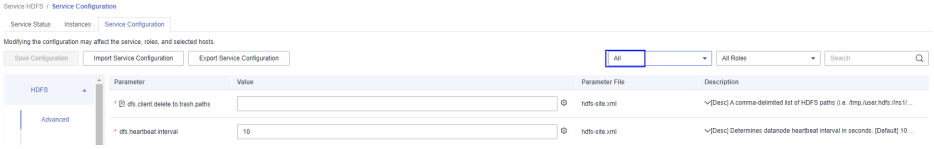

e. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

The parameters under the service nodes and role nodes are service configuration parameters and role configuration parameters respectively.

f. If you select **Non-default** from the **--Select--** drop-down list, parameters whose values are not default ones are displayed on the page. (This value option is available in versions earlier than MRS 3.*x*.)

- Query role instance configurations.

a. On the MRS cluster details page, click **Components**.

Figure 3-33 Components tab page (using MRS 1.9.2 as an example)

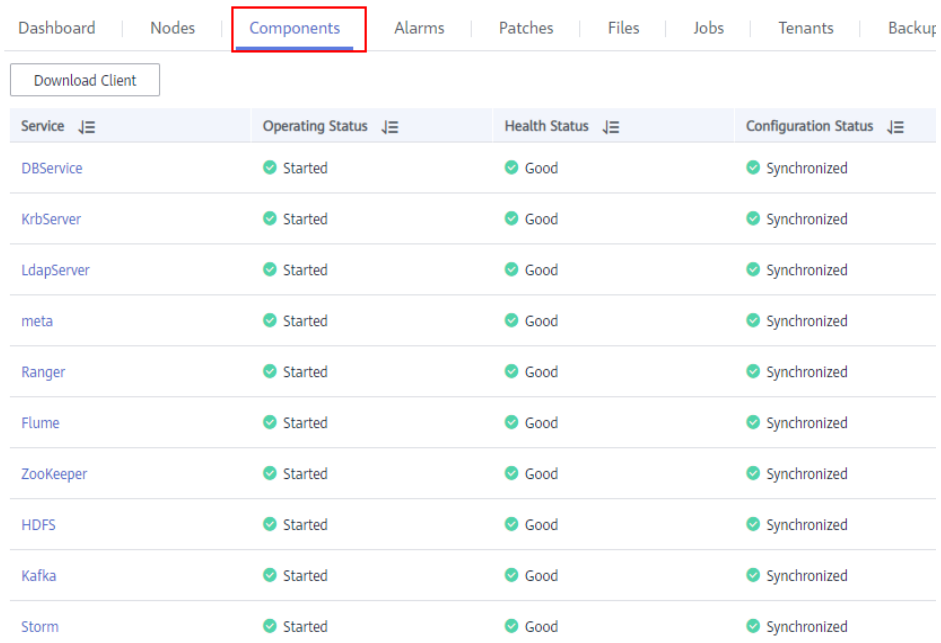

b. Select the target service from the service list.

c. Click the **Instances** tab.

d. Click the target role instance from the role instance list.

e. Click **Instance Configuration**.

f. Switch **Basic** to **All** on the right of the page. All configuration parameters of the role instance are displayed in the navigation tree.

g. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

h. If you select **Non-default** from the **--Select--** drop-down list, parameters whose values are not default ones are displayed on the page. (This value option is available in versions earlier than MRS 3.*x*.)

3.7.3 Managing Services

You can perform the following operations on MRS:

- Add or delete services. This is available only for custom clusters of MRS 3.1.2 and later versions.
- Start the service in the **Stopped**, **Stop Failed**, or **Failed to Start** state to use the service.
- Stop the services or stop abnormal services.
- Restart abnormal services or configure expired services to restore or enable the services.

Prerequisites

- **You have configured permissions for the user group to which the IAM users belong.**

Adding or deleting a service in a cluster is a high-risk operation. Bind the MRS FullAccess, MRS Administrator, Server Administrator, Tenant Guest, MRS Administrator, or Tenant Administrator policy to the user group before you perform this operation. For details about the permissions, see **Synchronizing IAM Users to MRS**.

- **You have synchronized IAM users.** (On the **Dashboard** tab page, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.)

Figure 3-34 Synchronizing IAM users

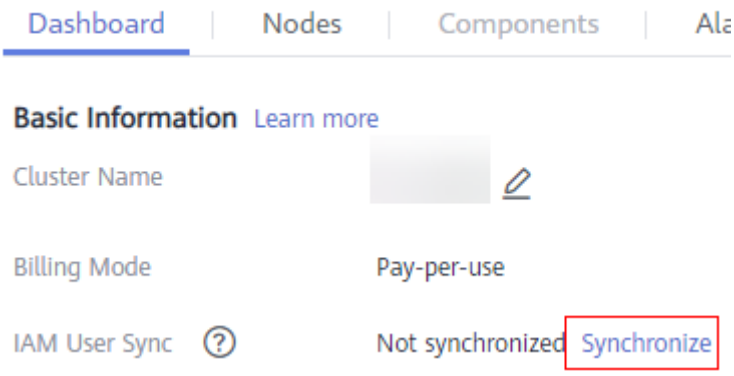

Impact on the System

- The stateful component cannot be added to the task node group.

Adding a Service

NOTE

Services can be added and deleted in MRS 3.1.2-LTS.3 or later.

Step 1 On the cluster details page, Choose **Components** and click **Add Service**.

Step 2 In the service list, select the services to be added and click **Next**.

NOTE

- When you add a service, the underlying services on which the service depends are automatically selected. You can add multiple services at the same time.
- You can add a service only on a node in normal state.
- If you add Hadoop to a cluster with no Hadoop before, you need to refresh the cluster details page on the MRS console and synchronize IAM users so that jobs can be successfully submitted.
- A single component of the Hadoop service cannot be added to the cluster. Only the Hadoop service can be added. The Hadoop service includes MapReduce, Yarn, and HDFS.
- After the Spark2x component is added, if you need to operate SparkSQL on the Hue web UI, restart the Hue service first.

Step 3 On the **Topology Adjustment** page, select the nodes where the service is to be deployed. For details about the deployment scheme, see **Table 2-9**.

Step 4 Click **OK**. After the service is added, you can view the added service on the **Components** page.

NOTE

The services added on the console are automatically synchronized to Manager.

----End

Deleting a Service

NOTE

Services can be added and deleted in MRS 3.1.2-LTS.3 or later.

Step 1 On the cluster details page, click **Components**.

Step 2 Locate the row that contains the target service and click **Delete**.

NOTE

- If the service to be deleted has upper-layer dependencies, the service cannot be deleted. Only one service can be deleted at a time.
- You can delete installed services except Hadoop (HDFS, Yarn, and MapReduce), Ranger, DBService, KrbServer, LdapServer, and meta services.

Step 3 Enter **DELETE** in the displayed **Delete Service** dialog box and click **OK** to confirm the deletion.

CAUTION

- The services deleted on the console are automatically synchronized to Manager.
- Before deleting a service, back up the service data to prevent data loss.

----End

Starting, Stopping, and Restarting a Service

Step 1 On the MRS cluster details page, click **Components**.

Step 2 Locate the row that contains the target service, **Start**, **Stop**, and **Restart** to start, stop, or restart the service.

Services are interrelated. If a service is started, stopped, and restarted, services dependent on it will be affected.

The services will be affected in the following ways:

- If a service is to be started, the lower-layer services dependent on it must be started first.
- If a service is stopped, the upper-layer services dependent on it are unavailable.
- If a service is restarted, the running upper-layer services dependent on it must be restarted.

----End

3.7.4 Configuring Service Parameters

On the MRS console, you can view and modify the default service configurations based on site requirements and export or import the configurations.

Impact on the System

- You need to download and update the client configuration files after configuring HBase, HDFS, Hive, Spark, Yarn, and MapReduce service properties.
- The parameters of DBService cannot be modified when only one DBService role instance exists in the cluster.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-35 Synchronizing IAM users

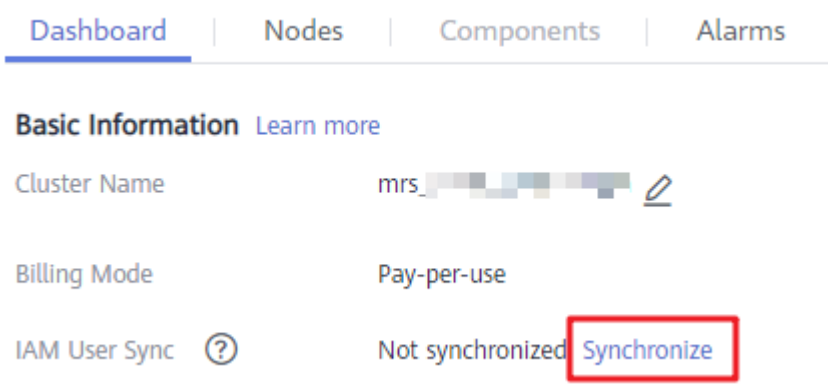

Modifying Service Parameters

1. On the MRS cluster details page, click **Components**.

Figure 3-36 Components tab page (using MRS 1.9.2 as an example)

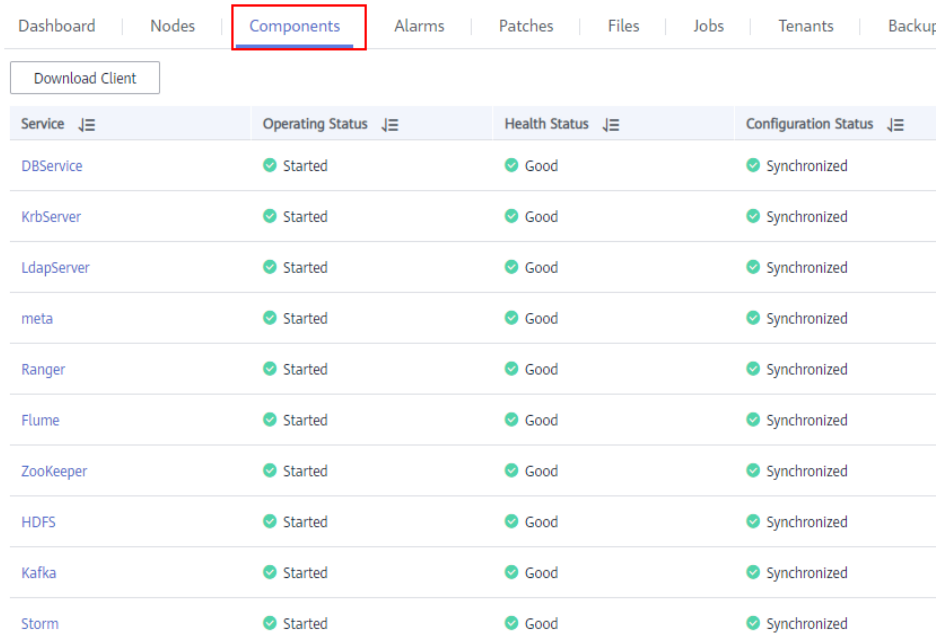

2. Select the target service from the service list.
3. Click **Service Configuration**.
4. Switch **Basic** to **All**. All configuration parameters of the service are displayed in the navigation tree. The service name and role names are displayed from upper to lower in the navigation tree.
5. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

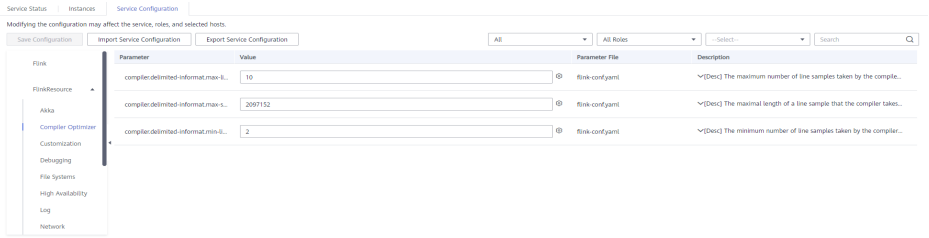

If you want to cancel the modification of a parameter value, click to restore it.

6. Click **Save Configuration**, save the parameters as prompted, and restart the service.

NOTE

In versions earlier than MRS 3.*x*, to update the queue configuration of YARN without restarting the service, choose **More** > **Refresh Queue** on the **Service Status** tab page.

3.7.5 Configuring Customized Service Parameters

Each component of MRS supports all open-source parameters. MRS supports the modification of some parameters for key application scenarios. Some component clients may not include all parameters with open-source features. To modify the

component parameters that are not directly supported by MRS, you can add new parameters for components by using the configuration customization function on MRS. Newly added parameters are saved in component configuration files and take effect after restart.

Impact on the System

- After the service attributes are configured, the service needs to be restarted. The service cannot be accessed during restart.
- You need to download and update the client configuration files after configuring HBase, HDFS, Hive, Spark, Yarn, and MapReduce service properties.

Prerequisites

- You have understood the meanings of parameters to be added, configuration files that have taken effect, and the impact on components.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

  Figure 3-37 Synchronizing IAM users

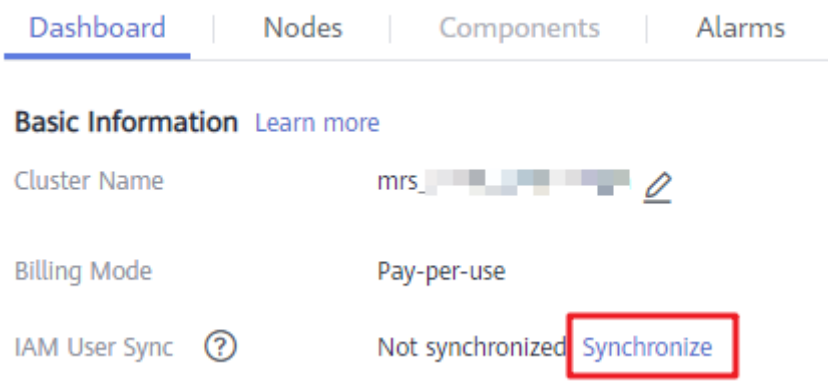

Procedure

Step 1 On the MRS cluster details page, click **Components**.

Figure 3-38 Components tab page (using MRS 1.9.2 as an example)

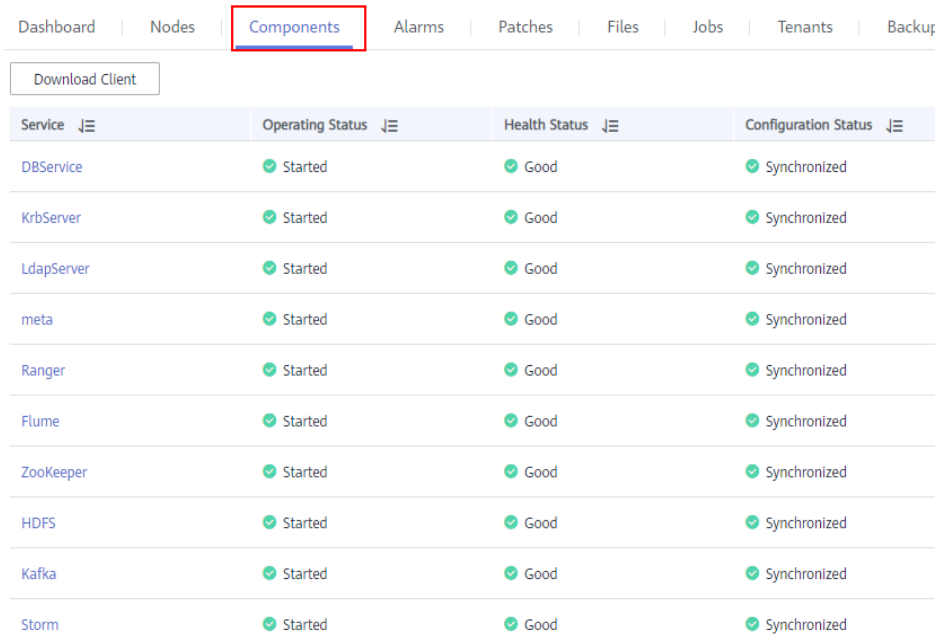

Step 2 Select the target service from the service list.

Step 3 Click **Service Configuration**.

Step 4 In the configuration type drop-down box on the right side, switch **Basic** to **All**.

Figure 3-39 All configurations

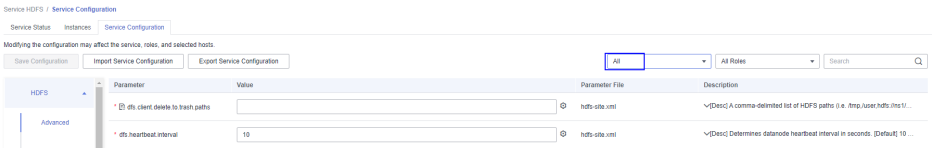

Step 5 In the navigation tree, select **Customization**. The customized parameters of the current component are displayed on MRS.

The configuration files that save the newly added customized parameters are displayed in the **Parameter File** column. Different configuration files may have same open-source parameters. After the parameters in different files are set to different values, whether the configuration takes effect depends on the loading sequence of the configuration files by components. You can customize parameters for services and roles as required. Adding customized parameters for a single role instance is not supported.

Step 6 Based on the configuration files and parameter functions, locate the row where a specified parameter resides, enter the parameter name supported by the component in the **Parameter** column and enter the parameter value in the **Value** column.

- You can click or to add or delete a custom parameter. You can delete a customized parameter only after you click for the first time.
- If you want to cancel the modification of a parameter value, click to restore it.

Step 7 Click **Save Configuration** and operate as prompted.

----End

Task Example

Configuring Customized Hive Parameters

Hive depends on HDFS. By default, Hive accesses the HDFS client. The configuration parameters to take effect are controlled by HDFS in a unified manner. For example, the HDFS parameter **ipc.client.rpc.timeout** affects the RPC timeout period for all clients to connect to the HDFS server. If you need to modify the timeout period for Hive to connect to HDFS, you can use the configuration customization function. After this parameter is added to the **core-site.xml** file of Hive, this parameter can be identified by the Hive service and its configuration overwrites the parameter configuration in HDFS.

Step 1 On the MRS cluster details page, click **Components**.

Step 2 Choose **Hive** > **Service Configuration**.

Step 3 In the configuration type drop-down box on the right side, switch **Basic** to **All**.

Step 4 In the navigation tree on the left, select **Customization** for the Hive service. The system displays the customized service parameters supported by Hive.

Step 5 In **core-site.xml**, locate the row that contains the **core.site.customized.configs** parameter, enter **ipc.client.rpc.timeout** in the **Parameter** column, and enter a new value in the **Value** column, for example, **150000**. The unit is millisecond.

Figure 3-40 Configuring custom parameters (using MRS 1.9.2 as an example)

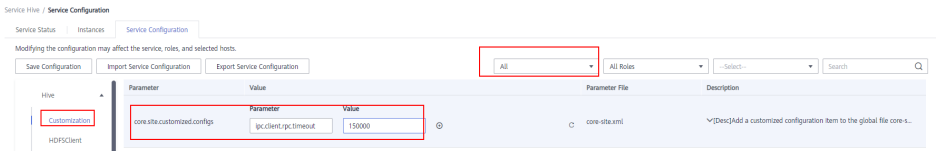

Step 6 Click **Save Configuration** and operate as prompted.

Figure 3-41 Saving custom configurations

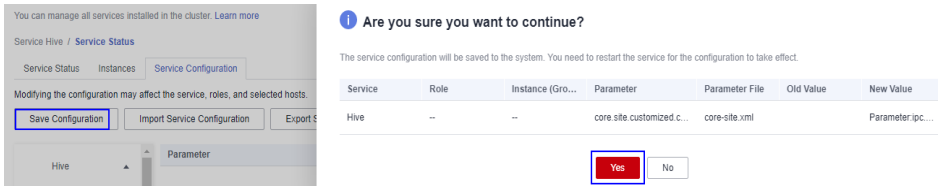

----End

3.7.6 Synchronizing Service Configuration

Scenario

If **Configuration Status** of some services is **Configuration expired** or **Configuration failed**, synchronize configuration for the cluster or service to

restore its configuration status. If all services in the cluster are in the **Configuration failed** state, synchronize the cluster configuration with the background configuration.

Impact on the System

After synchronizing service configurations, you need to restart the services whose configurations have expired. These services are unavailable during restart.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-42 Synchronizing IAM users

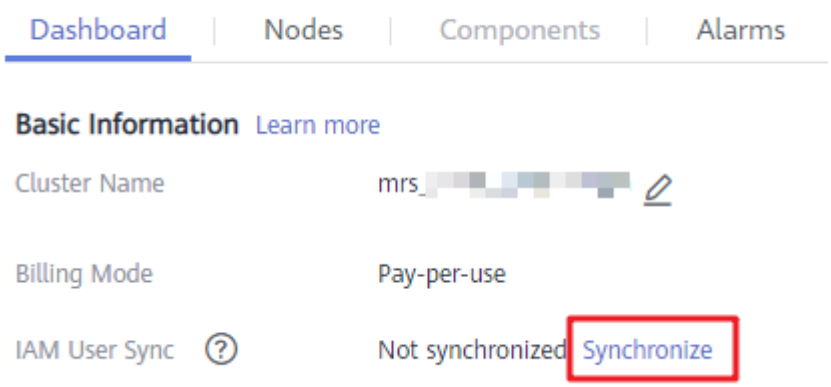

Procedure

Step 1 On the MRS cluster details page, click **Components**.

Figure 3-43 Components tab page (using MRS 1.9.2 as an example)

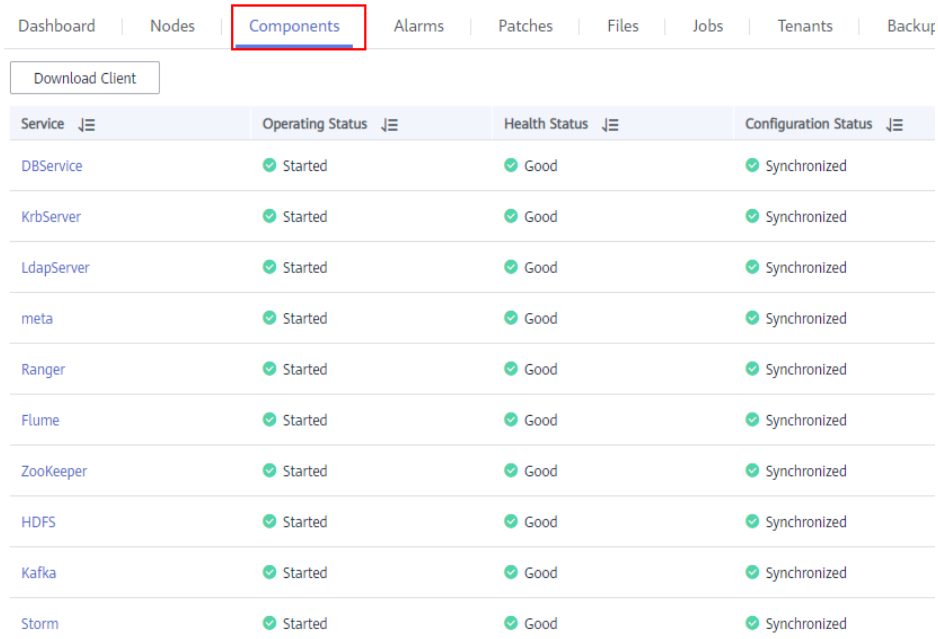

Step 2 Select the target service from the service list.

Step 3 In the **Service Status** tab, choose **More** > **Synchronize Configuration** and operate as prompted.

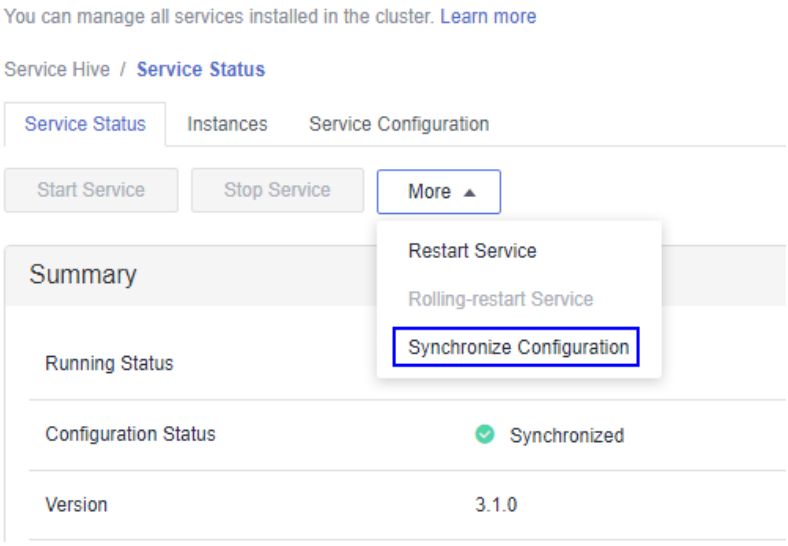

----End

3.7.7 Managing Role Instances

Scenario

You can start a role instance that is in the **Stopped**, **Failed to stop** or **Failed to start** status, stop an unused or abnormal role instance or restart an abnormal role instance to recover its functions.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-44 Synchronizing IAM users

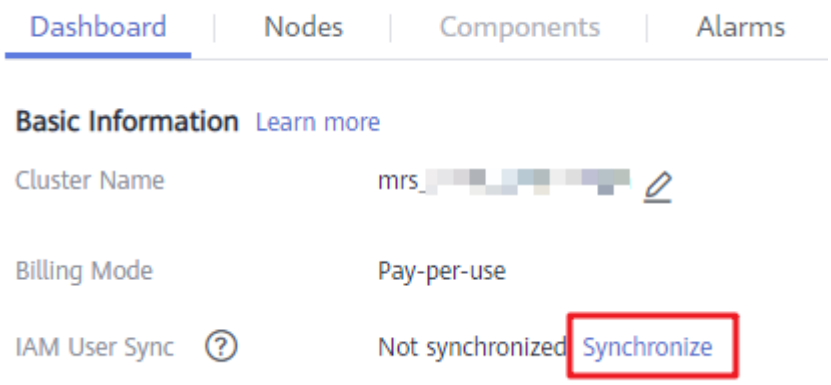

Procedure

Step 1 On the cluster details page, click the **Components** tab.

Figure 3-45 Components tab page (using MRS 1.9.2 as an example)

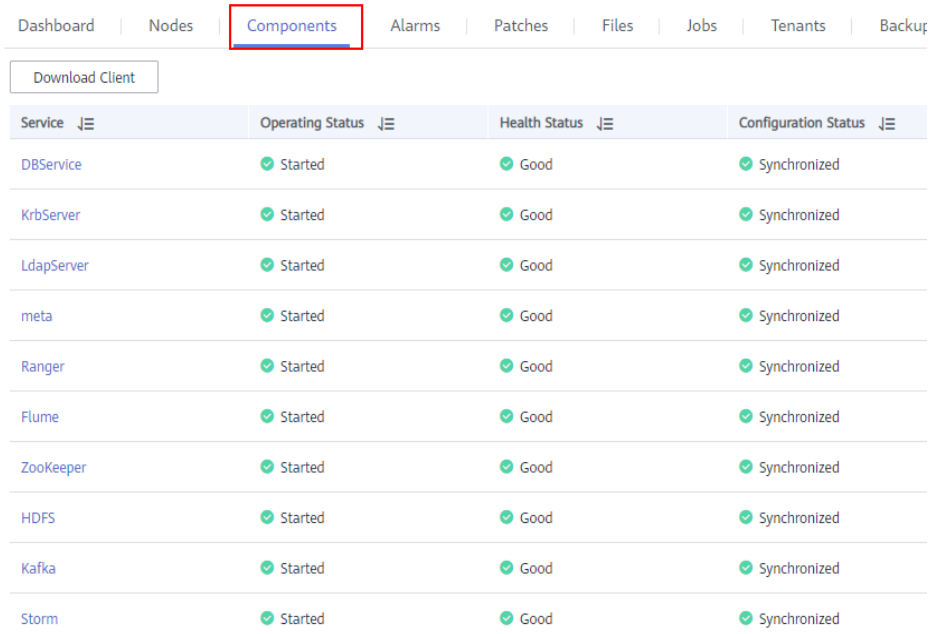

Step 2 Select the target service from the service list.

Step 3 Click the **Instances** tab.

Step 4 Select the check box on the left of the target role instance.

Step 5 Click **More**, select operations such as **Start Instance**, **Stop Instance**, **Restart Instance**, **Rolling-restart Instance**, or **Delete Instance** based on site requirements.

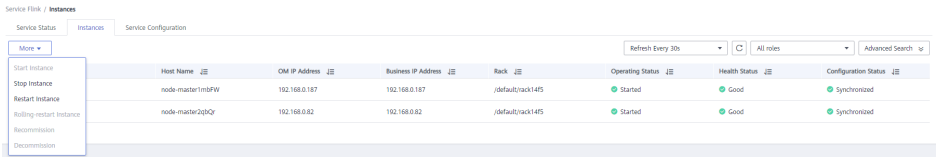

----End

3.7.8 Configuring Role Instance Parameters

Scenario

You can view and modify default role instance configuration on MRS based on site requirements. The configurations can be imported and exported.

Impact on the System

You need to download and update the client configuration files after configuring HBase, HDFS, Hive, Spark, Yarn, and MapReduce service properties.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-46 Synchronizing IAM users

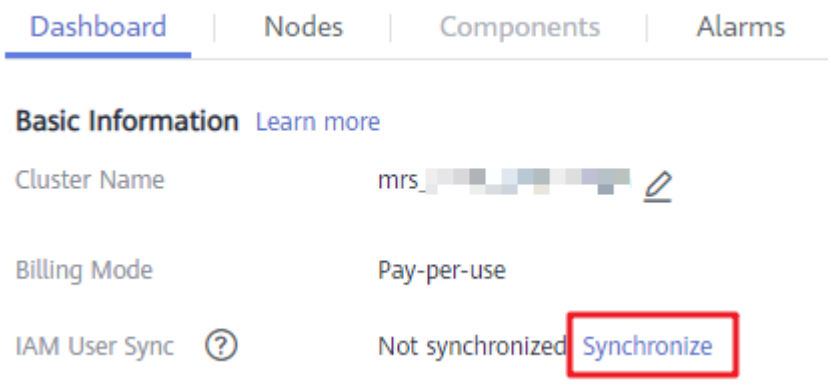

Modifying Role Instance Parameters

1. On the cluster details page, click the **Components** tab.

Figure 3-47 Components tab page (using MRS 1.9.2 as an example)

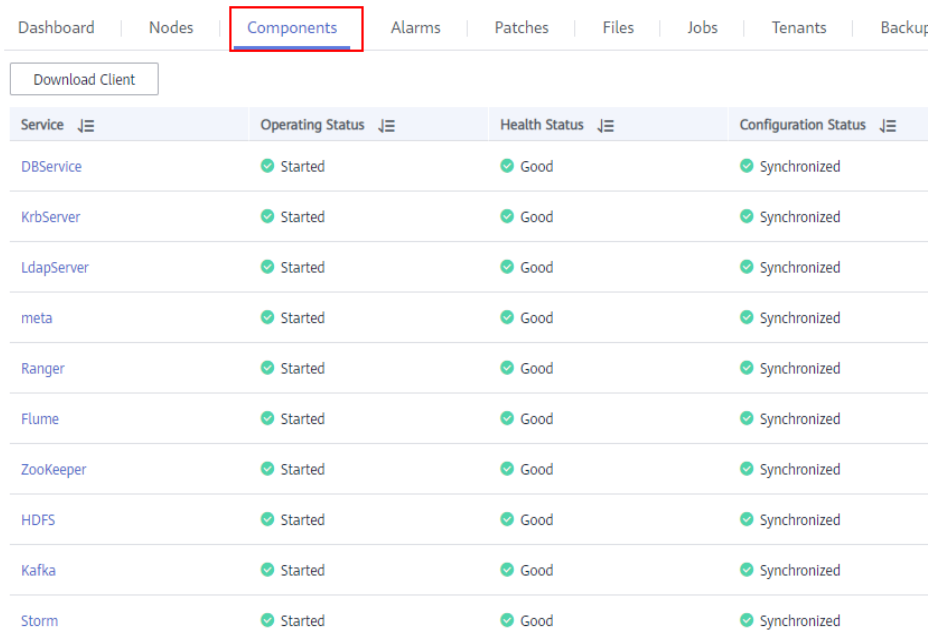

2. Select the target service from the service list.
3. Click the **Instances** tab.
4. Click the target role instance from the role instance list.
5. Click the **Instance Configuration** tab.
6. Switch **Basic** to **All** from the drop-down list on the right of the page. All configuration parameters of the role instance are displayed in the navigation tree.
7. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

   If you want to cancel the modification of a parameter value, click to restore it.
8. Click **Save Configuration** and operate as prompted.

3.7.9 Synchronizing Role Instance Configuration

Scenario

When **Configuration Status** of a role instance is **Configuration expired** or **Configuration failed**, you can synchronize the configuration data of the role instance with the background configuration.

Impact on the System

After synchronizing a role instance configuration, you need to restart the role instance whose configuration has expired. The role instance is unavailable during restart.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-48 Synchronizing IAM users

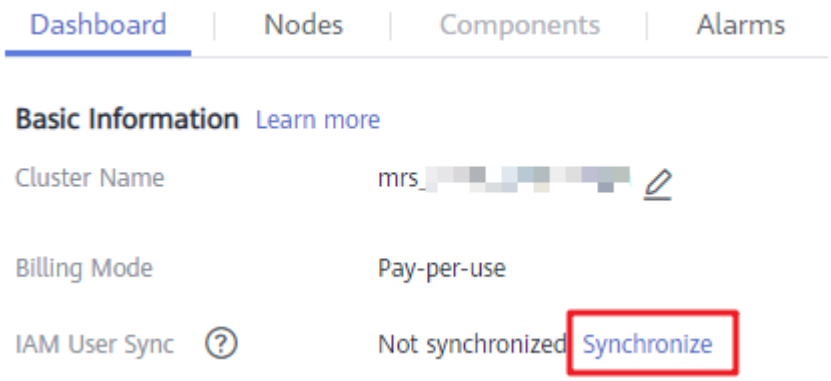

Procedure

Step 1 On the cluster details page, click the **Components** tab.

Figure 3-49 Components tab page (using MRS 1.9.2 as an example)

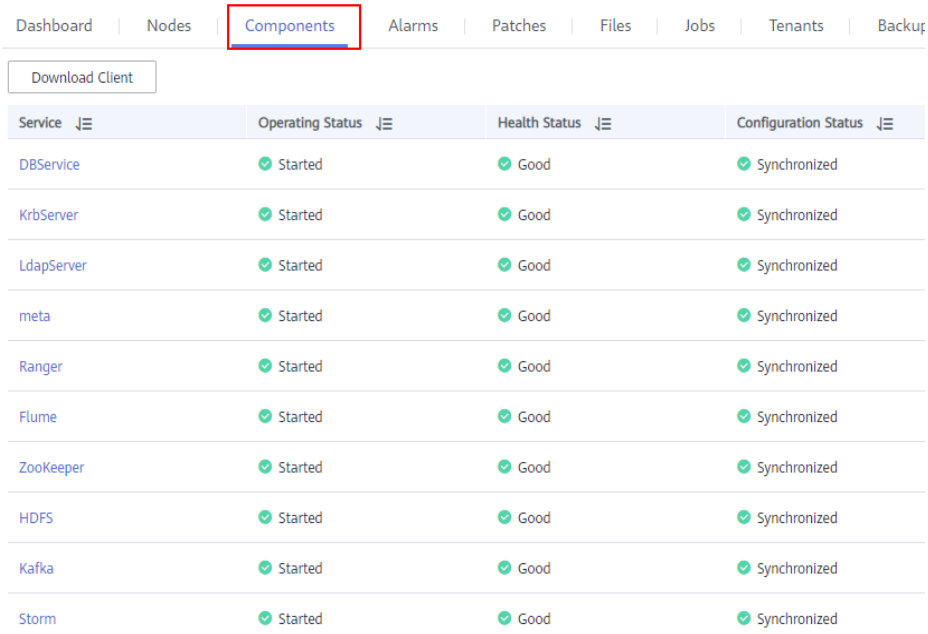

Step 2 Select a service name.

Step 3 Click the **Instances** tab.

Step 4 Click the target role instance from the role instance list.

Step 5 Click **More** and select **Synchronize Configuration** above the role instance status and indicator information.

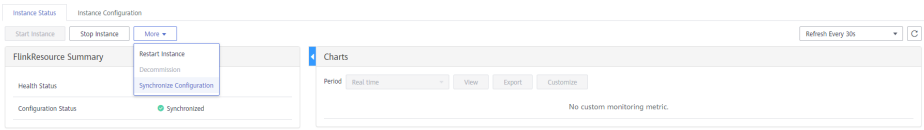

Step 6 In the dialog box that is displayed, select **Restart the service or instances whose configurations have expired** and click **Yes** to restart the role instance.

----End

3.7.10 Decommissioning and Recommissioning a Role Instance

Scenario

If a Core or Task node is faulty, the cluster status may be displayed as **Abnormal**. In an MRS cluster, data can be stored on different Core nodes. You can decommission the specified role instance on MRS to stop the role instance from providing services. After fault rectification, you can recommission the role instance.

The following role instances can be decommissioned or recommissioned:

- DataNode role instance on HDFS
- NodeManager role instance on Yarn
- RegionServer role instance on HBase
- ClickHouseServer role instance on ClickHouse

  NOTE

  ClickHouseServer role instances can be decommissioned only in MRS 3.1.2 or later.

Restrictions:

- If the number of the DataNodes is less than or equal to that of HDFS copies, decommissioning cannot be performed. If the number of HDFS copies is three and the number of DataNodes is less than four in the system, decommissioning cannot be performed. In this case, an error will be reported and force MRS to exit the decommissioning 30 minutes after MRS attempts to perform the decommissioning.
- If a role instance is out of service, you must recommission the instance to start it before using it again.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-50 Synchronizing IAM users

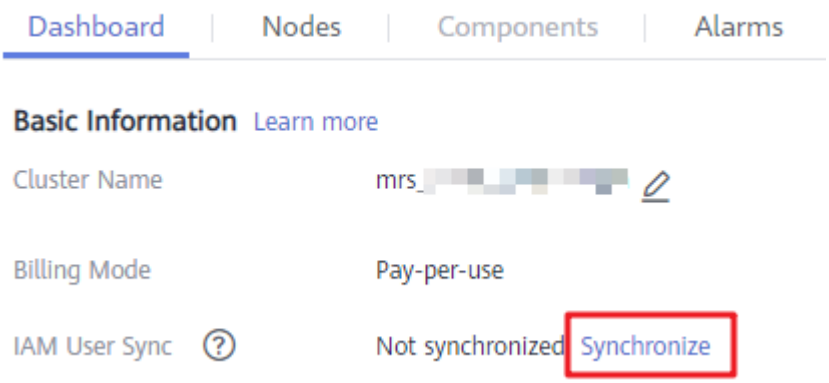

Procedure

Step 1 On the cluster details page, click the **Components** tab.

Figure 3-51 Components tab page (using MRS 1.9.2 as an example)

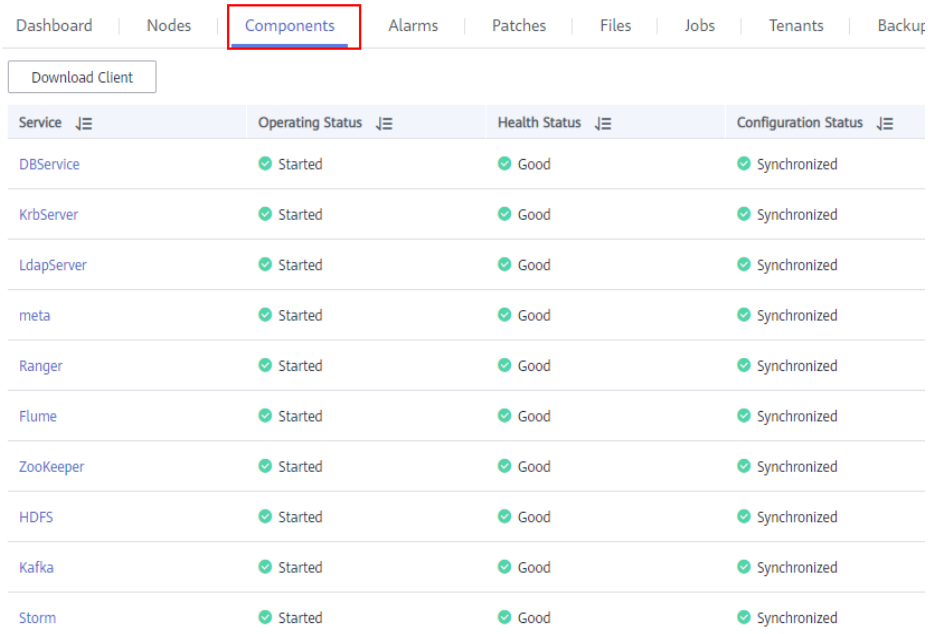

Step 2 Click a service in the service list.

Step 3 Click the **Instances** tab.

Step 4 Select an instance.

Step 5 Choose **More** > **Decommission** or **Recommission** to perform the corresponding operation.

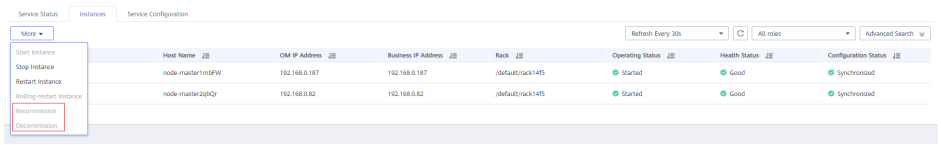

NOTE

During the instance decommissioning, if the service corresponding to the instance is restarted in the cluster using another browser, MRS displays a message indicating that the instance decommissioning is stopped, but the **Operating Status** of the instance is displayed as **Started**. In this case, the instance has been decommissioned on the background. You need to decommission the instance again to synchronize the operating status.

----End

3.7.11 Starting and Stopping a Cluster

A cluster is a collection of service components. You can start or stop all services in a cluster.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-52 Synchronizing IAM users

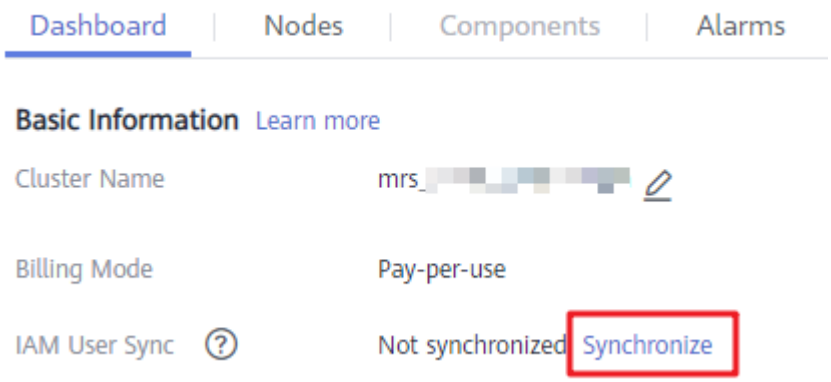

Procedure

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

3.7.12 Synchronizing Cluster Configuration

Scenario

If **Configuration Status** of all services or some services is **Configuration expired** or **Configuration failed**, synchronize configuration for the cluster or service to restore its configuration status.

- If all services in the cluster are in the **Configuration failed** status, synchronize the cluster configuration with the background configuration.
- If all services in the cluster are in the **Configuration failed** status, synchronize the service configuration with the background configuration.

NOTE

In **MRS 3.x**, you cannot perform operations in this section on the management console.

Impact on the System

After synchronizing cluster configurations, you need to restart the services whose configurations have expired. These services are unavailable during restart.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-53 Synchronizing IAM users

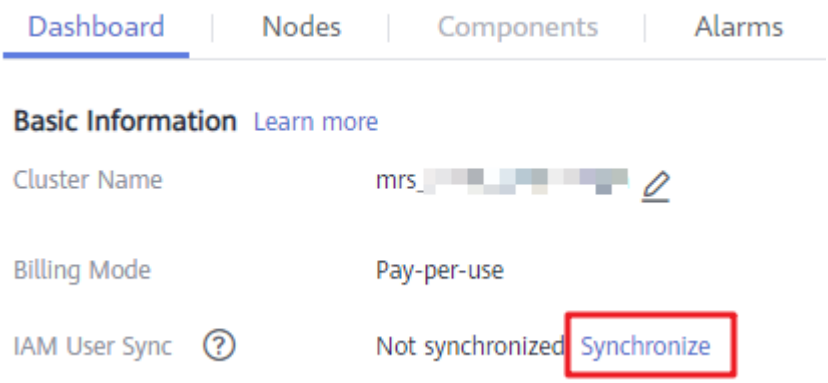

Procedure

Step 1 On the cluster details page, choose **Configuration** > **Synchronize Configuration** in the upper right corner.

Figure 3-54 Synchronizing configurations (using MRS 1.9.2 as an example)

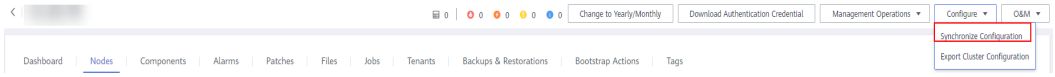

Step 2 In the displayed dialog box, select **Restart services and instances whose configuration have expired**, and click **OK** to restart the service whose configuration has expired.

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

----End

3.7.13 Exporting Cluster Configuration

Scenario

You can export all configuration data of a cluster using MRS to meet site requirements. The exported configuration data is used to rapidly update service configuration.

NOTE

In **MRS 3.x**, you cannot perform operations in this section on the management console.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-55 Synchronizing IAM users

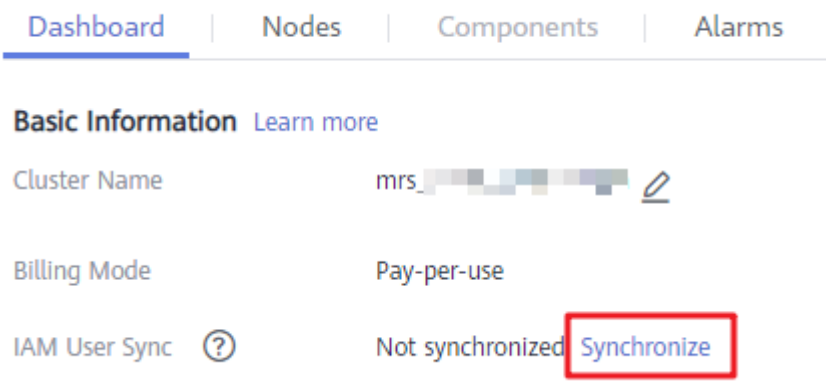

Procedure

On the cluster details page, choose **Configuration** > **Export Cluster Configuration** in the upper right corner.

Figure 3-56 Exporting cluster configurations (using MRS 1.9.2 as an example)

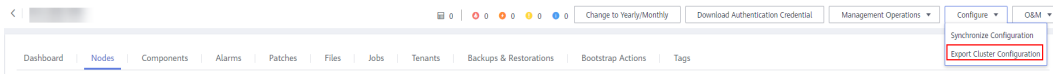

The exported file is used to update service configurations. For details, see **Importing Service Configuration Parameters** in **Configuring Service Parameters**.

3.7.14 Performing Rolling Restart

After modifying the configuration items of a big data component, you need to restart the corresponding service to make new configurations take effect. If you use a normal restart mode, all services or instances are restarted concurrently,

which may cause service interruption. To ensure that services are not affected during service restart, you can restart services or instances in batches by rolling restart. For instances in active/standby mode, a standby instance is restarted first and then an active instance is restarted. Rolling restart takes longer than normal restart.

Table 3-53 provides services and instances that support or do not support rolling restart in the MRS cluster.

Table 3-53 Services and instances that support or do not support rolling restart

Service	Instance	Support Rolling Restart
Alluxio	AlluxioJobMaster	Yes
	AlluxioMaster	
ClickHouse	ClickHouseServer	Yes
	ClickHouseBalancer	
CDL	CDLConnector	Yes
	CDLService	
Flink	FlinkResource	No
	FlinkServer	
Flume	Flume	Yes
	MonitorServer	
Guardian	TokenServer	Yes
HBase	HMaster	Yes
	RegionServer	
	ThriftServer	
	RESTServer	
HetuEngine	HSBroker	Yes
	HSConsole	
	HSFabric	
	QAS	
HDFS	NameNode	Yes
	Zkfc	
	JournalNode	
	HttpFS	
	DataNode	

Service	Instance	Support Rolling Restart
Hive	MetaStore	Yes
	WebHCat	
	HiveServer	
Hue	Hue	No
Impala	Impalad	No
	StateStore	
	Catalog	
IoTDB	IoTDBServer	Yes
Kafka	Broker	Yes
	KafkaUI	No
Kudu	KuduTserver	Yes
	KuduMaster	
Loader	Sqoop	No
MapReduce	JobHistoryServer	Yes
Oozie	oozie	No
Presto	Coordinator	Yes
	Worker	
Ranger	RangerAdmin	Yes
	UserSync	
	TagSync	
Spark	JobHistory	Yes
	JDBCServer	
	SparkResource	
Storm	Nimbus	Yes
	UI	
	Supervisor	
	Logviewer	
Tez	TezUI	No
Yarn	ResourceManager	Yes
	NodeManager	

Service	Instance	Support Rolling Restart
Zookeeper	Quorumpeer	Yes

Restrictions

- Perform a rolling restart during off-peak hours.
  - Otherwise, a rolling restart failure may occur. For example, if the throughput of Kafka is high (over 100 MB/s) during the Kafka rolling restart, the Kafka rolling restart may fail.
  - For example, if the requests per second of each RegionServer on the native interface exceed 10,000 during the HBase rolling restart, you need to increase the number of handles to prevent a RegionServer restart failure caused by heavy loads during the restart.
- Before the restart, check the number of current requests of HBase. If the number of requests of each RegionServer on the native interface exceeds 10,000, increase the number of handles to prevent a failure.
- If the number of Core nodes in a cluster is less than six, services may be affected for a short period of time.
- Preferentially perform a rolling instance or service restart and select **Only restart instances whose configurations have expired**.

Performing a Rolling Service Restart

Step 1 Choose **Clusters** > **Active Clusters** and click a cluster name to go to the cluster details page.

Step 2 Click **Components** and select a service for which you want to perform a rolling restart.

Step 3 On the **Service Status** tab page, click **More** and select **Rolling-restart Service**.

Figure 3-57 Service status (MRS 1.9.2 is used as an example)

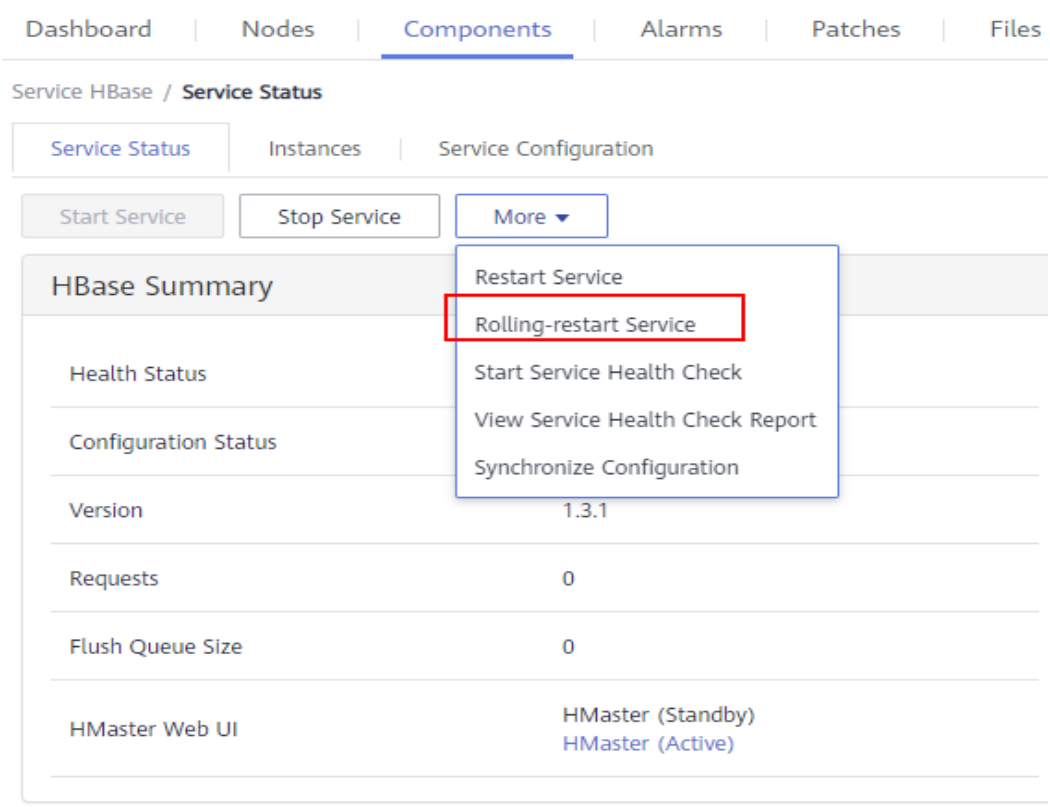

Step 4 The **Rolling-restart Service** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the service.

Figure 3-58 Performing a rolling service restart

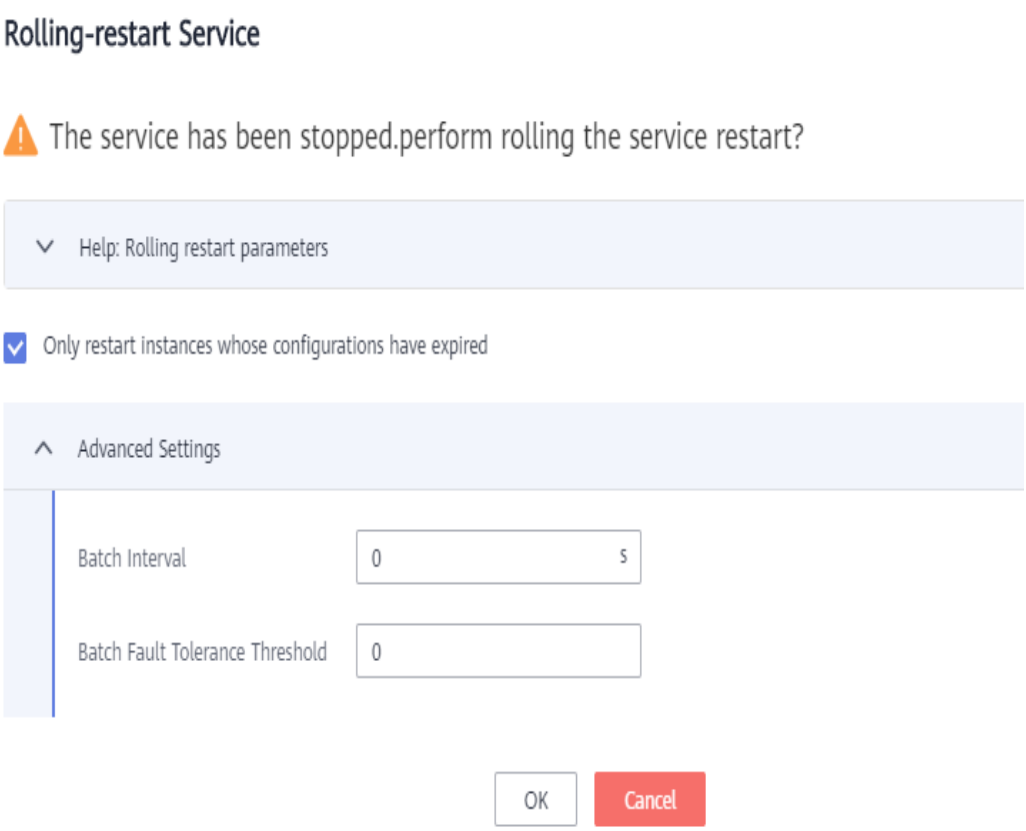

Step 5 After the rolling restart task is complete, click **Finish**.

Figure 3-59 Finishing the rolling service restart

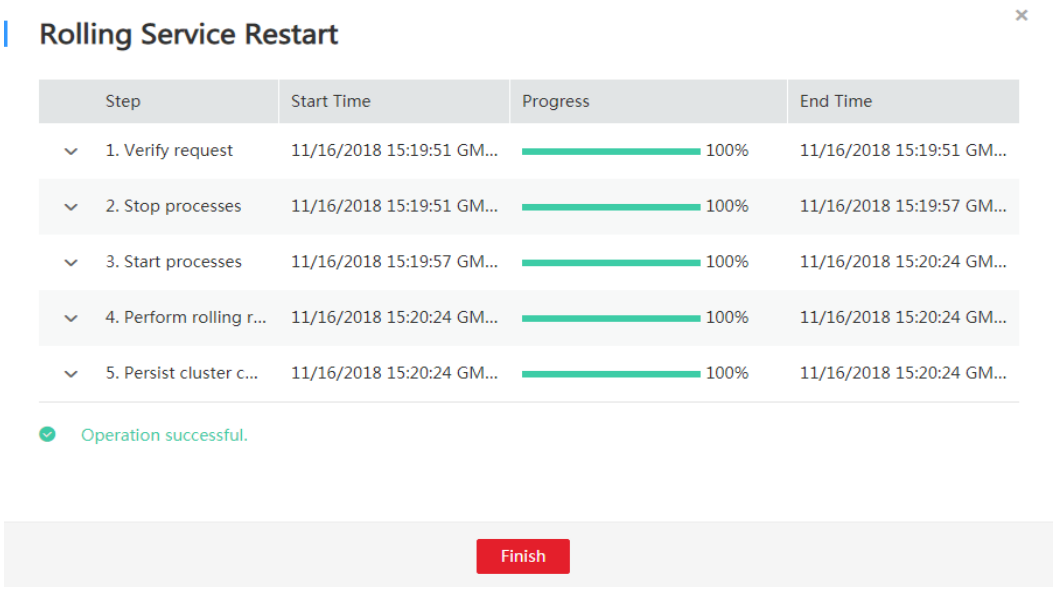

----End

Performing a Rolling Instance Restart

Step 1 Choose **Clusters** > **Active Clusters** and click a cluster name to go to the cluster details page.

Step 2 Click **Components** and select a service for which you want to perform a rolling restart.

Step 3 On the **Instance** tab page, select the instance to be restarted. Click **More** and select **Rolling-restart Instance**.

Figure 3-60 Performing a rolling instance restart

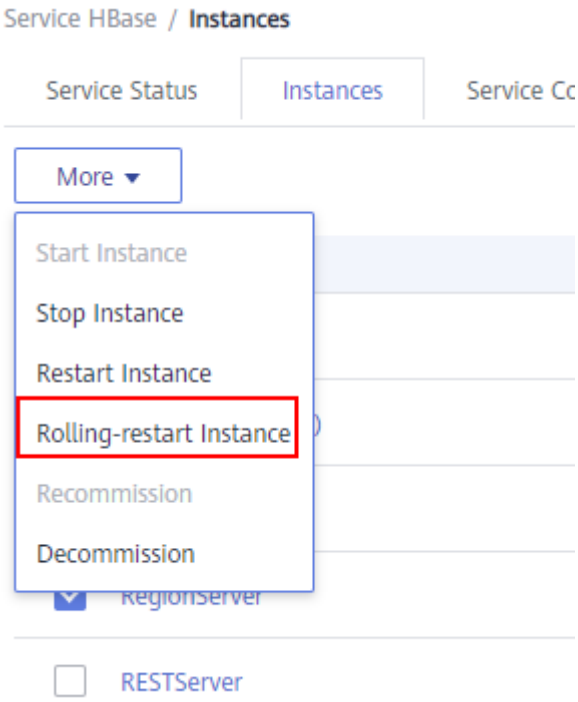

Step 4 After you enter the administrator password, the **Rolling-restart Instance** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the instance.

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

----End

Perform a Rolling Cluster Restart

Step 1 Choose **Clusters** > **Active Clusters** and click a cluster name to go to the cluster details page.

Step 2 In the upper right corner of the page, choose **Management Operations** > **Perform Rolling Cluster Restart**.

Figure 3-61 Performing a Rolling Restart of a Cluster (Using MRS 1.9.2 as an Example)

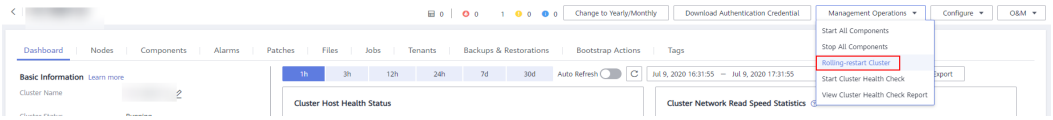

Step 3 The **Rolling-restart Cluster** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the cluster.

Step 4 After the rolling restart task is complete, click **Finish**.

----End

Rolling Restart Parameter Description

Table 3-54 describes rolling restart parameters.

Table 3-54 Rolling restart parameter description

Parameter	Description
Only restart instances whose configurations have expired	Specifies whether to restart only the modified instances in a cluster.
Enable rack strategy	Whether to enable the concurrent rack rolling restart strategy. This parameter takes effect only for roles that meet the rack rolling restart strategy. (The roles support rack awareness, and instances of the roles belong to two or more racks.) **NOTE** This parameter is configurable only when a rolling restart is performed on HDFS and YARN in MRS 3.*x* or later.
Data Node Instances to Be Batch Restarted	Specifies the number of instances that are restarted in each batch when the batch rolling restart strategy is used. The default value is **1**. The value ranges from 1 to 20. This parameter is valid only for data nodes.
Batch Interval	Specifies the interval between two batches of instances for rolling restart. The default value is **0**. The value ranges from 0 to 2147483647. The unit is second. Note: Setting the batch interval parameter can increase the stability of the big data component process during the rolling restart. You are advised to set this parameter to a non-default value, for example, 10.
Decommissioning Timeout Interval	Decommissioning interval for role instances during a rolling restart.

Parameter	Description
Batch Fault Tolerance Threshold	Specifies the tolerance times when the rolling restart of instances fails to be executed in batches. The default value is **0**, which indicates that the rolling restart task ends after any batch of instances fails to be restarted. The value ranges from 0 to 2147483647.

Procedure in a Typical Scenario

Step 1 Choose **Clusters** > **Active Clusters** and click a cluster name to go to the cluster details page.

Step 2 Click **Components** and select **HBase**. The **HBase** service page is displayed.

Step 3 Click the **Service Configuration** tab, modify an HBase parameter, and save the configuration as prompted.

NOTE

In versions earlier than MRS 3.*x*, do not select **Restart the affected services or instances**. This option indicates a normal restart. If you select this option, all services or instances will be restarted, which may cause service interruption.

Step 4 After saving the configurations, click **Finish**.

Step 5 Click the **Service Status** tab.

Step 6 On the **Service Status** tab page, click **More** and select **Rolling-restart Service**.

Figure 3-62 Service status - rolling restart (using MRS 1.9.2 as an example)

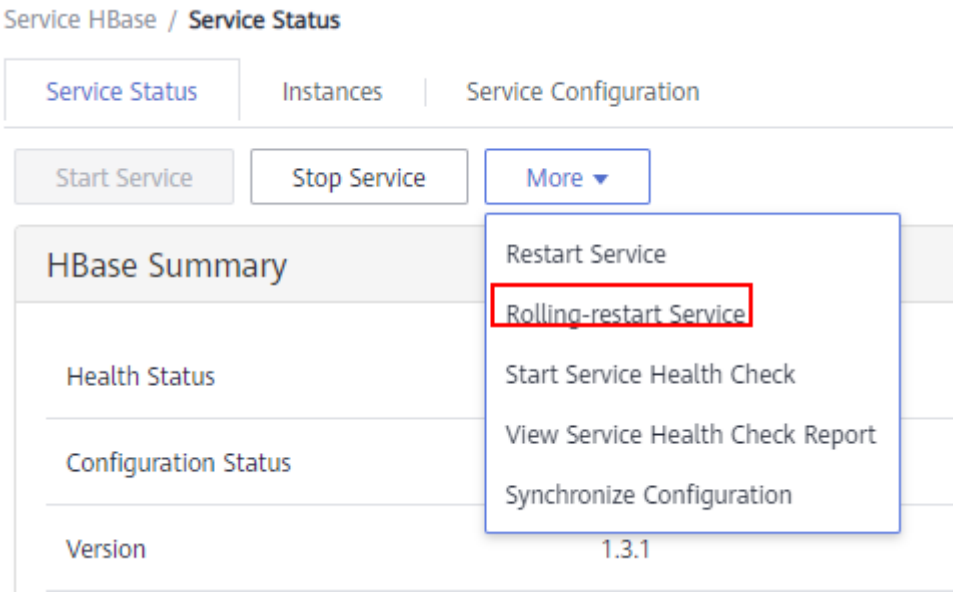

Step 7 After you enter the administrator password, the **Rolling-restart Service** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart.

Figure 3-63 Configuring the rolling service restart

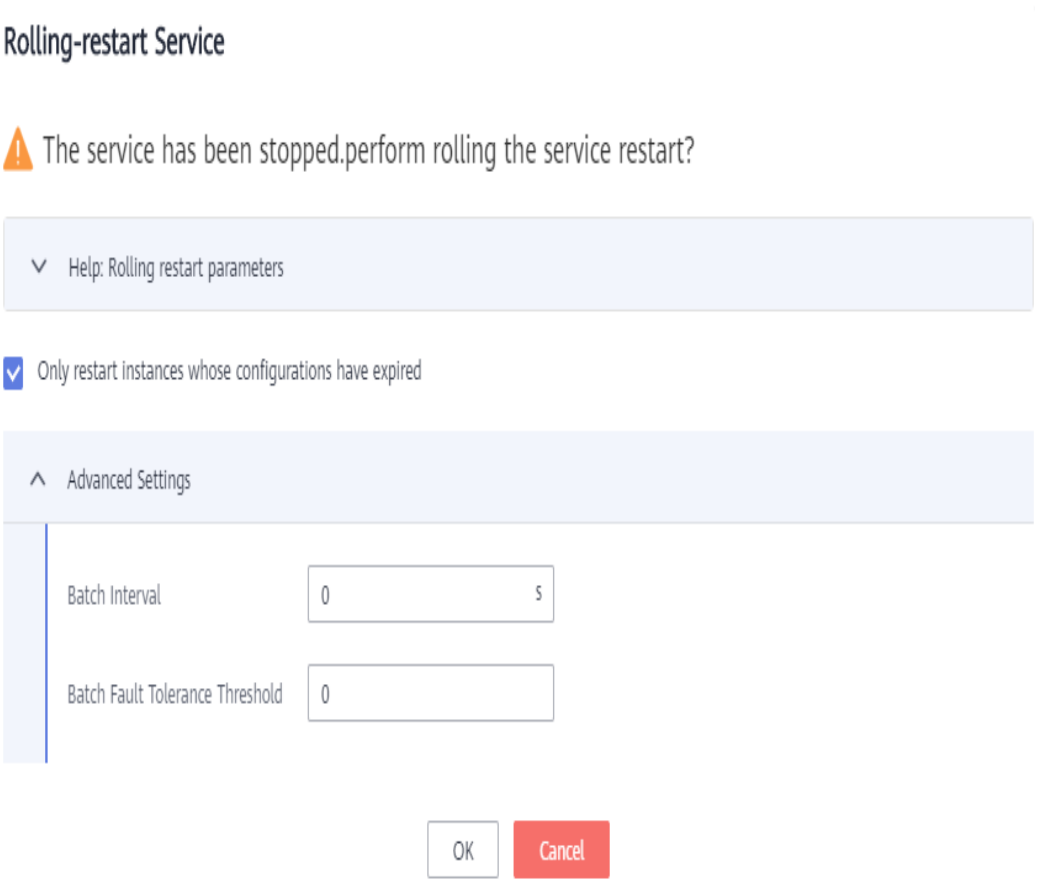

Step 8 After the rolling restart task is complete, click **Finish**.

Figure 3-64 Finishing the rolling service restart

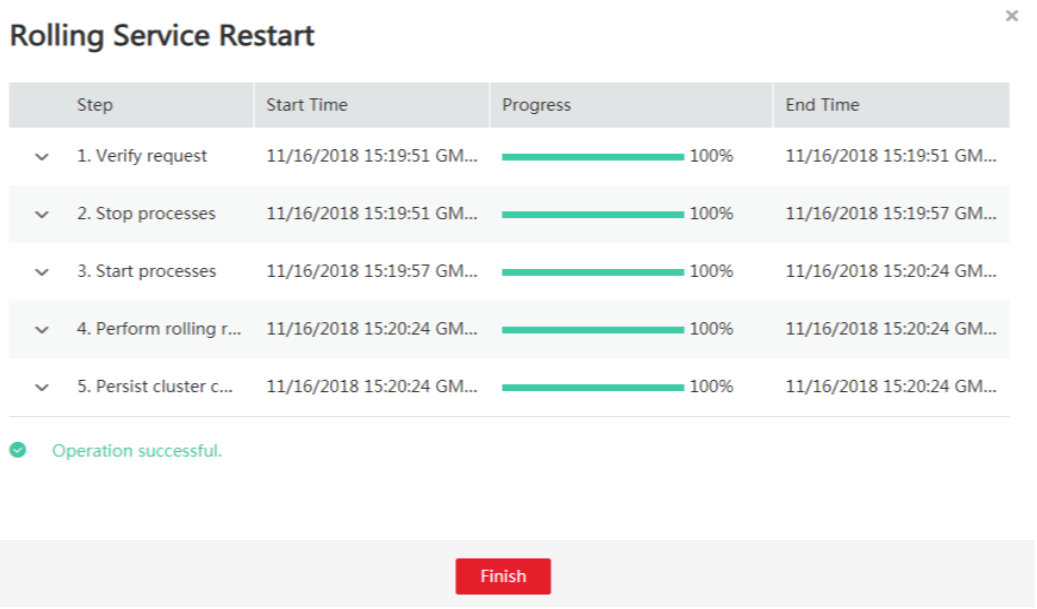

----End

3.8 Alarm Management

3.8.1 Viewing the Alarm List

The alarm list displays all alarms in the MRS cluster. The MRS page displays the alarms that need to be handled in a timely manner and the events.

On the MRS management console, you can only query basic information about uncleared MRS alarms on the **Alarms** tab page. For details about how to view alarm details or manage alarms, see **Viewing and Manually Clearing an Alarm**.

Alarms are listed in chronological order by default in the alarm list, with the most recent alarms displayed at the top.

Table 3-55 describes various fields in an alarm.

Table 3-55 Alarm description

Parameter	Description
Alarm ID	ID of an alarm.
Alarm Name	Name of an alarm.

Parameter	Description
Severity	Alarm severity. In versions earlier than MRS 3.x, the cluster alarm severity is as follows: • Critical Indicates alarms reporting errors that affect cluster running, such as unavailable cluster services, node faults, data inconsistency between the active and standby GaussDB databases, and abnormal LdapServer data synchronization. You need to check the cluster status based on the alarms and rectify the faults in a timely manner. • Major Indicates alarms reporting errors that affect some cluster functions, including process faults, periodic backup task failures, and abnormal key file permissions. Check the objects for which the alarms are generated based on the alarms and clear the alarms in a timely manner. • Minor Indicates alarms reporting errors that do not affect major functions of the current cluster, including alarms indicating that the certificate file is about to expire, audit logs fail to be dumped, and the license file is about to expire. • Warning Indicates an alarm of the lowest severity. It is used for information display or prompt and indicates that an event occurs in the scenarios when you stop a service, delete a service, stop an instance, delete an instance, delete a node, restart a service, restart an instance, perform an active/standby switchover for MRS Manager, scale in a host, or restore an instance. Additionally, this type of alarms also occurs when an instance is faulty, a job executed successfully, or a job failed to be executed. In MRS 3.x or later, the alarm severity of a cluster is as follows: • Critical Indicates alarms reporting errors that affect cluster running, such as unavailable cluster services, node faults, data inconsistency between the active and standby GaussDB databases, and abnormal LdapServer data synchronization. You need to check the cluster status based on the alarms and rectify the faults in a timely manner. • Major Indicates alarms reporting errors that affect some cluster functions, including process faults, periodic backup task failures, and abnormal key file permissions. Check the objects for which the alarms are generated based on the alarms and clear the alarms in a timely manner. • Minor Indicates alarms reporting errors that do not affect major functions of the current cluster, including alarms indicating

Parameter	Description
	that the certificate file is about to expire, audit logs fail to be dumped, and the license file is about to expire. • Suggestion Indicates an alarm of the lowest severity. It is used for information display or prompt and indicates that an event occurs in the scenarios when you stop a service, delete a service, stop an instance, delete an instance, delete a node, restart a service, restart an instance, perform an active/standby switchover for MRS Manager, scale in a host, or restore an instance. Additionally, this type of alarms also occurs when an instance is faulty, a job executed successfully, or a job failed to be executed.
Generated	Time when the alarm is generated.
Location	Details about the alarm.
Operation	If the alarm can be manually cleared, click **Clear Alarm**. To view details about an alarm, click **View Help**. (This function is available in MRS 3.*x* or later).

Table 3-56 Button description

Button	Description
	Select an interval for refreshing the alarm list from the drop-down list. • Refresh every 30s • Refresh every 60s • Stop refreshing
	Select an alarm severity from the drop-down list box to filter alarms. For versions earlier than MRS 3.x, the following alarms can be filtered: All, Critical, Major, Minor, and Warning. (For MRS 3.x or later) You can filter the following alarms: All, Critical, Major, Minor, and Warning.
	Click  and manually refresh the alarm list.
Advanced Search	Click **Advanced Search**. In the displayed alarm search area, set search criteria and click **Search** to view the information about specified alarms. You can click **Reset** to clear the search criteria.

3.8.2 Viewing the Event List

The event list displays information about all events in a cluster, such as service restart and service termination.

Events are listed in the event list in chronological order by default, with the most recent events displayed at the top.

NOTE

You can view the event list of an MRS 2.x or later cluster when IAM user synchronization is complete.

Table 3-57 describes various fields for an event.

Table 3-57 Event description

Parameter	Description
Event ID	Specifies the ID of an event.
Event Severity	Specifies the event severity. In versions earlier than MRS 3.x, the cluster event level is as follows: • Critical • Major • Minor • Suggestion In MRS 3.x or later, the event level of a cluster is as follows: • Critical • Major • Minor • Suggestion
Event Name	Name of the generated event.
Generated	Time when the event is generated.
Location	Specifies the detailed information for locating the event,

Table 3-58 Icon description

Icon	Description
	Select an interval for refreshing the event list from the drop-down list. • Refresh every 30s • Refresh every 60s • Stop refreshing

Icon	Description
	Click  to manually refresh the event list.
Advanced Search	Click **Advanced Search**. In the displayed event search area, set search criteria and click **Search** to view the information about specified events. Click **Reset** to clear the search criteria.

Exporting events

Step 1 Choose **Clusters > Active Clusters** and click a cluster name to go to the cluster details page.

Step 2 Click **Alarm Management** > **Events**.

Step 3 Click **Export All**.

Step 4 In the displayed dialog box, select the type and click **OK**.

----End

Common Events

Table 3-59 Common events

Event ID	Event Name
12019	Stop Service
12020	Delete Service
12021	Stop RoleInstance
12022	Delete RoleInstance
12023	Delete Node
12024	Restart Service
12025	Restart RoleInstance
12026	Manager Switchover
12065	Restart Process
12070	Job Running Succeeded
12071	Job Running Failed
12072	Job Killed
12086	Restart Agent
12152	Start Periodic Replication

Event ID	Event Name
12153	Periodic Replication Completed
12154	Start Streaming Replication
12155	Restart Streaming Replication
12156	Stop Streaming Replication
12157	Skip Periodic Synchronization
14005	NameNode Switchover
14028	HDFS DiskBalancer Task
14029	Active NameNode Entered Security Mode and Generated New Fsimage
17001	Oozie Workflow Execution Failure
17002	Oozie Scheduled Job Execution Failure
18001	ResourceManager Switchover
18004	JobHistoryServer Switchover
19001	HMaster Failover
20003	Hue Failover
24002	Flume Channel Overflow
25001	LdapServer Failover
27000	DBServer Switchover
29001	Impala HaProxy Active/Standby Switchover
29002	Impala StateStoreCatalog Active/Standby Switchover
38003	Adjust Topic Data Storage Period
43014	Spark2x Data Skew
43015	Spark2x SQL Large Query Results
43016	Spark2x SQL Execution Timeout
43024	Start JDBCServer
43025	Stop JDBCServer
43026	ZooKeeper Connection Succeeded
43027	Zookeeper Connection Failed
44003	Coordinator Switchover

3.8.3 Viewing and Manually Clearing an Alarm

Scenario

You can view and clear alarms on MRS.

Generally, the system automatically clears an alarm when the fault is rectified. If the fault has been rectified and the alarm cannot be automatically cleared, you can manually clear the alarm.

You can view the latest 100,000 alarms (including uncleared, manually cleared, and automatically cleared alarms) on MRS. If the number of cleared alarms exceeds 100,000 and is about to reach 110,000, the system automatically dumps the earliest 10,000 cleared alarms to the dump path.

3. In versions earlier than x, the value is the same as that of ${BIGDATA_HOME}/OMSV100R001C00x8664/workspace/data for the active management node.

(For 3.x and later versions) The path is **${BIGDATA_HOME}/om-server/OMS/workspace/data** of the active management node.

A directory is automatically generated when alarms are dumped for the first time.

NOTE

Set an automatic refresh interval or click for an immediate refresh.

The following refresh interval options are supported:

- Refresh every 30 seconds
- Refresh every 60 seconds
- Stop refreshing

Procedure

Step 1 Choose **Clusters > Active Clusters** and click a cluster name to go to the cluster details page.

Step 2 Click **Alarms** and view the alarm information in the alarm list.

- By default, the alarm list page displays the latest 10 alarms.
- By default, data is sorted in descending order based on the generation time. For MRS 3.x or earlier, you can click the alarm ID, severity, and generation time to modify the sorting mode. For clusters of MRS 3.x or later, you can click the severity and generation time to modify the sorting mode.
- You can filter all alarms of the same severity. The results include cleared and uncleared alarms.
- For clusters of MRS 3.x and earlier versions, you can click , , or in the upper right corner of the page to quickly filter **Critical**, **Major**, **Minor**, or **Suggestion** alarms that are uncleared.
- For clusters of MRS 3.x or later: You can click , , or in the upper right corner of the page to quickly filter uncleared **Critical**, **Major**, **Minor** or **Warning** alarms.

Step 3 Click **Advanced Search**. In the displayed alarm search area, set search criteria and click **Search** to view the information about specified alarms. You can click **Reset** to clear the search criteria.

NOTE

The start time and end time are specified in **Time Range**. You can search for alarms generated within the time range.

Handle the alarm by referring to **Alarm Reference**. If the alarms in some scenarios are generated due to other cloud services that MRS depends on, you need to contact maintenance personnel of the corresponding cloud services.

Step 4 Click **Clear Alarm** if you need to. In the displayed dialog box, click **OK**.

NOTE

If multiple alarms have been handled, you can select one or more alarms to be cleared and click **Clear Alarm** to clear the alarms in batches. A maximum of 300 alarms can be cleared in each batch.

----End

Exporting Alarms

Step 1 Choose **Clusters > Active Clusters** and click a cluster name to go to the cluster details page.

Step 2 Click **Alarm Management** > **Alarms**.

Step 3 Click **Export All**.

Step 4 In the displayed dialog box, select the type and click **OK**.

----End

3.9 Patch Management

3.9.1 Installing an Online Patch

Install the cluster version patch as required if you obtain patch information from:

- The message center service
- On a cluster details page of the MRS console, choose **Patches** > **Cluster Component Patches** to view the patches that can be installed.

Preparing for Patch Installation

CAUTION

- For details about how to check the cluster status, see section **Performing a Health Check**. Exceptions such as cluster node faults and hard disk faults may cause patch installation and uninstallation failures. Before you install or uninstall the patch, ensure that the cluster is healthy.
- Click **Patches** then **Cluster Component Patches**, view the **Patch Description** column of the target patch, and read the patch description carefully to understand the patch installation procedure and impact.
- MRS 2.x and earlier versions, MRS 3.1.5 and later versions, and MRS 3.2.0-LTS and later versions support online patch installation.

Installing a Patch

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of the desired cluster.

Step 3 In the **Patches** tab, click **Cluster Component Patches**. In the operation list, click **Install** next to the patch you want to install.

Step 4 In the displayed dialog box, select **I have read Patch Description and understood that this operation may restart services.** Click **Yes** and wait until the patch is successfully installed.

Step 5 Check the patch status. Restart the components and install the client patch according to the patch description.

NOTE

If there is an isolated host in the cluster, the patch will not be installed on the isolated host. In this case, when the installation completes, the patch is partially installed. After an isolated node is restored and the isolation is canceled, you can install the patch again. In this case, the patch is installed only on the node where isolation is canceled. For versions earlier than MRS 3.x, perform operations by referring to **Restoring Patches for the Isolated Hosts**.

----End

Uninstalling a Patch

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of the desired cluster.

Step 3 In the **Patches** tab, click **Cluster Component Patches**. In the operation list, click **Uninstall** next to the patch you want to uninstall.

Step 4 In the displayed dialog box, select the confirmation check box and click **Yes**. Wait until the patch is successfully uninstalled.

Step 5 Restart the component and uninstall the client patch according to the patch description.

NOTE

If there is an isolated host in the cluster, the patch will not be uninstalled on the isolated host. In this case, when the uninstallation completes, the patch is partially uninstalled. After the isolated node is restored and the isolation is canceled, you can uninstall the patch again. In this case, the patch is uninstalled only on the node whose isolation is canceled. For versions earlier than MRS 3.x, perform operations by referring to **Restoring Patches for the Isolated Hosts**.

----End

3.9.2 Installing a Rolling Patch

The rolling patch function indicates that patches are installed or uninstalled for one or more services in a cluster by performing a rolling service restart (restarting services or instances in batches), without interrupting the services or within a minimized service interruption interval. Services in a cluster are divided into the following three types based on whether they support rolling patch:

- Services supporting rolling patch installation or uninstallation: All businesses or part of them (varying depending on different services) of the services are not interrupted during patch installation or uninstallation.
- Services not supporting rolling patch installation or uninstallation: Businesses of the services are interrupted during patch installation or uninstallation.
- Services with some roles supporting rolling patch installation or uninstallation: Some businesses of the services are not interrupted during patch installation or uninstallation.

NOTE

In **MRS 3.x**, you cannot perform operations in this section on the management console.

Table 3-60 provides services and instances that support or do not support rolling restart in the MRS cluster.

Table 3-60 Services and instances that support or do not support rolling restart

Service	Instance	Support Rolling Restart
Alluxio	AlluxioJobMaster	Yes
	AlluxioMaster	
ClickHouse	ClickHouseServer	Yes
	ClickHouseBalancer	
CDL	CDLConnector	Yes
	CDLService	
Flink	FlinkResource	No
	FlinkServer	
Flume	Flume	Yes
	MonitorServer	

Service	Instance	Support Rolling Restart
Guardian	TokenServer	Yes
HBase	HMaster	Yes
	RegionServer	
	ThriftServer	
	RESTServer	
HetuEngine	HSBroker	Yes
	HSConsole	
	HSFabric	
	QAS	
HDFS	NameNode	Yes
	Zkfc	
	JournalNode	
	HttpFS	
	DataNode	
Hive	MetaStore	Yes
	WebHCat	
	HiveServer	
Hue	Hue	No
Impala	Impalad	No
	StateStore	
	Catalog	
IoTDB	IoTDBServer	Yes
Kafka	Broker	Yes
	KafkaUI	No
Kudu	KuduTserver	Yes
	KuduMaster	
Loader	Sqoop	No
MapReduce	JobHistoryServer	Yes
Oozie	oozie	No
Presto	Coordinator	Yes

Service	Instance	Support Rolling Restart
	Worker	
Ranger	RangerAdmin	Yes
	UserSync	
	TagSync	
Spark	JobHistory	Yes
	JDBCServer	
	SparkResource	
Storm	Nimbus	Yes
	UI	
	Supervisor	
	Logviewer	
Tez	TezUI	No
Yarn	ResourceManager	Yes
	NodeManager	
Zookeeper	Quorumpeer	Yes

Installing a Rolling Patch

Step 1 Log in to the MRS console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's basic information.

Step 3 On the **Patches** page, click **Install** in the **Operation** column.

Step 4 On the **Warning** page, enable or disable **Rolling Patch**.

Figure 3-65 Rolling patch installation

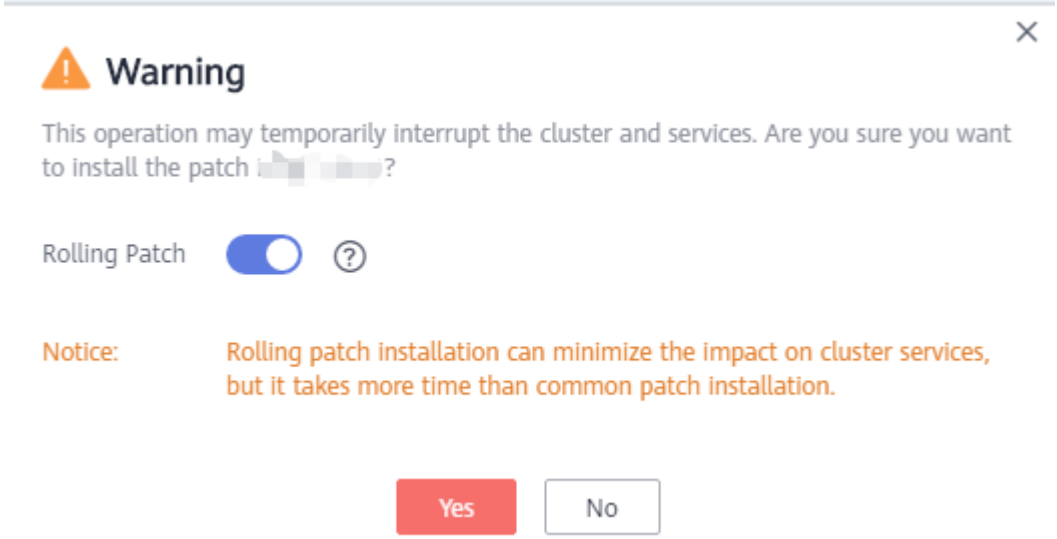

NOTE

- Enabling the rolling patch installation function: Services are not stopped before patch installation, and rolling service restart is performed after the patch installation. This minimizes the impact on cluster services but takes more time than common patch installation.
- Disabling the rolling patch uninstallation function: All services are stopped before patch uninstallation, and all services are restarted after the patch uninstallation. This temporarily interrupts the cluster and the services but takes less time than rolling patch uninstallation.
- The rolling patch installation function is not available in clusters with less than two Master nodes and three Core nodes.

Step 5 Click **Yes** to install the target patch.

Step 6 View the patch installation progress.

1. Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.
2. Choose **System** > **Manage Patch**. On the **Manage Patch** page, you can view the patch installation progress.

NOTE

For the isolated host nodes in the cluster, follow instructions in **Restoring Patches for the Isolated Hosts** to restore the patch.

----End

Uninstalling a Rolling Patch

Step 1 Log in to the MRS console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's basic information.

Step 3 On the **Patches** page, click **Uninstall** in the **Operation** column.

Step 4 On the **Warning** page, enable or disable **Rolling Patch**.

Figure 3-66 Rolling patch uninstallation

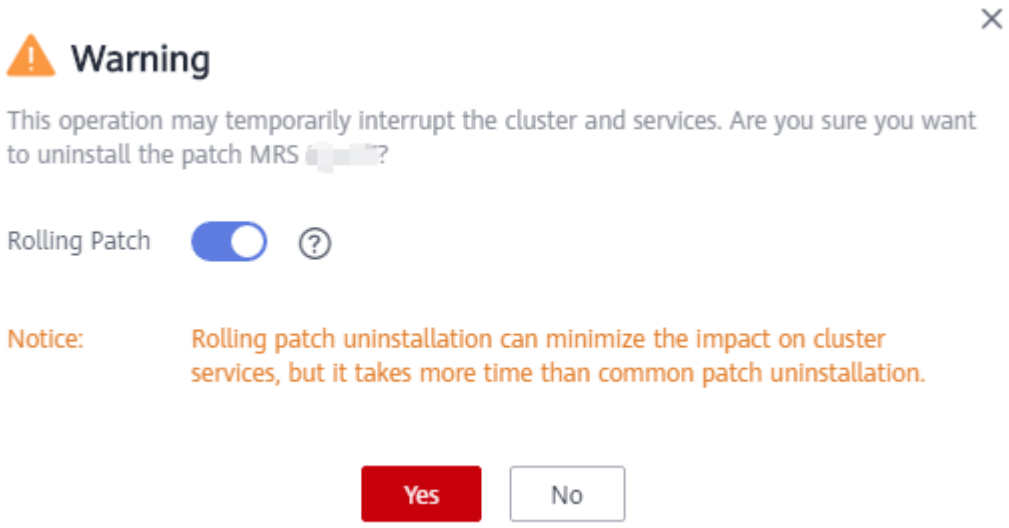

NOTE

- Enabling the rolling patch uninstallation function: Services are not stopped before patch uninstallation, and rolling service restart is performed after the patch uninstallation. This minimizes the impact on cluster services but takes more time than common patch uninstallation.
- Disabling the rolling patch uninstallation function: All services are stopped before patch uninstallation, and all services are restarted after the patch uninstallation. This temporarily interrupts the cluster and the services but takes less time than rolling patch uninstallation.
- Only patches that are installed in rolling mode can be uninstalled in the same mode.

Step 5 Click **Yes** to uninstall the target patch.

Step 6 View the patch uninstallation progress.

1. Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.
2. Choose **System** > **Manage Patch**. On the **Manage Patch** page, you can view the patch uninstallation progress.

NOTE

For the isolated host nodes in the cluster, follow instructions in **Restoring Patches for the Isolated Hosts** to restore the patch.

----End

3.9.3 Restoring Patches for the Isolated Hosts

If some hosts are isolated in a cluster, perform the following operations to restore patches for these isolated hosts after patch installation on other hosts in the cluster. After patch restoration, versions of the isolated host nodes are consistent with those are not isolated.

NOTE

Operations in this section cannot be performed on the management console of MRS 3.x. This section applies only to versions earlier than 3.x.

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 Choose **System** > **Manage Patch**. The **Manage Patch** page is displayed.

Step 3 In the **Operation** column, click **View Details**.

Step 4 On the patch details page, select host nodes whose **Status** is **Isolated**.

Step 5 Click **Select and Restore** to restore the isolated host nodes.

Figure 3-67 Restoring patches for the isolated hosts

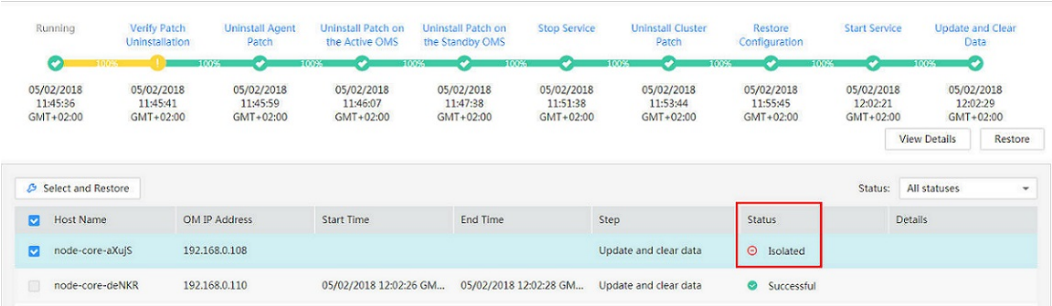

----End

3.9.4 MRS Patch Description

3.9.4.1 Fixed the Privilege Escalation Vulnerability of User omm

Applicable Version

All MRS versions

Resolved Issue

Rectify the issue that user **omm** can use the **installSudoExecute.sh** script to obtain the permission of user **root**.

Structure of the Patch Package

- **install.sh**: patch installation script.
- **ips.ini**: stores the IP addresses of all nodes in the cluster. Modify this file based on the actual IP addresses of the nodes in the cluster. Each IP address occupies a line. No blank line is allowed between IP addresses. Leave a blank line at the end of the file.
- **scp-util.exp**: SCP tool script.
- **ssh-util.exp**: SSH tool script.
- **Sudo_Vulnerability_20210330**: directory for storing the **sudo_repair.sh** script. You can copy the directory to the folder for running the script on each node.

- **sudo_repair.sh**: script for fixing the vulnerability.
- **README.md**: describes how to use the patch tool.

Installing the Patch

Step 1 Click the address in the corresponding area of the cluster to download the patch package.

- **CN-Hong Kong**: **https://mrs-container1-patch-ap-southeast-1.obs.ap-southeast-1.myhuaweicloud.com/MRS_Common_Script/MRS_All_Sudo_Vulnerability_20210330.tar.gz**
- **AP-Bangkok**: **https://mrs-container1-patch-ap-southeast-2.obs.ap-southeast-2.myhuaweicloud.com/MRS_Common_Script/MRS_All_Sudo_Vulnerability_20210330.tar.gz**
- **AP-Singapore**: **https://mrs-container1-patch-ap-southeast-3.obs.ap-southeast-3.myhuaweicloud.com/MRS_Common_Script/MRS_All_Sudo_Vulnerability_20210330.tar.gz**

Step 2 Log in to the active master node of the cluster as the **root** user.

Step 3 Upload the patch package to **/root/**.

Step 4 Run the following command to decompress the patch tool package **MRS_All_Sudo_Vulnerability_20210330.tar.gz** to the current directory (**/root**).

tar -zxf MRS_All_Sudo_Vulnerability_20210330.tar.gz

Step 5 Run the following command to go to the directory where the **ips.ini** file is located.

cd /root/MRS_All_Sudo_Vulnerability_20210330/

Step 6 Configure the IP addresses of all nodes in the cluster in the **ips.ini** file. Each IP address occupies a line. No blank line is allowed between IP addresses. Leave a blank line at the end of the file.

Step 7 Run the following script to install the patch.

After the script is run, you need to enter the correct password of user **root**. If the password is incorrect, the account may be locked for 5 minutes.

cd /root/MRS_All_Sudo_Vulnerability_20210330/

dos2unix ./*

chmod +x ./* -R

sh install.sh "install"

----End

Uninstalling a Patch

Run the following script to uninstall the patch. After the script is run, you need to enter the correct password of user **root**. If the password is incorrect, the account may be locked for 5 minutes during the SSH process of the script.

cd /root/MRS_All_Sudo_Vulnerability_20210330/

sh install.sh "uninstall"

3.9.4.2 MRS 3.2.0-LTS.1 Patch Description

Basic information about MRS 3.2.0-LTS.1.6

Table 3-61 Basic information

Patch Version	MRS 3.2.0-LTS.1.6
Release Date	2024-02-04
Pre-Installation Operations	If an MRS cluster node is faulty or the network is disconnected, isolate the node first. Otherwise, the patch installation will fail.
New Features	• CDL supports uppercase letters of table fields. • Sharding key is specified when flink writes data to the NetEase DDB. • Flink supports writing upsertkafka monitoring data into influxdb. • Flink stream read Hudi supports monitoring of the message retention time and message stacking time. • Flink supports the ignoreDelete feature. • Yarn NodeManager supports graceful decommissioning. • Kafka supports data encryption. • Spark supports subquery fields without aggregate functions. (Set **spark.sql.legacy.correlated.scalar.query.enabled** to **true**.) • Spark supports view and chart permission control. (Add **spark.ranger.plugin.viewaccesscontrol.enable** to the custom parameter of JDBCserver and set it to **true**.) Add **spark.ranger.plugin.viewaccesscontrol.enable=true** to the **Spark2x/spark/conf/spark-defaults.conf** configuration file in the client directory and restart the JDBCserver instance.

Resolved Issues	List of resolved issues in MRS 3.2.0-LTS.1.6: ● No data is written into the Hudi table when the start time parameter is added to a CDL job. ● Residual connector threads cannot be killed when CDL is restarted. ● The error logs of the Clickhouse balance are not rolled back. ● It takes a long time to import ClickHouse data to the database. ● The ClickHouse TTL does not take effect. ● The Hive client cannot be connected due to too many ClickHouse and ZooKeeper connections. ● OOM is reported when columns are added to the base table of the ClickHouse materialized view. ● Executor optimizes the job aging sequence to aging based on the job end time. ● The UDF fails to be updated because the package name of the UDF uploaded on Flink UI is not changed. ● The UI paging display is abnormal after UDFs are uploaded to Flink. ● The time displayed when Flink streams read Hudi data is 8 hours shorter than the actual time. ● An Error Is Reported When FlinkServer Is Used to Create a Kafka Connection with SASL_SSL Authentication ● When Flink interconnects with Guardian and OBS is enabled, Flink jobs occasionally fail to be restarted. ● Resolved the problem that jobs submitted by DGC cannot be restored from checkpoints. ● When FlinkServer is used to submit a job, a message is displayed indicating that the job fails to be submitted, but the job on Yarn is in RUNNING state. ● The StackOverFlow error is reported when a Flink-SQL job is submitted using DGC. ● The value of Flink security.ssl.encrypt.enabled is changed to TRUE, which affects the task startup speed. ● When a DGC job fails and is retried, the job is consumed from the beginning instead of being started from the checkpoint. ● An error is reported when a common cluster uses REST APIs to invoke FlinkServer jobs or actions ● The ClickHouse job cannot be submitted when the Flink JAR package is written. ● The function of deleting hint parameters in the Flink Join state does not take effect. ● The Flink job error information is incorrect, and the number of error lines displayed in the log is inconsistent with the number of SQL error lines of the job.

	• The Flink stream fails to read the Hudi job when data is written to the table, and the Hudi job fails to be submitted and executed. • The Flinkserver SQL verification fails. • Keytab and principal cannot be configured by **-yD** or **-D** on the Flink client. • When the Flink operator chain is opened, the sink operator exception is captured by the source operator, but the TM job does not fail. • The specific job name is not displayed in the MRS real-time task alarm notification. • The Flume customized time interceptor does not take effect. • Single-Node capacity expansion fails occasionally due to NodeManager startup failure during autoscaling. • In security mode, the HDFS user cannot be used to deliver balance tasks. • The available resource metrics of the Yarn resource pool are abnormal. As a result, auto scaling is triggered abnormally. • Rectify the fault that the disk of the NodeManager node is full. As a result, members in the resource pool are migrated to the default resource pool. • The alarm indicating that the number of dead DataNodes exceeds the threshold cannot be automatically cleared. • A false alarm indicating that the service is unavailable is reported when invoking the HBase health check script times out. • After a column is added using HetuEngine alter, no data is available in the old partition when data is inserted. • When HetuEngine or Flink writes data, there is a possibility that HetuEngine reads no data from the RO table after compaction. • The HDFS log collection package is missing. • HetuEngine fails to execute SQL statements when using HSFabric to connect to JDBC. • An error is reported when the Hive on tez overwrite partition table result is empty. • Hive integrates the DataAtrs metadata synchronization plug-in package. • Hive Jobs frequently report errors after the last access time of Hive metadata is configured. • The batch deletion parameter does not take effect in "Hive alter table test drop partition (partition <'xxxx');". • Parameters such as **&useSSL=false** do not take effect when Hive JDBC is connected. • In high-concurrency scenarios, Hive on Spark jobs fail due to JAR package submission timeout.

	• Duplicate bucket IDs exist when CDL writes data to the Hudi table. • Hudi compaction runs faster than clean. As a result, data cannot be read. • The "xx is not a Parquet file" exception is reported when Spark reads Hudi because abnormal files are not cleared after a compaction task fails. • When a Spark job reads an upstream database table, Excutor reports an error indicating that files in the OBS **.schema** directory of the table cannot be found. • The Hudi compaction schedule is optimized. A plan is generated based on the last compaction action. • By default, Hudi retains 5 GB archived compressed files. • By default, Hudi OBS is not moved to the recycle bin. • Hudi archive archives more clean and rollback operations to reduce the number of metadata files. • Archiving cannot be triggered if clean is not executed for Hudi Cow. • The table directory is deleted when the Hudi uses Spark to generate data in the batch supplement table. • The namespaces generated for the Hudi table are inconsistent. As a result, the update fails. • The call command is added to Hudi to clear invalid metadata files. • Instant data is deleted when the Hudi DELETE_EMPTY_INSTANT function is abnormal. • Data Is Lost in Flink append mode • Historical data is not cleared after being stored for 30 days by default. • The real-time monitoring data on FusionInsight Manager is empty after the time zone is changed to UTC. • The global user policy content under Tenant Resources on FusionInsight Manager is not displayed on multiple pages. • The **yarn on cce** button is added to the Manager page. • The Tomcat configuration needs to be manually modified when a new cluster accesses the Manager page through a private line. • Interconnecting Manager with the O&M Page • A false alarm is generated in an MRS cluster. • The monitoring metrics are displayed abnormally when the Manager queries data across 00:00. • The CPU usage of the kernel space on FusionInsight Manager is incorrect. • After a user switches to the distribution chart on the host cable management monitoring page of FusionInsight Manager, the

	default value range of the icon is incorrect when the user switches to the distribution chart. • The GaussDB process of the active OMS occupies a large amount of memory. • After the customized configuration of the Manager component is complete, the customized configuration of the added instance is not displayed. • The PMS process keeps restarting. • The alarm indicating that the trust relationship between nodes is invalid is falsely reported during node scale-out. • Data synchronization between the active and standby Manager nodes is abnormal. • The alarm indicating that the number of dead DataNodes exceeds the threshold cannot be automatically cleared. • Alarms Reported in the Manager Auto Scaling Scenario Are Optimized • Optimizing the Node Network Communication Exception Alarm in the Manager Scale-in Scenario • An Alarm Indicating that the Service Is Unavailable Is Reported When the Controller Is Restarted • The administrator account of MRS Manager forgets to add the redirection support document. • RangeAdmin Instance Fails to Be Started After External Ranger Metadata Is Configured for a Cluster • An error is reported when the quit command is executed to exit the Spark client. • The driver port conflicts with the thriftserver port occasionally in Spark multi-tenant mode. • Resources Are Not Released for Idle Spark JDBC Tasks That Have Been Running for More Than 30 Minutes • After the SSL of ZooKeeper is enabled and a Spark job is submitted, ZooKeeper fails to be connected. As a result, the task execution times out. • Spark Fails to Connect to JDBCServer • When Spark grants only the view permission but not the table permission, Hive can query views but SparkSQL cannot query views.
Compatibility with Other Patches	The MRS 3.2.0-LTS.1.6 patch package contains all patches for fixing single-point issues in MRS 3.2.0-LTS.1.

Impact of Patch Installation	See **Impact of Patch Installation**.

Impact of Patch Installation

- During MRS 3.2.0-LTS.1 patch installation, OMS automatically restarts, which affects cluster management operations, such as job submission and cluster scaling. Install patch in off-peak hours.
- After the MRS 3.2.0-LTS.1.6 patch is installed or uninstalled, restart the Flink, Yarn, HDFS, MapReduce, Ranger, HetuEngine, Flume, Hive, Kafka, and Spark2x services on FusionInsight Manager for the patch to take effect. During restart, some services may be unavailable for a short period. To ensure service continuity, restart the components in off-peak hours. Before uninstalling the patch, log in to FusionInsight Manager, and choose **System** > **Third-Party AD** to disable AD interconnection.
- If a client is manually installed inside or outside the cluster, you need to upgrade or roll back the client.

  a. Log in to the active node of the cluster.

     cd /opt/Bigdata/patches/{*Patch version*}**/download/**

     NOTE

     In all operations, replace {*Patch version*} with that used in the real-life project. For example, if the installed patch is MRS_3.2.0-LTS.1.1, the value of {*Patch version*} is **MRS_3.2.0-LTS.1.1**.

  b. Copy the patch installation package to the **/opt/** directory on the client node.

     scp patch.tar.gz *{IP address of the client node}***:/opt/**

     The following shows an example.

     scp patch.tar.gz 127.0.0.1:/opt/

  c. Log in to the node where the client is deployed.

     The following shows an example.

     ssh 127.0.0.1

  d. Run the following commands to create a patch directory and decompress the patch package:

     mkdir /opt/{*Patch version*}

     tar -zxf /opt/patch.tar.gz -C /opt/{*Patch version*}

  e. Upgrade/Rollback a patch.

     - Upgrade the patch on the client node.

       Log in to the node where the client is deployed.

       cd /opt/{*Patch version*}**/client**

       sh upgrade_client.sh upgrade *{Client installation directory}*

       The following shows an example.

sh upgrade_client.sh upgrade /opt/client/

- Roll back the patch on the client node (after the patch is uninstalled).

  Log in to the node where the client is deployed.

  cd /opt/{*Patch version*}**/client**

  sh upgrade_client.sh rollback *{Client installation directory}*

  The following shows an example.

  sh upgrade_client.sh rollback /opt/client/

- If the Spark service is installed on MRS 3.2.0-LTS.1, upgrade the ZIP package in HDFS on the active OMS node after the patch is installed.

  a. Log in to the active node of the cluster.

  su - omm

  cd /opt/Bigdata/patches/{*Patch version*}**/client/**

  source /opt/Bigdata/client/bigdata_env

  b. Authenticate users who have permissions on HDFS on a cluster in security mode.

  kinit *{Service user}*

  c. Upgrade the package in HDFS.

  sh update_hdfs_file.sh

  d. (Optional) Roll back upgrade after the patch is uninstalled.

  sh rollback_hdfs_file.sh

3.9.4.3 MRS 2.1.0.11 Patch Description

Basic Information

Table 3-62 Basic information

Patch Version	MRS 2.1.0.11
Release Date	2020-12-30

Resolved Issues	**List of resolved issues in MRS 2.1.0.11:** **MRS Manager** Executor, KNOX, and OS logs can be rolled back. Executor GC logs are now added. The Knox restart failure is resolved. Resolve the problem that jobs fail to be submitted when a node is faulty. Full-link monitoring is supported. Job Status can be updated when switching over the ResourceManager active and standby nodes. Backup and restoration fail in some scenarios is resolved. The process fault alarm that frequently generated on the HMaster is resolved. **Big data components** Resolved the issue of JobHistory memory leakage. Resolved the issue of the Hive truncate table times out and fails to be truncated. Resolved the issue that the table data file does not exist after an incremental Hive task fails. The Hive SQL statement is not running properly. After a Carbon table is created in a security cluster and the hive group does not have the permission to create the Carbon table, other users can create the Carbon table. Resolved the problem that the spark JDBCServer process is abnormal.
	List of resolved issues in MRS 2.1.0.10: **MRS Manager** New queue configurations in the **capacity-schedule.xml** file will not be lost during cluster scale-out after the patch is installed. Full-link monitoring can be rolled back. **Big data components** Hive permission assignment failure on Spark is resolved. If no queue is specified, tasks are submitted to the launcher-job queue by default. Task running will not be affected.

	List of resolved issues in MRS 2.1.0.9: **MRS Manager** The MRS Executor memory overflow is resolved. The cluster scale-out process is optimized. The problem that the SQL statement is incorrectly combined when the value of SparkSQL contains spaces is resolved. The problem that HiveSQL jobs fail to be submitted occasionally is resolved. The permission control for downloading the keytab file is optimized. **Big data components** When the Presto role name contains uppercase letters, the permission model can take effect. The problem that Hive partitions are deleted slowly is resolved. The problem that the token expires after Spark runs for a long time is resolved.
	List of resolved issues in MRS 2.1.0.8: **MRS Manager** The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved. Multiple users can log in to MRS Manager at the same time. Full-link monitoring is supported. **MRS big data components** Carbon 2.0 has been upgraded. The HBASE-18484 issue has been solved.
	List of the resolved issues in MRS 2.1.0.7: **MRS Manager** The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved. The Presto query result can be saved as a file.

	List of resolved issues in MRS 2.1.0.6: **MRS Manager** The problem that the disk I/O usage of monitoring data is inaccurate has been solved. The problem that the Spark job status is not updated occasionally has been solved. The problem that the job running failure has been solved. The patch mechanism has been optimized. **MRS big data components** The HBase exceptions are rectified. The problem that the system responds slowly when Hive roles are bound to permissions has been solved.
	List of resolved issues in MRS 2.1.0.5: **MRS big data components** Impala supports the ObsFileSystem function. The timeout period of the MRS Manager page and the native pages of components can be configured. The Hive privilege binding freezing problem has been solved. The data connection failure has been solved.
	List of resolved issues in MRS 2.1.0.3: **MRS Manager** Problems of Manager executor's high concurrent job submission have been solved. **MRS big data components** Data insertion failure in hive on tez has been fixed.

	List of the resolved issues in MRS 2.1.0.2: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. **manager executor** can be used to configure high concurrency. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.
	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.
Compatibility with Other Patches	The MRS 2.1.0.11 patch package contains all patches released for MRS 2.1.0.
Vulnerability Disclosure	The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see **CVE-2020-9480**

Impact of Patch Installation

- During the installation of the MRS 2.1.0.11 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, Yarn, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are

temporarily unavailable but services are not interrupted during the rolling restart.

- After the MRS 2.1.0.11 patch is installed, log in to the **standby Master node** (Log in to MRS Manager. The Master node with a hollow pentagon on the **Host Management** page is the standby Master node), switch to user **omm**, and run the **sh /opt/knox/bin/restart-knox.sh** command to restart the Knox process. This operation is not required for a cluster with only one Master node.

  You can run the **ps -ef |grep knox** command to check whether the knox process is started. If the knox process ID is displayed, the knox process is started successfully.

- (Optional) After installing the MRS 2.1.0.11 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:

  a. Change the session timeout interval of the web and CAS services on all Master nodes.

     i. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml**. The unit is minute.

     ii. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml**. The unit is minute.

  b. Change the TGT validity period of the CAS on all Master nodes.

     Change **1200** in **p:maxTimeToLiveInSeconds="${tgt.maxTimeToLiveInSeconds:1200}** and **p:timeToKillInSeconds="${tgt.timeToKillInSeconds:1200}"** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml** to the corresponding timeout interval, in seconds.

  c. Restart the Tomcat service on the active Master node.

     i. On the active Master node, run the **netstat -anp |grep 28443 |grep LISTEN** command as user **omm** to query the Tomcat process ID.

ii. Run the **kill -9** *{pid}* command, in which *{pid}* indicates the process ID obtained in the previous step.

iii. Wait for the process to automatically restart. You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is started. If the command output is displayed, the process is started successfully.

d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.

- HDFS/MapReduce/Yarn: Add the customized configuration item **http.server.session.timeout.secs**.
- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

3.9.4.4 MRS 3.0.5.1 Patch Description

Basic Information

Table 3-63 Basic information

Patch Version	MRS 3.0.5.1
Release Date	2021-08-14

Resolved Issues	**List of resolved issues in MRS 3.0.5.1:** **MRS Manager** ● Resolved the failure to submit SparkSQL jobs on the job management page due to long SQL statements. ● Resolved the failure to execute SQL statements with comments. **Big data components** ● Resolved the failure to synchronize IAM users with ClickHouse clusters. ● Resolved the issue that the Flume client in the cluster cannot use an agency to access OBS. ● Resolved the issue that the value of **% of Queue** is not displayed for a specified job on the native Yarn web UI. ● Resolved the issue that job logs are incompletely displayed on the native Yarn web UI. ● Resolved the issue that temporary files reside in HDFS after execution of Hive jobs. ● Resolved the incompatibility in the interconnection between open-source Sqoop 1.4.7 and MRS Hive. ● Resolved the failure to query Avro tables through Hive on MR. ● Resolved the memory leak issue caused when HiveServer loads user-defined functions (UDFs). ● Resolved the issue that the execution results of Hive and SparkSQL time functions are inconsistent. ● Resolved the issue (HIVE-20187) that the result is incorrect when Hive on Tez uses MapJoin to achieve performance tuning. ● Resolved the issue that an error occurs when the **beeline -p** command is executed. ● Resolved the issue that Hue fails to format SQL statements. ● Resolved the failure to submit Oozie jobs due to the incompatibility between Hue and Oozie time zones. ● Resolved the unavailability of the variable drop-down list when a variable-declared Hive SQL statement is executed on the Hue web UI. ● Resolved the query failure caused by incorrectly closed sessions when Hue connects to Hive for queries. ● Resolved slow responses to Kunpeng servers' queries of Kudu tables using Impala. ● Resolved the failure to install the Kudu client. ● Resolved the unexpected restarts of KuduMaster instance on Kunpeng servers. ● Resolved the search exceptions on the Ranger web UI. ● Resolved the failure to redirect users to the login page after the logout from the Ranger web UI.

Compatibility with Other Patches	The MRS 3.0.5.1 patch can resolve all the issues detected in MRS 3.0.5.

Impact of Patch Installation

- During the installation of MRS 3.0.5.1, the executor and controller processes are automatically restarted and cluster functions on the management plane, such as job submission and cluster scaling, will be affected. Therefore, install the patch at an appropriate time.
- After you install the patch, restart the Spark2x, Hive, Yarn, Impala, Kudu, and Hue components on FusionInsight Manager for the patch to take effect. During the restart, some services may be unavailable for a short period of time. To minimize the impact on service continuity, perform the restart at a proper time.
- To install the MRS 3.0.5.1 patch, you need to manually download the patch file and install it on any master node in the cluster. For details, see the **README.md** file in the patch package.
- This patch must be also installed for any new node subsequently added to the cluster. To install the patch for this new node, install the patch on the master node and restart the corresponding service.

Patch Download Addresses

- **RU-Moscow2**: **https://mrs-container1-patch-ru-northwest-2.obs.ru-northwest-2.myhuaweicloud.com/MRS_Common_Script/MRS_3.0.5.1_Patch_All_20210724.tar.gz**

3.9.4.5 MRS 2.1.0.10 Patch Description

Basic Information

Table 3-64 Basic information

Patch Version	MRS 2.1.0.10
Release Date	2020-09-21

Resolved Issues	**List of resolved issues in MRS 2.1.0.10:** **MRS Manager** New queue configurations in the **capacity-schedule.xml** file will not be lost during cluster scale-out after the patch is installed. Full-link monitoring can be rolled back. **Big data components** Hive permission assignment failure on Spark is resolved. If no queue is specified, tasks are submitted to the launcher-job queue by default. Task running will not be affected.
	List of resolved issues in MRS 2.1.0.9: **MRS Manager** The MRS Executor memory overflow is resolved. Optimized the cluster scale-out process. The problem that the SQL statement is incorrectly combined when the value of SparkSQL contains spaces is resolved. The problem that HiveSQL jobs fail to be submitted occasionally is resolved. The permission control for downloading the keytab file is optimized. **Big data components** When the Presto role name contains uppercase letters, the permission model can take effect. The problem that Hive partitions are deleted slowly is resolved. The problem that the token expires after Spark runs for a long time is resolved.
	List of resolved issues in MRS 2.1.0.8: **MRS Manager** The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved. Multiple users can log in to MRS Manager at the same time. Full-link monitoring is supported. **MRS big data components** Carbon 2.0 has been upgraded. The HBASE-18484 issue has been solved.
	List of the resolved issues in MRS 2.1.0.7: **MRS Manager** The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved. The Presto query result can be saved as a file.

	List of resolved issues in MRS 2.1.0.6: **MRS Manager** The problem that the disk I/O usage of monitoring data is inaccurate has been solved. The problem that the Spark job status is not updated occasionally has been solved. The problem that the job running failure has been solved. The patch mechanism has been optimized. **MRS big data components** The HBase exceptions are rectified. The problem that the system responds slowly when Hive roles are bound to permissions has been solved.
	List of resolved issues in MRS 2.1.0.5: **MRS big data components** Impala supports the ObsFileSystem function. The timeout period of the MRS Manager page and the native pages of components can be configured. The Hive privilege binding freezing problem has been solved. The data connection failure has been solved.
	List of resolved issues in MRS 2.1.0.3: **MRS Manager** Problems of Manager executor's high concurrent job submission have been solved. **MRS big data components** Data insertion failure in **hive on tez** has been fixed.

	List of the resolved issues in MRS 2.1.0.2: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. **manager executor** can be used to configure high concurrency. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.
	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.
Compatibility with Other Patches	The MRS 2.1.0.10 patch package contains all patches released for MRS 2.1.0.
Vulnerability Disclosure	The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see **CVE-2020-9480**

Impact of Patch Installation

- During the installation of the MRS 2.1.0.10 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, Yarn, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are

temporarily unavailable but services are not interrupted during the rolling restart.

- After installing the MRS 2.1.0.10 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:

  a. Change the session timeout interval of the web and CAS services on all Master nodes.

     i. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml**. The unit is minute.

     ii. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml**. The unit is minute.

  b. Change the TGT validity period of the CAS on all Master nodes.

     Change **1200** in **p:maxTimeToLiveInSeconds="${tgt.maxTimeToLiveInSeconds:1200}** and **p:timeToKillInSeconds="${tgt.timeToKillInSeconds:1200}"** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml** to the corresponding timeout interval, in seconds.

  c. Restart the Tomcat service on the active Master node.

     i. On the active Master node, run the **netstat -anp |grep 28443 |grep LISTEN** command as user **omm** to query the Tomcat process ID.

     ii. Run the **kill -9** *{pid}* command, in which *{pid}* indicates the process ID obtained in the previous step.

     iii. Wait for the process to automatically restart. You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is started. If the command output is displayed, the process is started successfully.

  d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.

- HDFS/MapReduce/Yarn: Add the customized configuration item **http.server.session.timeout.secs**.
- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

3.9.4.6 MRS 2.1.0.9 Patch Description

Basic Information

Table 3-65 Basic information

Patch Version	MRS 2.1.0.9
Release Date	2020-08-21
Resolved Issues	**List of resolved issues in MRS 2.1.0.9:** **MRS Manager** The MRS Executor memory overflow is resolved. Optimized the disk scale-out process. The problem that the SQL statement is incorrectly combined when the value of SparkSQL contains spaces is resolved. The problem that HiveSQL jobs fail to be submitted occasionally is resolved. The permission control for downloading the keytab file is optimized. **Big data components** When the Presto role name contains uppercase letters, the permission model can take effect. Partitions are deleted slowly in Hive. The problem that the token expires after Spark runs for a long time is resolved.

	List of resolved issues in MRS 2.1.0.8: **MRS Manager** The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved. Multiple users can log in to MRS Manager at the same time. Full-link monitoring is supported. **MRS big data components** Carbon 2.0 has been upgraded. The HBASE-18484 issue has been solved.
	List of the resolved issues in MRS 2.1.0.7: **MRS Manager** The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved. The Presto query result can be saved as a file.
	List of resolved issues in MRS 2.1.0.6: **MRS Manager** The problem that the disk I/O usage of monitoring data is inaccurate has been solved. The problem that the Spark job status is not updated occasionally has been solved. The problem that the job running failure has been solved. The patch mechanism has been optimized. **MRS big data components** The HBase exceptions are rectified. The problem that the system responds slowly when Hive roles are bound to permissions has been solved.
	List of resolved issues in MRS 2.1.0.5: **MRS big data components** Impala supports the ObsFileSystem function. The timeout period of the MRS Manager page and the native pages of components can be configured. The Hive privilege binding freezing problem has been solved. The data connection failure has been solved.
	List of resolved issues in MRS 2.1.0.3: **MRS Manager** Problems of Manager executor's high concurrent job submission have been solved. **MRS big data components** Data insertion failure in **hive on tez** has been fixed.

	List of the resolved issues in MRS 2.1.0.2: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. **manager executor** can be used to configure high concurrency. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.
	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.
Compatibility with Other Patches	The MRS 2.1.0.9 patch package contains all patches released for MRS 2.1.0.
Vulnerability Disclosure	The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see **CVE-2020-9480**

Impact of Patch Installation

- During the installation of the MRS 2.1.0.9 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, Yarn, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are

temporarily unavailable but services are not interrupted during the rolling restart.

- After installing the MRS 2.1.0.9 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:

  a. Change the session timeout interval of the web and CAS services on all Master nodes.

     i. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml**. The unit is minute.

     ii. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml**. The unit is minute.

  b. Change the TGT validity period of the CAS on all Master nodes.

     Change **1200** in **p:maxTimeToLiveInSeconds="${tgt.maxTimeToLiveInSeconds:1200}** and **p:timeToKillInSeconds="${tgt.timeToKillInSeconds:1200}"** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml** to the corresponding timeout interval, in seconds.

  c. Restart the Tomcat service on the active Master node.

     i. On the active Master node, run the **netstat -anp |grep 28443 |grep LISTEN** command as user **omm** to query the Tomcat process ID.

     ii. Run the **kill -9** *{pid}* command, in which *{pid}* indicates the process ID obtained in the previous step.

     iii. Wait for the process to automatically restart. You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is started. If the command output is displayed, the process is started successfully.

  d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.

- HDFS/MapReduce/Yarn: Add the customized configuration item **http.server.session.timeout.secs**.
- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

3.9.4.7 MRS 2.1.0.8 Patch Description

Basic Information

Table 3-66 Basic information

Patch Version	MRS 2.1.0.8
Release Date	2020-08-04
Resolved Issues	**List of resolved issues in MRS 2.1.0.8:** **MRS Manager** The problem that the ECS API traffic is limited when OBS is accessed through an agency has been solved. Multiple users can log in to MRS Manager at the same time. Full-link monitoring is supported. **MRS big data components** Carbon 2.0 has been upgraded. The HBASE-18484 issue has been solved.
	List of the resolved issues in MRS 2.1.0.7: **MRS Manager** The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved. The Presto query result can be saved as a file.

	List of resolved issues in MRS 2.1.0.6: **MRS Manager** The problem that the disk I/O usage of monitoring data is inaccurate has been solved. The problem that the Spark job status is not updated occasionally has been solved. The problem that the job running failure has been solved. The patch mechanism has been optimized. **MRS big data components** The HBase exceptions are rectified. The problem that the system responds slowly when Hive roles are bound to permissions has been solved.
	List of resolved issues in MRS 2.1.0.5: **MRS big data components** Impala supports the ObsFileSystem function. The timeout period of the MRS Manager page and the native pages of components can be configured. The Hive privilege binding freezing problem has been solved. The data connection failure has been solved.
	List of resolved issues in MRS 2.1.0.3: **MRS Manager** Problems of Manager executor's high concurrent job submission have been solved. **MRS big data components** Data insertion failure in hive on tez has been fixed.

	List of the resolved issues in MRS 2.1.0.2: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. **manager executor** can be used to configure high concurrency. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.
	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.
Compatibility with Other Patches	The MRS 2.1.0.8 patch package contains all patches released for MRS 2.1.0.
Vulnerability Disclosure	The remote code execution vulnerability of the Spark has been fixed. For details about the vulnerability, see **CVE-2020-9480**

Impact of Patch Installation

- During the installation of the MRS 2.1.0.8 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, Yarn, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are

temporarily unavailable but services are not interrupted during the rolling restart.

- After installing the MRS 2.1.0.8 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:

  a. Change the session timeout interval of the web and CAS services on all Master nodes.

     i. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml**. The unit is minute.

     ii. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml**. The unit is minute.

  b. Change the TGT validity period of the CAS on all Master nodes.

     Change **1200** in **p:maxTimeToLiveInSeconds="${tgt.maxTimeToLiveInSeconds:1200}** and **p:timeToKillInSeconds="${tgt.timeToKillInSeconds:1200}"** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml** to the corresponding timeout interval, in seconds.

  c. Restart the Tomcat service on the active Master node.

     i. On the active Master node, run the **netstat -anp |grep 28443 |grep LISTEN** command as user **omm** to query the Tomcat process ID.

     ii. Run the **kill -9** *{pid}* command, in which *{pid}* indicates the process ID obtained in the previous step.

     iii. Wait for the process to automatically restart. You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is started. If the command output is displayed, the process is started successfully.

  d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.

- HDFS/MapReduce/Yarn: Add the customized configuration item **http.server.session.timeout.secs**.
- Spark: Change the value of **spark.session.maxAge**.
- Hive: Add the customized configuration item **http.server.session.timeout.secs**.

When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

3.9.4.8 MRS 2.1.0.7 Patch Description

Basic Information

Table 3-67 Basic information

Patch Version	MRS 2.1.0.7
Release Date	2020-07-15
Resolved Issues	**List of the resolved issues in MRS 2.1.0.7**: **MRS Manager** The problem that data and files are displayed incorrectly if a field contains a newline character in the DLF+Presto query has been solved. The Presto query result can be saved as a file.
	List of resolved issues in MRS 2.1.0.6: **MRS Manager** The problem that the disk I/O usage of monitoring data is inaccurate has been solved. The problem that the Spark job status is not updated occasionally has been solved. The problem that the job running failure has been solved. The patch mechanism has been optimized. **MRS big data components** The HBase exceptions are rectified. The problem that the system responds slowly when Hive roles are bound to permissions has been solved.

	List of resolved issues in MRS 2.1.0.5: **MRS big data components** Impala supports the ObsFileSystem function. The timeout period of the MRS Manager page and the native pages of components can be configured. The Hive privilege binding freezing problem has been solved. The data connection failure has been solved.
	List of resolved issues in MRS 2.1.0.3: **MRS Manager** Problems of Manager executor's high concurrent job submission have been solved. **MRS big data components** Data insertion failure in hive on tez has been fixed.
	List of the resolved issues in MRS 2.1.0.2: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. **manager executor** can be used to configure high concurrency. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.
	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.

Compatibility with Other Patches	The MRS 2.1.0.7 patch package contains all patches released for MRS 2.1.0.

Impact of Patch Installation

- During the installation of the MRS 2.1.0.7 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, Yarn, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 2.1.0.7 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:

  a. Change the session timeout interval of the web and CAS services on all Master nodes.

     i. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml**. The unit is minute.

     ii. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml**. The unit is minute.

  b. Change the TGT validity period of the CAS on all Master nodes.

     Change **1200** in **p:maxTimeToLiveInSeconds="${tgt.maxTimeToLiveInSeconds:1200}** and **p:timeToKillInSeconds="${tgt.timeToKillInSeconds:1200}"** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml** to the corresponding timeout interval, in seconds.

c. Restart the Tomcat service on the active Master node.

   i. On the active Master node, run the **netstat -anp |grep 28443 |grep LISTEN** command as user **omm** to query the Tomcat process ID.

   ii. Run the **kill -9** *{pid}* command, in which *{pid}* indicates the process ID obtained in the previous step.

   iii. Wait for the process to automatically restart. You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is started. If the command output is displayed, the process is started successfully.

d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.

   - HDFS/MapReduce/Yarn: Add the customized configuration item **http.server.session.timeout.secs**.
   - Spark: Change the value of **spark.session.maxAge**.
   - Hive: Add the customized configuration item **http.server.session.timeout.secs**.

   When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

3.9.4.9 MRS 2.1.0.6 Patch Description

Basic Information

Table 3-68 Basic information

Patch Version	MRS 2.1.0.6
Release Date	2020-06-10
Resolved Issues	**List of resolved issues in MRS 2.1.0.6:** **MRS Manager** The problem that the disk I/O usage of monitoring data is inaccurate has been solved. The problem that the Spark job status is not updated occasionally has been solved. The problem that the job running failure has been solved. The patch mechanism has been optimized. **MRS big data components** The HBase exceptions are rectified. The problem that the system responds slowly when Hive roles are bound to permissions has been solved.

	List of resolved issues in MRS 2.1.0.5: **MRS big data components** Impala supports the ObsFileSystem function. The timeout period of the MRS Manager page and the native pages of components can be configured. The Hive privilege binding freezing problem has been solved. The data connection failure has been solved.
	List of resolved issues in MRS 2.1.0.3: **MRS Manager** Problems of Manager executor's high concurrent job submission have been solved. **MRS big data components** Data insertion failure in hive on tez has been fixed.
	List of the resolved issues in MRS 2.1.0.2: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. **manager executor** can be used to configure high concurrency. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.
	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.

Compatibility with Other Patches	The MRS 2.1.0.6 patch package contains all patches released for MRS 2.1.0.

Impact of Patch Installation

- During the installation of the MRS 2.1.0.6 patch, MRS Manager will be restarted, and the components such as Hive, Impala, Spark, HDFS, Yarn, MapReduce, Presto, HBase, Tez, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 2.1.0.6 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.
- (Optional) The timeout interval of the MRS Manager page and the native page of the component can be configured. You need to manually modify the following configuration:

  a. Change the session timeout interval of the web and CAS services on all Master nodes.

     i. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/web.xml**. The unit is minute.

     ii. Change the value of **<session-timeout>20</session-timeout>** in **/opt/Bigdata/tomcat/webapps/web/WEB-INF/web.xml**. The unit is minute.

  b. Change the TGT validity period of the CAS on all Master nodes.

     Change **1200** in **p:maxTimeToLiveInSeconds="${tgt.maxTimeToLiveInSeconds:1200}** and **p:timeToKillInSeconds="${tgt.timeToKillInSeconds:1200}"** in **/opt/Bigdata/tomcat/webapps/cas/WEB-INF/spring-configuration/ticketExpirationPolicies.xml** to the corresponding timeout interval, in seconds.

c. Restart the Tomcat service on the active Master node.

   i. On the active Master node, run the **netstat -anp |grep 28443 |grep LISTEN** command as user **omm** to query the Tomcat process ID.

   ii. Run the **kill -9** *{pid}* command, in which *{pid}* indicates the process ID obtained in the previous step.

   iii. Wait for the process to automatically restart. You can run the **netstat -anp |grep 28443 |grep LISTEN** command to check whether the process is started. If the command output is displayed, the process is started successfully.

d. Add or modify configuration items for each component. The values of the configuration items are the same as the timeout interval, in seconds.

   - HDFS/MapReduce/Yarn: Add the customized configuration item **http.server.session.timeout.secs**.
   - Spark: Change the value of **spark.session.maxAge**.
   - Hive: Add the customized configuration item **http.server.session.timeout.secs**.

   When saving the configuration items, you can choose not to restart the affected services or instances. Restart the services or instances when the service is not busy.

3.9.4.10 MRS 2.1.0.3 Patch Description

Basic Information

Table 3-69 Basic information

Patch Version	MRS 2.1.0.3
Release Date	2020-04-29
Resolved Issues	**List of resolved issues in MRS 2.1.0.3:** **MRS Manager** Problems of Manager executor's high concurrent job submission have been solved. **MRS big data components** Data insertion failure in hive on tez has been fixed.

	List of the resolved issues in MRS 2.1.0.2: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. **manager executor** can be used to configure high concurrency. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.
	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275.
Compatibility with Other Patches	The MRS 2.1.0.3 patch package contains all patches released for MRS 2.1.0.

Impact of Patch Installation

- During the installation of the MRS 2.1.0.3 patch, MRS Manager will be restarted, and the components such as Hive, Spark, HDFS, Yarn, MapReduce, Presto, HBase, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 2.1.0.3 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).

- For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
- For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
- For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.11 MRS 2.1.0.2 Patch Description

Basic Information

Table 3-70 Basic information

Patch Version	MRS 2.1.0.2
Release Date	2020-04-22
Resolved Issues	**List of the resolved issues in MRS 2.1.0.2**: **MRS Manager** No monitoring information is displayed after NodeAgent is restarted. When a job is under submission for a long time, memory overflow occurs in the **manager executor** process. Job submission is supported. You can configure the concurrency for **manager executor**. New Kafka topics are not displayed on the MRS Manager management plane. When you call security cluster's APIs to submit the **Spark Submit** job and perform operations on an HBase table, the permission control on the HBase table does not take effect. The MRS Manager patch mechanism has been optimized. **MRS big data components** The slow running of the **load data inpath** command executed by Spark has been optimized. Column names containing the dollar sign ($) can be used in Spark table creation. OBS-related problems have been solved.

	List of the resolved issues in MRS 2.1.0.1: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275
Compatibility with Other Patches	The MRS 2.1.0.2 patch package contains all content of the MRS 2.1.0.1 patch package.

Impact of Patch Installation

- During the installation of the MRS 2.1.0.2 patch, MRS Manager will be restarted, and the components such as Hive, Spark, HDFS, Yarn, MapReduce, Presto, HBase, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable, and services are not interrupted during the rolling restart.
- After installing the MRS 2.1.0.2 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.12 MRS 2.1.0.1 Patch Description

Basic Information

Table 3-71 Basic information

Patch Version	MRS 2.1.0.1

Release Date	2020-02-12
Resolved Issues	**List of the resolved issues in MRS 2.1.0.1**: **MRS Manager** The return results of Hive SQL statements submitted by V2 jobs have been optimized, and the issue that V2 jobs fail to be submitted using an agency token has been solved. **MRS big data components** HiveServer out of memory (OOM) has been solved for MRS Hive: HIVE-10970 and HIVE-22275
Compatibility with Other Patches	None

Impact of Patch Installation

During the installation of MRS 2.1.0.1 patches, MRS Manager and Hive are restarted. During the restart, the services are temporarily unavailable.

After MRS 2.1.0.1 patches are installed, log in to the Master1 node of the MRS cluster and delete the job directory in HDFS.

- For a cluster with Kerberos authentication disabled, run the following command to delete the job directory in HDFS:

  hdfs dfs -rm -r /mrs/mrsjob/hive

- For a cluster with Kerberos authentication enabled, perform the following operations to delete the job directory in HDFS:

  a. Run the following command and enter the password to perform authentication.

     kinit hdfs

  b. Run the following command to delete the job directory in HDFS:

     hdfs dfs -rm -r /mrs/mrsjob/hive

NOTE

This step is not required for a new MRS cluster because the directory does not exist in HDFS.

3.9.4.13 MRS 2.0.6.1 Patch Description

Basic Information

Table 3-72 Basic information

Patch Version	MRS 2.0.6.1
Release Date	2020-07-06
Resolved Issues	**List of the resolved issues in MRS 2.1.0.1:** **MRS Manager** Patch mechanism The monitoring metrics are empty occasionally. In DLF+Presto query, if a field contains a newline character, data and files are displayed incorrectly.
Compatibility with Other Patches	None

Impact of Patch Installation

During the installation of the MRS 2.0.6.1 patch, MRS Manager will be restarted, and the components such as Hive and services with dependency will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable, and services are not interrupted during the rolling restart.

3.9.4.14 MRS 2.0.1.3 Patch Description

Basic Information

Table 3-73 Basic information

Patch Version	MRS 2.0.1.3
Release Date	2019-12-25

Resolved Issues	**List of the resolved issues in MRS 2.0.1.3**: **MRS Manager** The cluster scaling logic has been optimized and the TCP connection leak issues have been solved at the background of V1 job management APIs. **MRS big data components** The following issues have been solved for MRS Hive: HiveServer out of memory (OOM); slow **MergeFile** phase if a large number of small files exist; files not found in the **load partition** phase of **insert overwrite**, and file merging failure during **HIVE-22373:Container** reuse.
	List of the resolved issues in MRS 2.0.1.2: **MRS Manager** The following issue has been solved: Scale-out fails occasionally due to timeout that occurs when ResourceManager executes **refreshNodes**.
	List of the resolved issues in MRS 2.0.1.1: **MRS Manager** The following issue has been solved: OOM occurs on the executor of MRS Master nodes due to repeated node scale-in or scale-out. **MRS big data components** The following new function has been added: MRS Presto supports OBSFileSystem. The following issues have been solved for MRS Presto: The jstack is frequently printed and log files are too large to scroll.
Compatibility with Other Patches	The MRS 2.0.1.3 patch package contains all content of the MRS 2.0.1.2 and MRS 2.0.1.1 patch packages.

Impact of Patch Installation

During the installation of MRS 2.0.1.3 patches, MRS Manager and Presto are restarted. During the restart, the services are temporarily unavailable.

3.9.4.15 MRS 2.0.1.2 Patch Description

Basic Information

Table 3-74 Basic information

Patch Version	MRS 2.0.1.2

Release Date	2019-09-30
Resolved Issues	**List of the resolved issues in MRS 2.0.1.2**: **MRS Manager** The following issue has been solved: Scale-out fails occasionally due to timeout that occurs when ResourceManager executes **refreshNodes**.
	List of the resolved issues in MRS 2.0.1.1: **MRS Manager** The following issue has been solved: Out of memory (OOM) occurs on the executor of MRS Master nodes due to repeated node scale-in or scale-out. **MRS big data components** The following new function has been added: MRS Presto supports OBSFileSystem. The following issues have been solved for MRS Presto: The jstack is frequently printed and log files are too large to scroll.
Compatibility with Other Patches	The MRS 2.0.1.2 patch package contains all content of the MRS 2.0.1.1 patch package.

Impact of Patch Installation

During the installation of MRS 2.0.1.2 patches, MRS Manager and Presto are restarted. During the restart, the services are temporarily unavailable.

3.9.4.16 MRS 2.0.1.1 Patch Description

Basic Information

Table 3-75 Basic information

Patch Version	MRS 2.0.1.1
Release Date	2019-09-30

Resolved Issues	**MRS Manager** The following issue has been solved: Out of memory (OOM) occurs on the executor of MRS Master nodes due to repeated node scale-in or scale-out. **MRS big data components** The following new function has been added: MRS Presto supports OBSFileSystem. The following issues have been solved for MRS Presto: The jstack is frequently printed and log files are too large to scroll.
Compatibility with Other Patches	None

Impact of Patch Installation

During the installation of MRS 2.0.1.1 patches, MRS Manager and Presto are restarted. During the restart, the services are temporarily unavailable.

3.9.4.17 MRS 1.9.3.3 Patch Description

Basic Information

Table 3-76 Basic information

Patch Version	MRS 1.9.3.3
Release Date	2021-01-04
Resolved Issues	**List of resolved issues in MRS 1.9.3.3:** **MRS Manager** Resolved the node isolation problem. **MRS big data components** Resolved the memory leak issue when Hive loads hooks.
	List of resolved issues in MRS 1.9.3.2: **MRS big data components** When the insert overwrite operation is performed using Spark SQL and Beeline, old files cannot be trashed.

	List of resolved issues in MRS 1.9.3.1: **MRS Manager** Solved the problem that Task nodes fail to be removed from a custom cluster. **MRS big data components** Solved the problem that the version of the **adapter-hadoop-wrapper-file-system** package in the Hive and Spark paths is incorrect. Solved the problem that multiple namespaces saved on FusionInsight Manager of HBase do not take effect in the background. Added HDFSWrapper to support AbstractFileSystem.
Compatibility with Other Patches	The MRS 1.9.3.3 patch package contains all patches released for MRS 1.9.3.

Impact of Patch Installation

- During the installation of the MRS 1.9.3.3 patch, MRS Manager is restarted, and Hadoop, HDFS, Hive, Spark, and related dependent services are restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.3.3 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.18 MRS 1.9.3.1 Patch Description

Basic Information

Table 3-77 Basic information

Patch Version	MRS 1.9.3.1
Release Date	2020-09-04
Resolved Issues	**MRS Manager** Solved the problem that Task nodes fail to be removed from a custom cluster. **MRS big data components** Solved the problem that the version of the **adapter-hadoop-wrapper-file-system** package in the Hive and Spark paths is incorrect. Solved the problem that multiple namespaces saved on FusionInsight Manager of HBase do not take effect in the background. Added HDFSWrapper to support AbstractFileSystem.
Compatibility with Other Patches	None

Impact of Patch Installation

- During the installation of the MRS 1.9.3.1 patch, MRS Manager is restarted, and Hadoop, HDFS, Hive, Spark, and related dependent services are restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.3.1 patch, you need to download and install all clients again, including the original clients of the Master node and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.19 MRS 1.9.2.2 Patch Description

Basic Information

Table 3-78 Basic information

Patch Version	MRS 1.9.2.2
Release Date	2021-05-18
Resolved Issues	**MRS Manager** Resolved the sudo privilege escalation vulnerability. Resolved the issue of queue information loss due to queue update during capacity expansion. **MRS big data components** Resolved the issue that the Hive on Spark task is suspended because the block ID is displayed as garbled characters. Self-developed APIs are added to Hive. Resolved the issue that the **map.xml** file cannot be read. Optimized the Hive Har feature. Resolved the issue that Yarn is unavailable due to ZooKeeper dirty data. Upgraded the OBS packages. Upgraded the JDK version. Resolved the issue of ResourceManager memory leakage. Added the monitoring on the exception that occurs when the ECS getSecuritykey API is called. Optimized the temporary AK/SK process. Resolved the issue of ResourceManager memory leakage. Fixed the error reported when the Hive union statement is used to merge small files. Resolved the issue that the Hadoop task fails to be executed due to insufficient space. Resolved the issue that no data is generated after a Hive job is successfully executed.

Compatibility with Other Patches	None

Impact of Patch Installation

- During the installation of the MRS 1.9.2.2 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, Presto, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After the MRS 1.9.2.2 patch is installed, you need to restart the OMS service.

  NOTE

  - Log in to the active and standby OMS nodes as user **root**, switch to user **omm**, and run the **sh ${BIGDATA_HOME}/om-0.0.1/sbin/restart-oms.sh** command to restart the OMS service.
  - Both the active and standby OMS nodes need to be restarted.
- After installing the MRS 1.9.2.2 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

  NOTE

  - You are advised to back up the old clients before reinstalling the new ones.
  - If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.20 MRS 1.9.0.8, 1.9.0.9, and 1.9.0.10 Patch Description

Basic Information

Table 3-79 Basic information

Resolved Issues	Patch number: **MRS 1.9.0.10** Release date: January 17, 2023 **Resolved issues** **MRS big data components** OBSA supports flow control retry.
	Patch number: **MRS 1.9.0.9** Release date: August 10, 2022 **Resolved issues** **MRS big data components** Superior scheduling algorithm optimization
	Patch number: **MRS 1.9.0.8** Release date: February 20, 2021 **Resolved issues** **MRS big data components** Added the monitoring on the exception that occurs when the ECS getSecuritykey API is called. Optimized the temporary AK/SK process. Resolved the issue of ResourceManager memory leakage.
	List of resolved issues in MRS 1.9.0.7: **MRS Manager** Resolved the issue of queue information loss due to queue update during capacity expansion. **MRS big data components** Resolved the issue that the Hive on Spark task is suspended because the block ID is displayed as garbled characters. Solve the problem that the Hadoop task fails to be executed due to insufficient space. Self-developed APIs are added to Hive. Resolved the issue that the **map.xml** file cannot be read. Resolved the issue that Yarn is unavailable due to ZooKeeper dirty data. Resolved the issue of ResourceManager memory leakage. Optimized the Hive Har feature. Upgraded the OBS packages. Upgraded the JDK version.

	List of resolved issues in MRS 1.9.0.6: **MRS Manager** MRS Manager supports scale-in of specified nodes in a yearly/monthly cluster. **MRS big data components** Resolved the issue of slow response when HiveSE delivers SQL statements. Supported the JobHistory query failure information interface. Resolved the issue that fine-grained permissions do not take effect. Fixed the exception that occurs when Hive on Spark reads data. Resolved the issue that data volume increases when the Hive on mrs task is executed twice. Resolved the issue that the performance of some strings is poor when vector-based vectorized query is enabled in Hive.

	List of resolved issues in MRS 1.9.0.5: **MRS Manager** Optimized the service restart process during configuration saving on MRS Manager. Rectified the periodic backup failure on MRS Manager. **MRS big data components** Private patch of Ranger Resolved the JVM Create GC thread failed issue of Yarn. Added the HiveServer2 task stacking alarm. Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s. Added the alarm indicating that HiveServer2 uncomment ZooKeeper. Added the alarm indicating that Yarn tasks failed and the number of killed tasks exceeds 5 within 20 minutes. Corrected time zone of Spark JobHistory. Optimized MetaStore restart mechanism. Resolved the HIVE-22771 open-source issue. Resolved the Hive beeline log printing errors. Corrected the number of active nodes displayed on the Yarn page. Resolved the slow response of RM page display, which is caused by large number of RM threads. Supported OBS monitoring. Upgraded the OBS packages. Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc. Resolved the issue that Hive occasionally reports a Kryo deserialization failure. Resolved the issue of Spark JobHistory memory leakage. Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.
	List of resolved issues in MRS 1.9.0.3: **MRS Manager** Upgraded ARM JDK on MRS Manager. Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager. **MRS big data components** Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.

	List of resolved issues in MRS 1.9.0.2: **MRS Manager** Resolved the mutual trust lost between some Core nodes in the cluster. Resolved the issue that instances fail to be added after the patch is installed. Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager. **MRS big data components** Upgraded the OBS packages.
	List of resolved issues in MRS 1.9.0.1: **MRS Manager** Resolved the issue that MRS Manager does not support rolling patch installation without restarting services. **MRS big data components** Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes. Resolved the issue that Kafka does not support open-source access. Resolved the SPARK-27637 open-source issue. Optimized the Hive rolling restart. Upgraded the OBS packages.
Compatibility with Other Patches	The MRS 1.9.0.10 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.

Impact of Patch Installation

- During the installation of the MRS 1.9.0.10 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Presto, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable, and services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.10 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

- (Optional) In the scenario where a temporary AK/SK is obtained by using an agency to access OBS, configure the **fs.obs.auth.node-cache-short-circuit.enable** parameter to determine whether to allow access to the ECS metadata API, thereby determining whether to trigger ECS flow control.

  The MRS cluster can access OBS using a temporary AK/SK obtained by an agency. The temporary AK/SK is obtained using the ECS metadata API. The ECS metadata API has a flow control threshold of 140 times within 5 minutes for a single node. After the flow control is triggered, the node is added to the blacklist and cannot call the metadata API again within 30 minutes. To prevent flow control, MRS provides the node-level cross-process cache service meta to cache temporary AK/SK.

  Application scenario: Yarn jobs such as Spark and Hadoop that access OBS using a temporary AK/SK obtained by an agency. This parameter is configured in the **core-site.xml** file on the client.

  The default value is **true**. The Yarn application process in the MRS cluster obtains the temporary AK/SK from the node-level cache service meta. If meta is abnormal, obtain the temporary AK/SK from the ECS metadata API.

  If you do not want to directly access the ECS metadata API when the meta is abnormal, set this parameter to **false** to prevent the node from being added to the blacklist due to flow control.

3.9.4.21 MRS 1.9.0.7 Patch Description

Basic Information

Table 3-80 Basic information

Patch Version	MRS 1.9.0.7
Release Date	2021-01-15

Resolved Issues	**List of resolved issues in MRS 1.9.0.7:** **MRS Manager** Solved the problem of queue information loss due to queue update during capacity expansion. **MRS big data components** Solved the problem that the Hive on Spark task is suspended because the block ID is displayed as garbled characters. Solve the problem that the Hadoop task fails to be executed due to insufficient space. Self-developed APIs are added to Hive. Solved the problem that the **map.xml** file cannot be read. Solved the problem that Yarn is unavailable due to ZooKeeper dirty data. The ResourceManager memory leakage problem of Yarn is resolved. The Hive Har feature is optimized. Upgraded the OBS packages. The JDK version is upgraded.
	List of resolved issues in MRS 1.9.0.6: **MRS Manager** MRS Manager supports scale-in of specified nodes in a yearly/monthly cluster. **MRS big data components** Solved the problem of slow response when HiveSE delivers SQL statements. Supported the jobhistory query failure information interface. Solved the problem that fine-grained permissions do not take effect. Fixed the exception that occurs when Hive on Spark reads data. Solved the problem that data volume increases when the Hive on mrs task is executed twice. Solved the problem that the performance of some strings is poor when vector-based vectorized query is enabled in Hive.

	List of resolved issues in MRS 1.9.0.5: **MRS Manager** Optimized the service restart process during configuration saving on MRS Manager. Rectified the periodic backup failure on MRS Manager. **MRS big data components** Private patch of Ranger Resolved the JVM Create GC thread failed issue of Yarn. Added the HiveServer2 task stacking alarm. Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s. Added the alarm indicating that HiveServer2 uncomment ZooKeeper. Added the alarm indicating that Yarn tasks failed and the number of killed tasks exceeds 5 within 20 minutes. Corrected time zone of Spark JobHistory. Optimized MetaStore restart mechanism. Resolved the HIVE-22771 open-source issue. Resolved the Hive beeline log printing errors. Corrected the number of active nodes displayed on the Yarn page. Resolved the slow response of RM page display, which is caused by large number of RM threads. Supported OBS monitoring. Upgraded the OBS packages. Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc. Resolved the issue that Hive occasionally reports a Kryo deserialization failure. Resolved the issue of Spark JobHistory memory leakage. Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.
	List of resolved issues in MRS 1.9.0.3: **MRS Manager** Upgraded ARM JDK on MRS Manager. Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager. **MRS big data components** Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.

	List of resolved issues in MRS 1.9.0.2: **MRS Manager** Resolved the mutual trust lost between some Core nodes in the cluster. Resolved the issue that instances fail to be added after the patch is installed. Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager. **MRS big data components** Upgraded the OBS packages.
	List of resolved issues in MRS 1.9.0.1: **MRS Manager** Resolved the issue that MRS Manager does not support rolling patch installation without restarting services. **MRS big data components** Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes. Resolved the issue that Kafka does not support open-source access. Resolved the SPARK-27637 open-source issue. Optimized the Hive rolling restart. Upgraded the OBS packages.
Compatibility with Other Patches	The MRS 1.9.0.7 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.

Impact of Patch Installation

- During the installation of the MRS 1.9.0.7 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.7 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.22 MRS 1.9.0.6 Patch Description

Basic Information

Table 3-81 Basic information

Patch Version	MRS 1.9.0.6
Release Date	2020-05-20
Resolved Issues	**List of resolved issues in MRS 1.9.0.6:** **MRS Manager** MRS Manager supports scale-in of specified nodes in a yearly/monthly cluster. **MRS big data components** Solved the problem of slow response when HiveSE delivers SQL statements. Supported the jobhistory query failure information interface. Solved the problem that fine-grained permissions do not take effect. Fixed the exception that occurs when Hive on Spark reads data. Solved the problem that data volume increases when the Hive on mrs task is executed twice. Solved the problem that the performance of some strings is poor when vector-based vectorized query is enabled in Hive.

	List of resolved issues in MRS 1.9.0.5: **MRS Manager** Optimized the service restart process during configuration saving on MRS Manager. Rectified the periodic backup failure on MRS Manager. **MRS big data components** Private patch of Ranger Resolved the JVM Create GC thread failed issue of Yarn. Added the HiveServer2 task stacking alarm. Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s. Added the alarm indicating that HiveServer2 uncomment ZooKeeper. Added the alarm indicating that Yarn tasks failed and the number of killed tasks exceeds 5 within 20 minutes. Corrected time zone of Spark JobHistory. Optimized MetaStore restart mechanism. Resolved the HIVE-22771 open-source issue. Resolved the Hive beeline log printing errors. Corrected the number of active nodes displayed on the Yarn page. Resolved the slow response of RM page display, which is caused by large number of RM threads. Supported OBS monitoring. Upgraded the OBS packages. Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc. Resolved the issue that Hive occasionally reports a Kryo deserialization failure. Resolved the issue of Spark JobHistory memory leakage. Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.
	List of resolved issues in MRS 1.9.0.3: **MRS Manager** Upgraded ARM JDK on MRS Manager. Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager. **MRS big data components** Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.

	List of resolved issues in MRS 1.9.0.2: **MRS Manager** Resolved the mutual trust lost between some Core nodes in the cluster. Resolved the issue that instances fail to be added after the patch is installed. Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager. **MRS big data components** Upgraded the OBS packages.
	List of resolved issues in MRS 1.9.0.1: **MRS Manager** Resolved the issue that MRS Manager does not support rolling patch installation without restarting services. **MRS big data components** Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes. Resolved the issue that Kafka does not support open-source access. Resolved the SPARK-27637 open-source issue. Optimized the Hive rolling restart. Upgraded the OBS packages.
Compatibility with Other Patches	The MRS 1.9.0.6 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.

Impact of Patch Installation

- During the installation of the MRS 1.9.0.6 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.6 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.23 MRS 1.9.0.5 Patch Description

Basic Information

Table 3-82 Basic information

Patch Version	MRS 1.9.0.5
Release Date	2020-03-21

Resolved Issues	**List of resolved issues in MRS 1.9.0.5:** **MRS Manager** Optimized the service restart process during configuration saving on MRS Manager. Rectified the periodic backup failure on MRS Manager. **MRS big data components** Private patch of Ranger Resolved the JVM Create GC thread failed issue of Yarn. Added the HiveServer2 task stacking alarm. Added the alarm indicating that the GC time of HiveServer HiveMetastore exceeds 5s. Added the alarm indicating that HiveServer2 uncomment ZooKeeper. Added the alarm indicating that Yarn tasks failed and the number of killed tasks exceeds 5 within 20 minutes. Corrected time zone of Spark JobHistory. Optimized MetaStore restart mechanism. Resolved the HIVE-22771 open-source issue. Resolved the Hive beeline log printing errors. Corrected the number of active nodes displayed on the Yarn page. Resolved the slow response of RM page display, which is caused by large number of RM threads. Supported OBS monitoring. Upgraded the OBS packages. Resolved the issue that some data is not inserted when 10 data records are concurrently inserted into hive-jdbc. Resolved the issue that Hive occasionally reports a Kryo deserialization failure. Resolved the issue of Spark JobHistory memory leakage. Resolved the issue that the application list cannot be displayed occasionally in Spark JobHistory.
	List of resolved issues in MRS 1.9.0.3: **MRS Manager** Upgraded ARM JDK on MRS Manager. Resolved the issue that the system disk is fully occupied by logs of the Core node on MRS Manager. **MRS big data components** Resolved the issue that the number of Ranger logs cannot be set, which may cause full disk occupation.

	List of resolved issues in MRS 1.9.0.2: **MRS Manager** Resolved the mutual trust lost between some Core nodes in the cluster. Resolved the issue that instances fail to be added after the patch is installed. Resolved the issue that the rolling restart timeout interval of HiveServer cannot be modified on MRS Manager. **MRS big data components** Upgraded the OBS packages.
	List of resolved issues in MRS 1.9.0.1: **MRS Manager** Resolved the issue that MRS Manager does not support rolling patch installation without restarting services. **MRS big data components** Resolved the issue that the OBS entrusted access frequency is not limited to 140 times within 5 minutes. Resolved the issue that Kafka does not support open-source access. Resolved the SPARK-27637 open-source issue. Optimized the Hive rolling restart. Upgraded the OBS packages.
Compatibility with Other Patches	The MRS 1.9.0.5 patch can resolve all problems that have been resolved by the MRS 1.9.0 patch.

Impact of Patch Installation

- During the installation of the MRS 1.9.0.5 patch, MRS Manager will be restarted, and the components such as Hadoop, Hive, Spark, Kafka, Ranger, and related dependent services will be restarted in rolling mode. During the restart of MRS Manager, services are temporarily unavailable but services are not interrupted during the rolling restart.
- After installing the MRS 1.9.0.5 patch, you need to download and install all clients again, including the original clients of Master nodes and the clients used by other nodes of VPC (that is, the clients that you set up).
  - For details about how to fully update the original client of the active Master node, see **Fully Updating the Original Client of the Active Master Node**.
  - For details about how to fully update the original client of the standby Master node, see **Fully Updating the Original Client of the Standby Master Node**.
  - For details about how to fully install the clients you set up, see **Installing a Client (Versions Earlier Than 3.x)**.

NOTE

- You are advised to back up the old clients before reinstalling the new ones.
- If you have modified client configurations based on the service scenario, modify them again after reinstalling the clients.

3.9.4.24 MRS 1.8.10.1 Patch Description

Basic Information

Table 3-83 Basic information

Patch Version	MRS 1.8.10.1
Release Date	2020-01-07
Resolved Issues	**MRS big data components** The health check and rolling restart logic of MRS Kafka have been optimized.
Compatibility with Other Patches	None

Impact of Patch Installation

During the installation of MRS 1.8.10.1 patches, MRS Manager and Kafka are restarted. During the restart, the services are temporarily unavailable.

3.10 Tenant Management

3.10.1 Before You Start

This section describes how to manage tenants on the MRS console.

Tenant management operations on the console apply only to clusters of versions earlier than MRS 3.x.

Tenant management operations on FusionInsight Manager apply to all versions. For MRS 3.x and later versions, see **Overview**. For versions earlier than MRS 3.x, see **Overview**.

3.10.2 Overview

Definition

An MRS cluster provides various resources and services for multiple organizations, departments, or applications to share. The cluster provides tenants as a logical entity to use these resources and services. A mode involving different tenants is called multi-tenant mode. Currently, only the analysis cluster supports tenant management.

Principles

The MRS cluster provides the multi-tenant function. It supports a layered tenant model and allows dynamic adding or deleting of tenants to isolate resources. It dynamically manages and configures tenants' computing and storage resources.

The computing resources indicate tenants' Yarn task queue resources. The task queue quota can be modified, and the task queue usage status and statistics can be viewed.

The storage resources can be stored on HDFS. You can add and delete the HDFS storage directories of tenants, and set the quotas of file quantity and the storage space of the directories.

Tenants can create and manage tenants in a cluster based on service requirements.

- Roles, computing resources, and storage resources are automatically created when tenants are created. By default, all permissions of the new computing resources and storage resources are allocated to a tenant's roles.
- Permissions to view the current tenant's resources, add a subtenant, and manage the subtenant's resources are granted to the tenant's roles by default.
- After you have modified the tenant's computing or storage resources, permissions of the tenant's roles are automatically updated.

MRS supports a maximum of 512 tenants. The default tenants created by the system include **default**. Tenants that are in the topmost layer with the default tenant are called level-1 tenants.

Resource Pools

Yarn task queues support only the label-based scheduling policy. This policy enables Yarn task queues to associate NodeManagers that have specific node labels. In this way, Yarn tasks run on specified nodes so that tasks are scheduled and certain hardware resources are utilized. For example, Yarn tasks requiring a large memory capacity can run on nodes with a large memory capacity by means of label association, preventing poor service performance.

In an MRS cluster, the tenant logically divides Yarn cluster nodes to combine multiple NodeManagers into a resource pool. Yarn task queues can be associated with specified resource pools by configuring queue capacity policies, ensuring efficient and independent resource utilization in the resource pools.

MRS supports a maximum of 50 resource pools. By default, the system contains a **default** resource pool.

3.10.3 Creating a Tenant

Scenario

You can create a tenant on MRS Manager to specify the resource usage.

Prerequisites

- A tenant name has been planned. The name must not be the same as that of a role or Yarn queue that exists in the current cluster.
- If a tenant requires storage resources, a storage directory has been planned based on service requirements, and the planned directory does not exist under the HDFS directory.
- The resources that can be allocated to the current tenant have been planned and the sum of the resource percentages of direct sub-tenants under the parent tenant at every level does not exceed 100%.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-68 Synchronizing IAM users

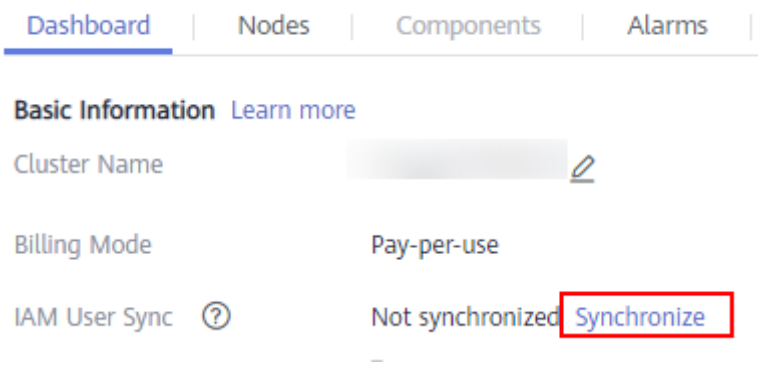

Procedure

Step 1 On the MRS cluster details page, click **Tenants**.

Figure 3-69 Tenants tab page

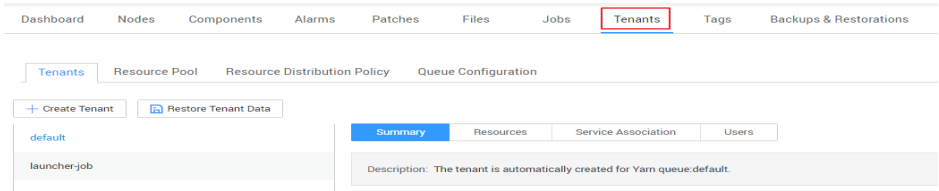

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 Click **Create Tenant**. On the page that is displayed, configure tenant properties. The following table takes MRS 3.x versions as an example.

Table 3-84 Tenant parameters

Parameter	Description
Name	Name of the current tenant. The value consists of 3 to 50 characters, and can contain letters, digits, and underscores (_).
Tenant Type	**Leaf** or **Non-leaf** tenant. If **Leaf** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. If **Non-leaf** is selected, sub-tenants can be added to the current tenant.
Compute Resource	Compute resources can be used by the tenant. The system automatically creates a task queue named after the tenant name in Yarn. If **Yarn** is not selected, the system does not automatically create a task queue.
Configuration Mode	If **Yarn** is selected for **Compute Resource**, this parameter can be set to **Basic** or **Advanced**. • **Basic**: Configure the percentage of compute resources used by the tenant in the default resource pool by specifying **Default Resource Pool Capacity (%)**. • **Advanced**: Configure the following parameters for advanced settings: – **Weight**: Tenant resource weight. The value ranges from 0 to 100. Tenant resource weight = Tenant weight/Total weight of tenants at the same level – **Minimum Resources**: resources preempted by the tenant. The value is a percentage or absolute value of the parent tenant's resources. When a tenant's workload is light, their resources are automatically lent to other tenants. When available resources are fewer than **Minimum Resources**, the tenant can preempt the resources that were lent out. – **Maximum Resources**: maximum resources that can be used by a tenant. The value is a percentage or absolute value of the parent tenant's resources. – **Reserved Resources**: resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is using resources of the tenant. The value is a percentage or absolute value of the parent tenant's resources.

Parameter	Description
Default Resource Pool Capacity (%)	Specifies the percentage of the computing resources used by the current tenant in the **default** resource pool. This parameter is required when **Configuration Mode** is **Basic**.
Storage Resource	Specifies storage resources for the current tenant. The system automatically creates a file folder named after the tenant name in the **/tenant** directory. When a tenant is created for the first time, the system automatically creates the **/tenant** directory in the HDFS root directory. If storage resources are not **HDFS**, the system does not create a storage directory under the root directory of HDFS.
Maximum Number of Files/Directories	Maximum number of files or directories that can be created in HDFS.
Storage Space Quota	Specifies the quota for HDFS storage space used by the current tenant. The value ranges from **1** to **8796093022208**. The unit is MB or GB. This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used. If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk. **NOTE** To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to **500** MB, the actual space for storing files is about 250 MB (500/2 = 250).
Storage Path	Specifies the tenant's HDFS storage directory. The system automatically creates a file folder named after the tenant name in the **/tenant** directory by default. For example, the default HDFS storage directory for **ta1** is **tenant/ta1**. When a tenant is created for the first time, the system automatically creates the **/tenant** directory in the HDFS root directory. The storage path is customizable.
Service	Specifies other service resources associated with the current tenant. HBase is supported. To configure this parameter, click **Associate Services**. In the dialog box that is displayed, set **Service** to **HBase**. If **Association Mode** is set to **Exclusive**, service resources are occupied exclusively. If **share** is selected, service resources are shared.
Description	Specifies the description of the current tenant.

Step 3 Click **OK** to save the settings.

It takes a few minutes to save the settings. If the **Tenant created successfully** is displayed in the upper-right corner, the tenant is added successfully. The tenant is created successfully.

NOTE

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. The role and its permissions are controlled by the system automatically and cannot be controlled manually under **Manage Role**.
- If you want to use the tenant, create a system user and assign the Manager_tenant role and the role corresponding to the tenant to the user. For details, see **Creating a User**.

----End

Related Tasks

View an added tenant.

Step 1 On the MRS details page, click **Tenants**.

Figure 3-70 Tenants tab page

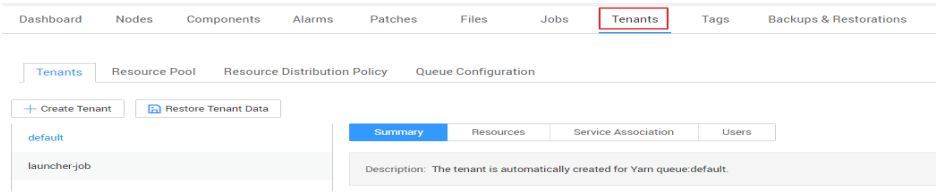

Step 2 In the tenant list on the left, click the name of the added tenant.

The **Summary** tab is displayed on the right by default.

Step 3 View **Basic Information**, **Resource Quota**, and **Charts** of the tenant.

If HDFS is in the **Stopped** state, **Available** and **Used** of **Space** in **Resource Quota** are **unknown.**

----End

3.10.4 Creating a Sub-tenant

Scenario

You can create a sub-tenant on MRS if the resources of the current tenant need to be further allocated.

Prerequisites

- A parent tenant has been added.
- A tenant name has been planned. The name must not be the same as that of a role or Yarn queue that exists in the current cluster.

- If a sub-tenant requires storage resources, a storage directory has been planned based on service requirements, and the planned directory does not exist under the storage directory of the parent tenant.
- The resources that can be allocated to the current tenant have been planned and the sum of the resource percentages of direct sub-tenants under the parent tenant at every level does not exceed 100%.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-71 Synchronizing IAM users

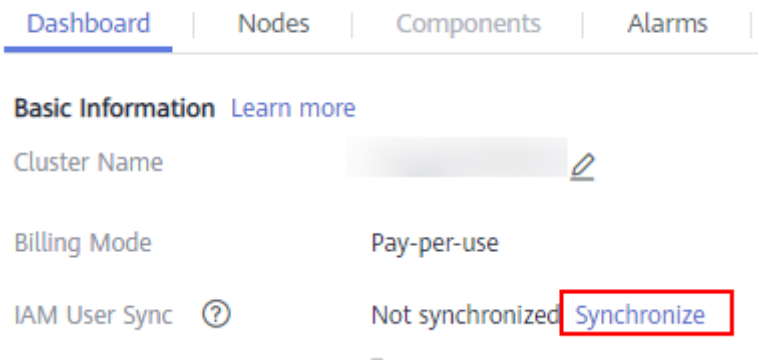

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-72 Tenants tab page

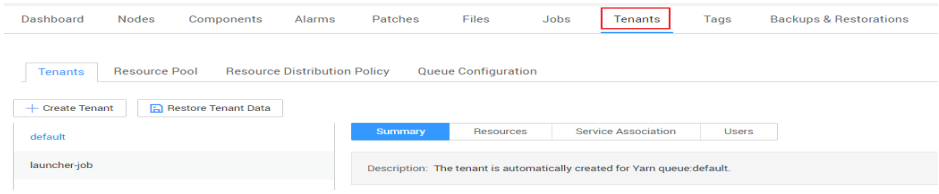

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 In the tenant list on the left, move the cursor to the tenant node to which a sub-tenant is to be added. Click **Create sub-tenant**. On the displayed page, configure the sub-tenant attributes according to the following table which takes MRS 3.x versions as an example.

Table 3-85 Sub-tenant parameters

Parameter	Description
Parent tenant	Specifies the name of the parent tenant.

Parameter	Description
Name	Specifies the name of the current tenant. The value consists of 3 to 20 characters, and can contain letters, digits, and underscores (_).
Tenant Type	The options include **Leaf** and **Non-leaf**. If **Leaf** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. If **Non-leaf** is selected, sub-tenants can be added to the current tenant.
Compute Resource	Specifies the dynamic computing resources for the current tenant. The system automatically creates a task queue named after the sub-tenant name in the Yarn parent queue. If **Yarn** is not selected, the system does not automatically create a task queue. If the parent tenant does not have compute resources, the sub-tenant cannot use compute resources.
Configuration Mode	If **Yarn** is selected for **Compute Resource**, this parameter can be set to **Basic** or **Advanced**. • **Basic**: Configure the percentage of compute resources used by the tenant in the default resource pool by specifying **Default Resource Pool Capacity (%)**. • **Advanced**: Configure the following parameters for advanced settings: – **Weight**: Tenant resource weight. The value ranges from 0 to 100. Tenant resource weight = Tenant weight/Total weight of tenants at the same level – **Minimum Resources**: resources preempted by the tenant. The value is a percentage or absolute value of the parent tenant's resources. When a tenant's workload is light, their resources are automatically lent to other tenants. When available resources are fewer than **Minimum Resources**, the tenant can preempt the resources that were lent out. – **Maximum Resources**: maximum resources that can be used by a tenant. The value is a percentage or absolute value of the parent tenant's resources. – **Reserved Resources**: resources reserved for the tenant. The value is a percentage or absolute value of the parent tenant's resources.
Default Resource Pool Capacity (%)	Specifies the percentage of the resources used by the current tenant. The base value is the total resources of the parent tenant. This parameter is required when **Configuration Mode** is **Basic**.

Parameter	Description
Storage Resource	Specifies storage resources for the current tenant. The system automatically creates a file in the HDFS parent tenant directory. The file is named the same as the name of the sub-tenant. If storage resources are not **HDFS**, the system does not create a storage directory under the root directory of HDFS. If the parent tenant does not have storage resources, the sub-tenant cannot use storage resources.
Maximum Number of Files/Directories	Maximum number of files or directories that can be created in HDFS.
Storage Space Quota	Specifies the quota for HDFS storage space used by the current tenant. The minimum value is **1**, and the maximum value is the total storage quota of the parent tenant. The unit is MB or GB. This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used. If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the HDFS physical disk. If the quota is greater than the quota of the parent tenant, the actual storage capacity is subject to the quota of the parent tenant. **NOTE** To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to **500** MB, the actual space for storing files is about 250 MB (500/2 = 250).
Storage Path	Specifies the tenant's HDFS storage directory. The system automatically creates a file folder named after the sub-tenant name in the directory of the parent tenant by default. For example, if the sub-tenant is **ta1s** and the parent directory is **tenant/ta1**, the system sets this parameter for the sub-tenant to **tenant/ta1/ta1s**. The storage path is customizable in the parent directory. The parent directory for the storage path must be the storage directory of the parent tenant.
Service	Specifies other service resources associated with the current tenant. HBase is supported. To configure this parameter, click **Associate Services**. In the dialog box that is displayed, set **Service** to **HBase**. If **Association Mode** is set to **Exclusive**, service resources are occupied exclusively. If **share** is selected, service resources are shared.
Description	Specifies the description of the current tenant.

Step 3 Click **OK** to save the settings.

It takes a few minutes to save the settings. If the **Tenant created successfully** is displayed in the upper-right corner, the tenant is added successfully. The tenant is created successfully.

NOTE

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. The role and its permissions are controlled by the system automatically and cannot be controlled manually under **Manage Role**.
- When using this tenant, create a system user and assign the user a related tenant role. For details, see **Creating a User**.

----End

3.10.5 Deleting a Tenant

Scenario

You can delete a tenant that is not required on MRS.

Prerequisites

- A tenant has been added.
- You have checked whether the tenant to be deleted has sub-tenants. If the tenant has sub-tenants, delete them; otherwise, you cannot delete the tenant.
- The role of the tenant to be deleted cannot be associated with any user or user group. For details about how to cancel the binding between a role and a user, see **Modifying User Information**.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-73 Synchronizing IAM users

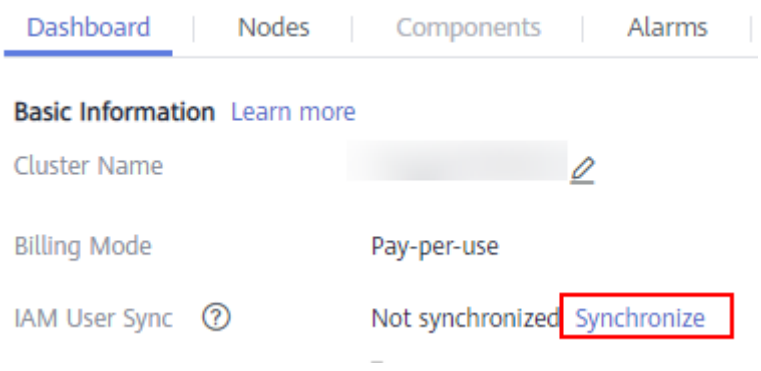

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-74 Tenants tab page

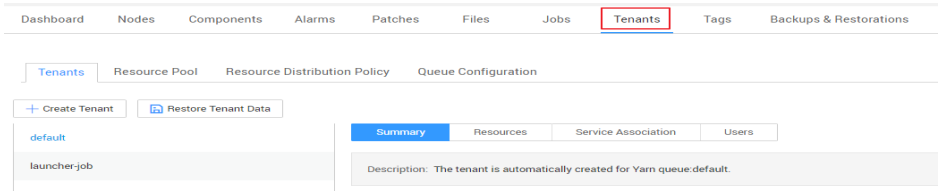

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 In the tenant list on the left, move the cursor to the tenant node to be deleted and click **Delete**.

The **Delete Tenant** dialog box is displayed. If you want to save the tenant data, select **Reserve the data of this tenant**. Otherwise, the tenant's storage space will be deleted.

Step 3 Click **OK**.

It takes a few minutes to save the configuration. After the tenant is deleted successfully, the role and storage space of the tenant are also deleted.

NOTE

- After the tenant is deleted, the task queue of the tenant still exists in Yarn.
- If you choose not to reserve data when deleting the parent tenant, data of sub-tenants is also deleted if the sub-tenants use storage resources.

----End

3.10.6 Managing a Tenant Directory

Scenario

You can manage the HDFS storage directory used by a specific tenant on MRS. The management operations include adding a tenant directory, modifying the directory file quota, modifying the storage space, and deleting a directory.

Prerequisites

- A tenant associated with HDFS storage resources has been added.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-75 Synchronizing IAM users

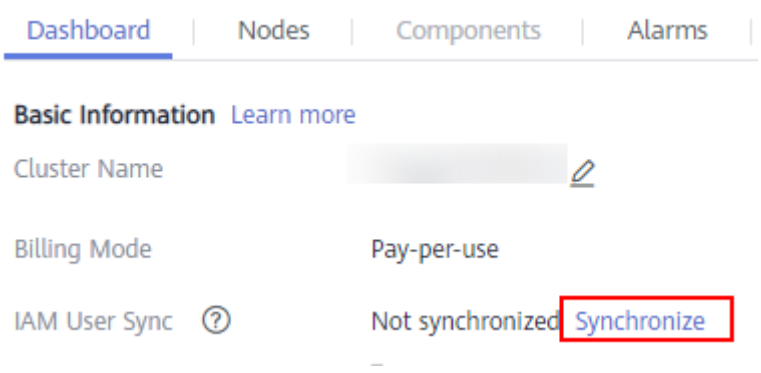

Procedure

- View a tenant directory.
  a. On the MRS details page, click **Tenants**.

  Figure 3-76 Tenants tab page

  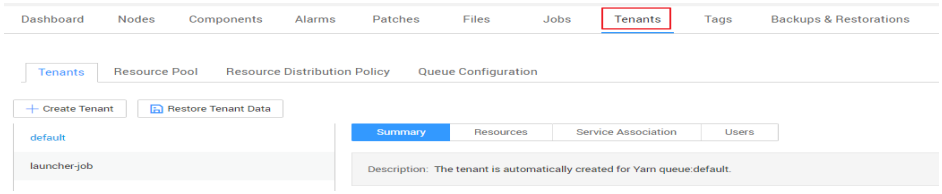

  NOTE

  For MRS 3.x or later, see **Overview**.

  b. In the tenant list on the left, click the target tenant.
  c. Click the **Resources** tab.
  d. View the **HDFS Storage** table.
     - The **Maximum Number of Files/Directories** column indicates the quotas for the file and directory quantity of the tenant directory.
     - The **Space Quota** column indicates storage space size of tenant directories.
- Add a tenant directory.
  a. On the MRS details page, click **Tenants**.

  NOTE

  For MRS 3.x or later, see **Overview**.

  b. In the tenant list on the left, click the tenant whose HDFS storage directory needs to be added.
  c. Click the **Resources** tab.
  d. In the **HDFS Storage** table, click **Create Directory**.
     - Set **Path** to a tenant directory path.

NOTE

- If the current tenant is not a sub-tenant, the new path is created in the HDFS root directory.
- If the current tenant is a sub-tenant, the new path is created in the specified directory.

A complete HDFS storage directory can contain a maximum of 1,023 characters. An HDFS directory name contains digits, letters, spaces, and underscores (_). The name cannot start or end with a space.

- Set **Maximum Number of Files/Directories** to the quotas of file and directory quantity.

  Maximum Number of Files/Directories is optional. Its value ranges from **1** to **9223372036854775806**.

- Set **Storage Space Quota** to the storage space size of the tenant directory.

  The value of **Storage Space Quota** ranges from **1** to **8796093022208**.

  NOTE

  To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value of **Storage Space Quota** is set to **500**, the actual space for storing files is about 250 MB (500/2 = 250).

e. Click **OK**. The system creates tenant directories in the HDFS root directory.

- Modify a tenant directory.

  a. On the MRS details page, click **Tenants**.

  NOTE

  For MRS 3.x or later, see **Overview**.

  b. In the tenant list on the left, click the tenant whose HDFS storage directory needs to be modified.

  c. Click the **Resources** tab.

  d. In the **HDFS Storage** table, click **Modify** in the **Operation** column of the specified tenant directory.

  - Set **Maximum Number of Files/Directories** to the quotas of file and directory quantity.

    Maximum Number of Files/Directories is optional. Its value ranges from **1** to **9223372036854775806**.

  - Set **Storage Space Quota** to the storage space size of the tenant directory.

    The value of **Storage Space Quota** ranges from **1** to **8796093022208**.

NOTE

To ensure data reliability, one backup is automatically generated for each file saved in HDFS, that is, two copies are generated in total. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value of **Storage Space Quota** is set to **500**, the actual space for storing files is about 250 MB (500/2 = 250).

e. Click **OK**.

- Delete a tenant directory.

  a. On the MRS details page, click **Tenants**.

  NOTE

  For MRS 3.x or later, see **Overview**.

  b. In the tenant list on the left, click the tenant whose HDFS storage directory needs to be deleted.

  c. Click the **Resources** tab.

  d. In the **HDFS Storage** table, click **Delete** in the **Operation** column of the specified tenant directory.

  The default HDFS storage directory set during tenant creation cannot be deleted. Only the newly added HDFS storage directory can be deleted.

  e. Click **OK**. The tenant directory is deleted.

3.10.7 Restoring Tenant Data

Scenario

Tenant data is stored on Manager and in cluster components by default. When components are restored from faults or reinstalled, some tenant configuration data may be abnormal. In this case, you can manually restore the tenant data.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-77 Synchronizing IAM users

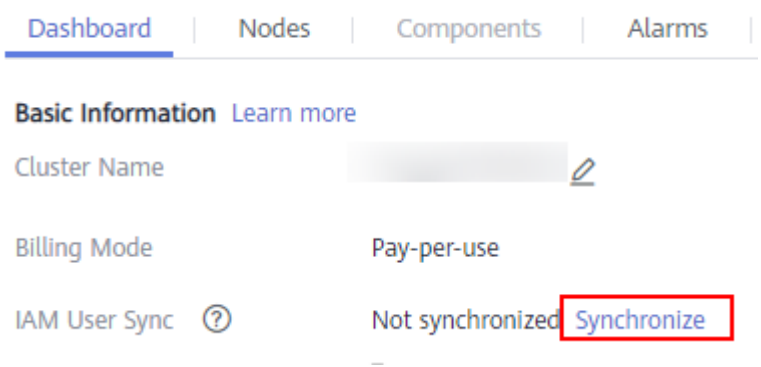

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-78 Tenants tab page

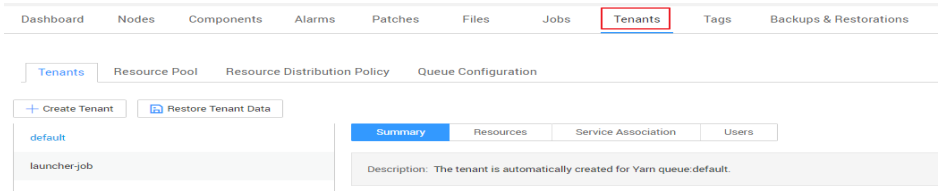

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 In the tenant list on the left, click a tenant node.

Step 3 Check the status of the tenant data.

1. In **Summary**, check the color of the circle on the left of **Basic Information**. Green indicates that the tenant is available and gray indicates that the tenant is unavailable.
2. Click **Resources** and check the circle on the left of **Yarn** or **HDFS Storage**. Green indicates that the resource is available, and gray indicates that the resource is unavailable.
3. Click **Service Association** and check the **Status** column of the associated service table. **Good** indicates that the component can provide services for the associated tenant. **Bad** indicates that the component cannot provide services for the tenant.
4. If any check result is abnormal, go to **Step 4** to restore tenant data.

Step 4 Click **Restore Tenant Data**.

Step 5 In the **Restore Tenant Data** window, select one or more components whose data needs to be restored. Click **OK**. The system automatically restores the tenant data.

----End

3.10.8 Creating a Resource Pool

Scenario

In an MRS cluster, users can logically divide Yarn cluster nodes to combine multiple NodeManagers into a Yarn resource pool. Each NodeManager belongs to one resource pool only. The system contains a **default** resource pool by default. All NodeManagers that are not added to customized resource pools belong to this resource pool.

You can create a customized resource pool on MRS and add hosts that have not been added to other customized resource pools to it.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-79 Synchronizing IAM users

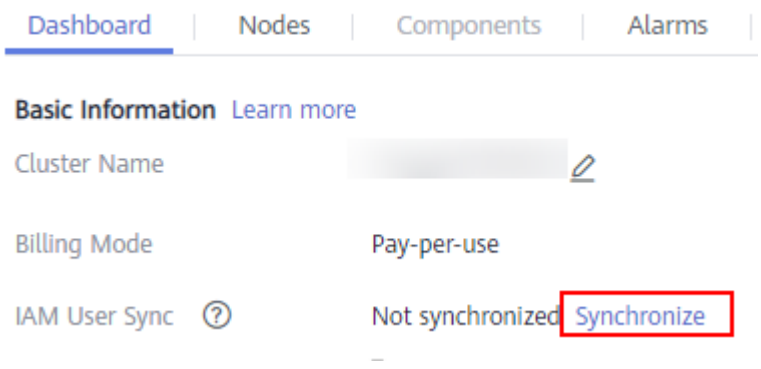

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-80 Tenants tab page

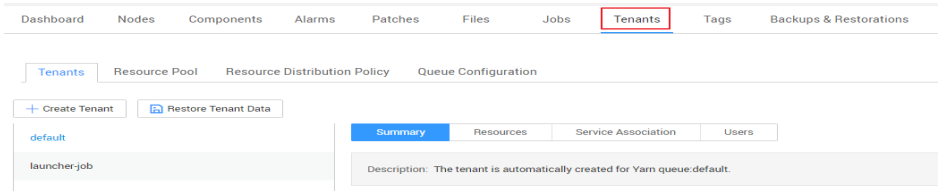

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 Click the **Resource Pools** tab.

Step 3 Click **Create Resource Pool**.

Step 4 In **Create Resource Pool**, set the properties of the resource pool.

- **Name**: Enter a name for the resource pool. The name of the newly created resource pool cannot be **default**.

  The name consists of 1 to 20 characters and can contain digits, letters, and underscores (_) but cannot start with an underscore (_).
- **Resource Label**: Enter the resource label of the resource pool. The value can contain 1 to 50 characters, including digits, letters, underscores (_), and hyphens (-), and must start with a digit or letter.
- **Available Hosts**: In the host list on the left, select a specified host name and add it to the resource pool. Only hosts in the cluster can be selected. The host list of a resource pool can be left blank.

Step 5 Click **OK**.

Step 6 After a resource pool is created, users can view the **Name**, **Members**, **Type**, **vCore** and **Memory** in the resource pool list. Hosts that are added to the customized resource pool are no longer members of the **default** resource pool.

----End

3.10.9 Modifying a Resource Pool

Scenario

You can modify members of an existing resource pool on MRS.

Prerequisites

You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-81 Synchronizing IAM users

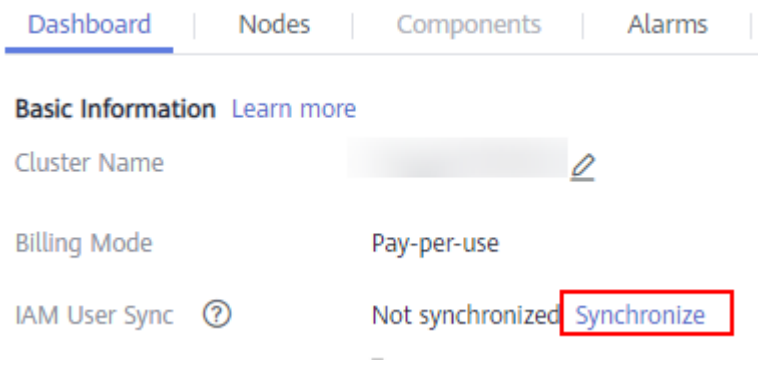

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-82 Tenants tab page

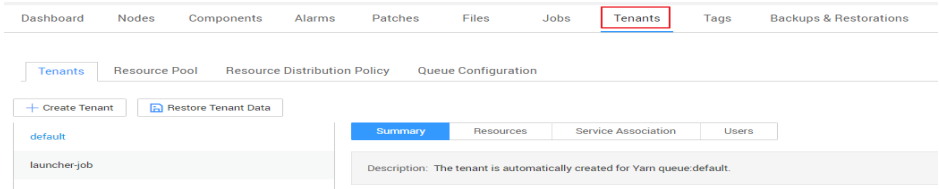

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 Click the **Resource Pools** tab.

Step 3 Locate the row that contains the specified resource pool, and click **Modify** in the **Operation** column.

Step 4 In **Modify Resource Pool**, modify **Added Hosts**.

- Adding a host: In the host list on the left, select the specified host name and add it to the resource pool.
- Deleting a host: In the host list on the right, click ✕ next to a host to remove the host from the resource pool. The host list of a resource pool can be left blank.

Step 5 Click **OK**.

----End

3.10.10 Deleting a Resource Pool

Scenario

You can delete an existing resource pool on MRS.

Prerequisites

- Any queue in a cluster cannot use the resource pool to be deleted as the default resource pool. Before deleting the resource pool, cancel the default resource pool. For details, see **Configuring a Queue**.
- Resource distribution policies of all queues have been cleared from the resource pool being deleted. For details, see **Clearing Configuration of a Queue**.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-83 Synchronizing IAM users

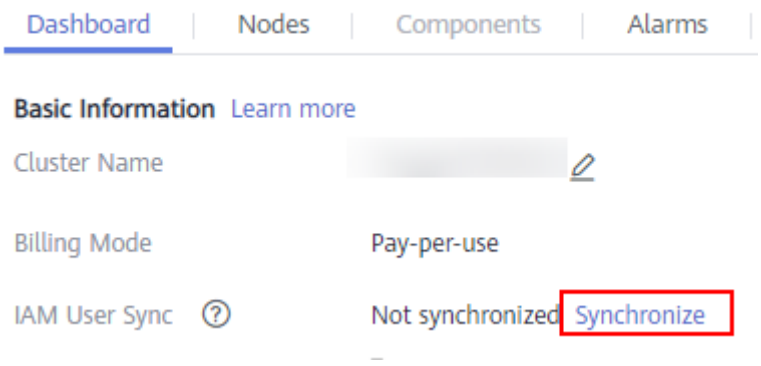

Procedure

Step 1 On the MRS details page, click **Tenant**.

Figure 3-84 Tenants tab page

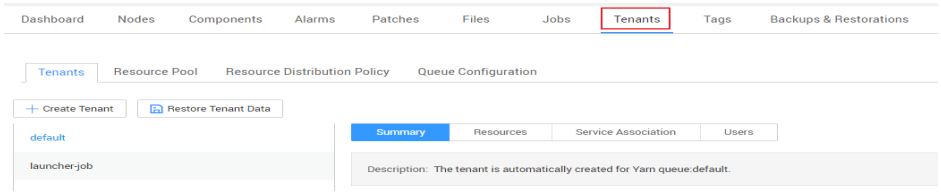

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 Click the **Resource Pools** tab.

Step 3 Locate the row that contains the specified resource pool, and click **Delete** in the **Operation** column.

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

----End

3.10.11 Configuring a Queue

Scenario

You can modify the queue configuration of a specified tenant on MRS based on service requirements.

Prerequisites

- A tenant associated with Yarn and allocated dynamic resources has been added.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-85 Synchronizing IAM users

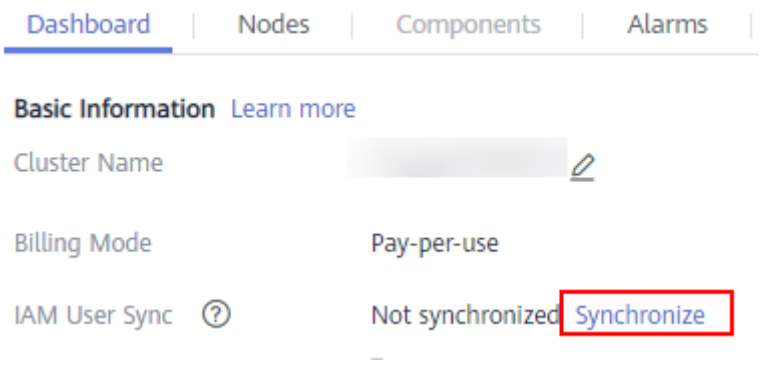

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-86 Tenants tab page

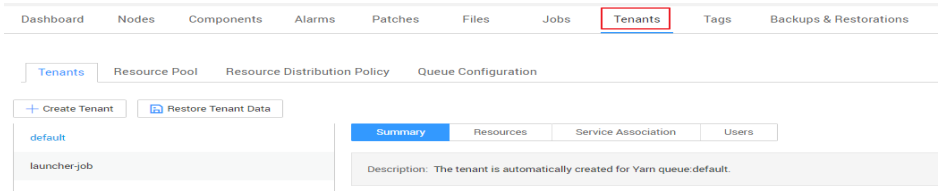

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 Click the **Queue Configuration** tab.

Step 3 In the tenant queue table, click **Modify** in the **Operation** column of the specified tenant queue.

NOTE

- In the tenant list on the left of the **Tenant Management** tab, click the target tenant. In the window that is displayed, choose **Resource**. On the page that is displayed, click to open the queue modification page.
- A queue can be bound to only one non-default resource pool.

Versions earlier than MRS 3.x:

Table 3-86 Queue configuration parameters

Parameter	Description
Maximum Applications	Specifies the maximum number of applications. The value ranges from 1 to 2147483647.
Maximum AM Resource Percent	Specifies the maximum percentage of resources that can be used to run the ApplicationMaster in a cluster. The value ranges from 0 to 1.
Minimum User Limit Percent (%)	Specifies the minimum percentage of resources consumed by a user. The value ranges from 0 to 100.
User Limit Factor	Specifies the limit factor of the maximum user resource usage. The maximum user resource usage percentage can be obtained by multiplying the limit factor with the percentage of the tenant's actual resource usage in the cluster. The minimum value is **0**.
Status	Specifies the current status of a resource plan. The values are **Running** and **Stopped**.
Default Resource Pool (Default Node Label Expression)	Specifies the resource pool used by a queue. The default value is **default**. If you want to change the resource pool, configure the queue capacity first. For details, see **Configuring the Queue Capacity Policy of a Resource Pool**.

MRS 3.x or later:

Table 3-87 Queue configuration parameters

Parameter	Description
Max Master Shares (%)	Indicates the maximum percentage of resources occupied by all ApplicationMasters in the current queue.
Max Allocated vCores	Indicates the maximum number of cores that can be allocated to a single YARN container in the current queue. The default value is **-1**, indicating that the number of cores is not limited within the value range.
Max Allocated Memory (MB)	Indicates the maximum memory that can be allocated to a single Yarn container in the current queue. The default value is **-1**, indicating that the memory is not limited within the value range.
Max Running Apps	Maximum number of tasks that can be executed at the same time in the current queue. The default value is **-1**, indicating that the number is not limited within the value range (the meaning is the same if the value is empty). The value 0 indicates that the task cannot be executed. The value ranges from -1 to 2147483647.
Max Running Apps per User	Maximum number of tasks that can be executed by each user in the current queue at the same time. The default value is **-1**, indicating that the number is not limited within the value range. If the value is **0**, the task cannot be executed. The value ranges from -1 to 2147483647.
Max Pending Apps	Maximum number of tasks that can be suspended at the same time in the current queue. The default value is **-1**, indicating that the number is not limited within the value range (the meaning is the same if the value is empty). The value **0** indicates that tasks cannot be suspended. The value ranges from -1 to 2147483647.
Resource Allocation Rule	Indicates the rule for allocating resources to different tasks of a user. The rule can be FIFO or FAIR. If a user submits multiple tasks in the current queue and the rule is FIFO, the tasks are executed one by one in sequential order. If the rule is FAIR, resources are evenly allocated to all tasks.
Default Resource Label	Indicates that tasks are executed on a node with a specified resource label. **NOTE** If you need to use a new resource pool, change the default label to the new resource pool label.

Parameter	Description
Active	• **ACTIVE**: indicates that the current queue can receive and execute tasks. • **INACTIVE**: indicates that the current queue can receive but cannot execute tasks. Tasks submitted to the queue are suspended.
Open	• **OPEN**: indicates that the current queue is opened. • **CLOSED**: indicates that the current queue is closed. Tasks submitted to the queue are rejected.

----End

3.10.12 Configuring the Queue Capacity Policy of a Resource Pool

Scenario

After a resource pool is added, the capacity policies of available resources need to be configured for Yarn task queues. This ensures that tasks in the resource pool are running properly. Each queue can be configured with the queue capacity policy of only one resource pool. Users can view the queues in any resource pool and configure queue capacity policies. After the queue policies are configured, Yarn task queues and resource pools are associated.

You can configure queue policies on MRS.

Prerequisites

- A resource pool has been added.
- The task queues are not associated with other resource pools. By default, all queues are associated with the **default** resource pool.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-87 Synchronizing IAM users

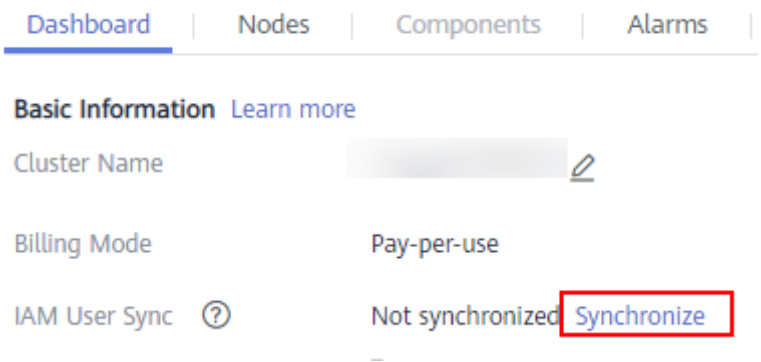

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-88 Tenants tab page

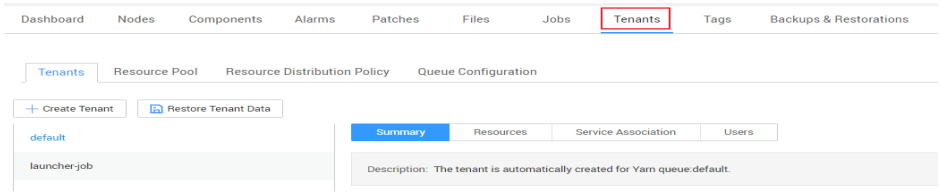

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 Click the **Resource Distribution Policies** tab.

Step 3 In **Resource Pools**, select a specified resource pool.

Available Resource Quota: indicates that all resources in each resource pool are available for queues by default.

Step 4 Locate the specified queue in the **Resource Allocation** table, and click **Modify** in the **Operation** column.

Step 5 In **Modify Resource Allocation**, configure the resource capacity policy of the task queue in the resource pool.

- **Capacity (%)**: specifies the percentage of the current tenant's computing resource usage.
- **Maximum Capacity (%)**: specifies the percentage of the current tenant's maximum computing resource usage.

Step 6 Click **OK** to save the settings.

----End

3.10.13 Clearing Configuration of a Queue

Scenario

Users can clear the configuration of a queue on MRS Manager when the queue does not need resources from a resource pool or if a resource pool needs to be disassociated from the queue. Clearing queue configurations means that the resource capacity policy of the queue is canceled.

Prerequisites

- If a queue is to be unbound from a resource pool, this resource pool cannot serve as the default resource pool of the queue. Therefore, you must first change the default resource pool of the queue to another one. For details, see **Configuring a Queue**.
- You have synchronized IAM users. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Figure 3-89 Synchronizing IAM users

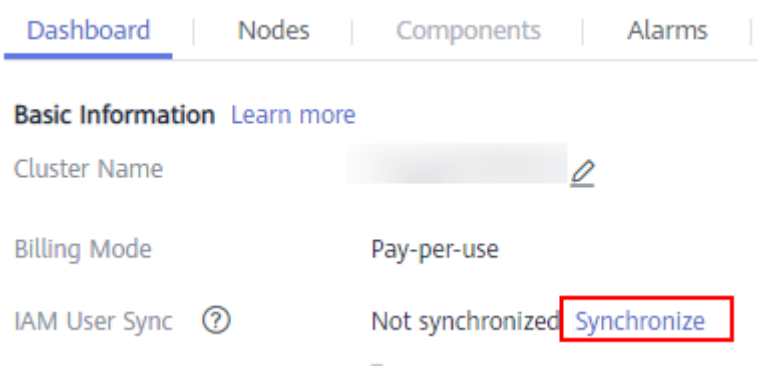

Procedure

Step 1 On the MRS details page, click **Tenants**.

Figure 3-90 Tenants tab page

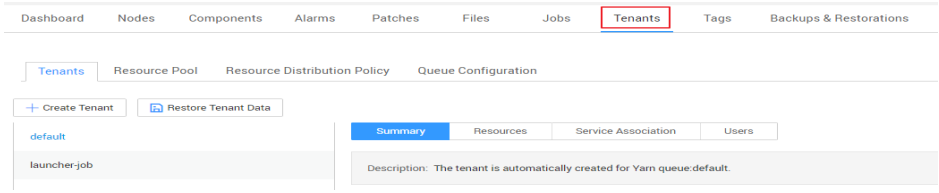

NOTE

For MRS 3.x or later, see **Overview**.

Step 2 Click the **Resource Distribution Policies** tab.

Step 3 In **Resource Pools**, select a specified resource pool.

Step 4 Locate the specified queue in the **Resource Allocation** table, and click **Clear** in the **Operation** column

In the **Clear Queue Configuration** dialog box, click **OK** to clear the queue configuration in the current resource pool.

NOTE

If no resource capacity policy is configured for a queue, the clearing function is unavailable for the queue by default.

----End

3.11 Bootstrap Actions

3.11.1 Introduction to Bootstrap Actions

You can run bootstrap actions to install additional third-party software, modify the cluster running environment, and perform other customizations. Bootstrap actions

can execute scripts on specified nodes before or after the first startup of cluster components. You can only manually run the third-party component installation script on the node to install a running cluster component.

If you choose to run bootstrap actions when scaling out a cluster, the bootstrap actions will be run on the newly added nodes in the same way. If auto scaling is enabled in a cluster, you can add an automation script in addition to configuring a resource plan. Then the automation script executes the corresponding script on the nodes that are scaled out or in to implement custom operations.

For versions earlier than MRS 3.*x*, your scripts are executed as user **root**. You can run the **su - XXX** command in a script to switch to another user.

For MRS 3.*x* or later, your scripts are executed as user **omm** by default. You can run the **su - XXX** command in a script to switch to another user.

NOTE

Versions earlier than MRS 3.*x*: Bootstrap action scripts must be executed as user **root**. Otherwise, your cluster may become unavailable.

MRS 3.*x* or later: Bootstrap action scripts must be executed as user **omm**. Otherwise, your cluster may become unavailable.

MRS determines the result based on the return code after the execution of the bootstrap action script. If the return code is **0**, the script is executed successfully. If the return code is not **0**, the execution fails. If a bootstrap action script fails to be executed on a node, the corresponding boot script will fail to be executed. In this case, you can set **Action upon Failure** to choose whether to continue to execute the subsequent scripts. Example 1: If you set **Action upon Failure** to **Continue** for all scripts during cluster creation, all the scripts will be executed regardless of whether they are successfully executed, and the startup process will be complete. Example 2: If a script fails to be executed and **Action upon Failure** is set to **Stop**, subsequent scripts will not be executed and cluster creation or scale-out will fail.

You can add a maximum of 18 bootstrap actions, which will be executed before or after the cluster component is started in the order you specified. The bootstrap actions performed before or after the component startup must be completed within 60 minutes. Otherwise, the cluster creation or scale-out will fail.

3.11.2 Preparing the Bootstrap Action Script

Currently, bootstrap actions support Linux shell scripts only. Script files must end with **.sh**.

Uploading the Installation Packages and Files to an OBS File System

Before compiling a script, you need to upload all required installation packages, configuration packages, and relevant files to the OBS file system in the same region. Because networks of different regions are isolated from each other, MRS VMs cannot download OBS files from other regions.

Compiling a Script for Downloading Files from the OBS File System

You can specify the file to be downloaded from OBS in the script. If you upload files to a private file system, you need to run the **hadoop fs** command to download the files. The following example shows that the **obs://yourbucket/**

myfile.tar.gz file will be downloaded to the local host and decompressed to the **/your-dir** directory.

```
#!/bin/bash
source /opt/Bigdata/client/bigdata_env;hadoop fs -D fs.obs.endpoint=<obs-endpoint> -D
fs.obs.access.key=<your-ak> -D fs.obs.secret.key=<your-sk> -copyToLocal obs://yourbucket/
myfile.tar.gz ./
mkdir -p /<your-dir>
tar -zxvf myfile.tar.gz -C /<your-dir>
```

NOTE

- The default client installation path for MRS 3.x or later is /opt/Bigdata/client, and for versions earlier than MRS 3.x is /opt/client. Configure the path based on site requirements.
- The Hadoop client has been preinstalled on the MRS node. You can run the **hadoop fs** command to download or upload data from or to OBS.
- Obtain the obs-endpoint of each region..
- Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.

Uploading the Script to the OBS File System

After script compilation, upload the script to the OBS file system in the same region. At the time you specify, each node in the cluster downloads the script from OBS and executes the script as user **root**.

3.11.3 View Execution Records

You can view the execution result of the bootstrap operation on the **Bootstrap Action** page.

Viewing the Execution Result

1. Log in to the MRS console.
2. In the left navigation pane, choose **Clusters** > **Active Clusters**. Click a cluster you want to query.

   The cluster details page is displayed.
3. On the cluster details page, click the **Bootstrap Action** tab. Information about the bootstrap actions added during cluster creation is displayed.

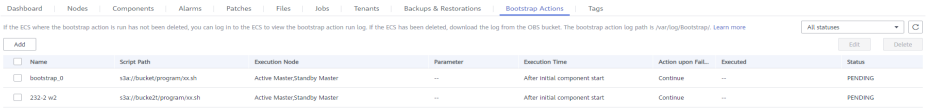

NOTE

- You select **Before initial component start** or **After initial component start** in the upper right corner to query information about the related bootstrap actions.
- The last execution result is listed here. For a newly created cluster, the records of bootstrap actions executed during cluster creation are listed. If a cluster is expanded, the records of bootstrap actions executed on the newly added nodes are listed.

Viewing Execution Logs

If you want to view the run logs of a bootstrap action, set **Action upon Failure** to **Continue** when adding the bootstrap action. And then, log in to each node to view the run logs in the **/var/log/Bootstrap** directory. If you add bootstrap actions before and after component start, you can distinguish bootstrap action logs of the two phases based on the timestamps.

You are advised to print logs in detail in the script so that you can view the detailed run result. MRS redirects the standard output and error output of the script to the log directory of the bootstrap action.

3.11.4 Adding a Bootstrap Action

Add a bootstrap action.

Procedure

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of your desired cluster.

Step 3 On the page that is displayed, click the **Bootstrap Actions** tab.

Step 4 Click **Add** and set parameters as prompted.

Figure 3-91 Add Bootstrap Action

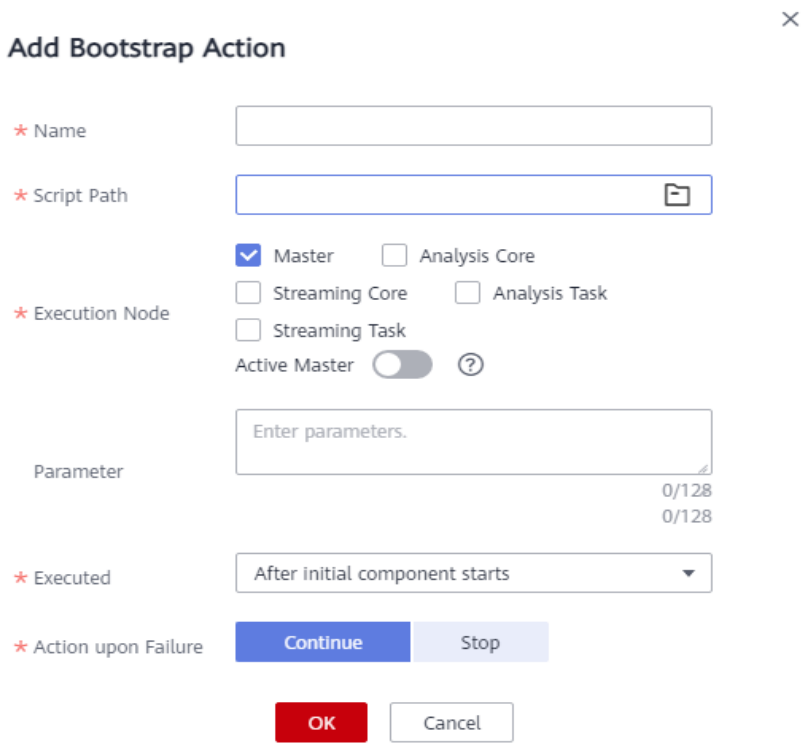

Table 3-88 Parameters

Parameter	Description
Name	Name of a bootstrap action script The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and must not start with a space. The value can contain 1 to 64 characters. **NOTE** A name must be unique in the same cluster. You can set the same name for different clusters.
Script Path	Script path. The value can be an OBS file system path or a local VM path. • An OBS file system path must start with **obs://** and end with **.sh**, for example, **obs://mrs-samples/***xxx***.sh**. • A local VM path must start with a slash (/) and end with **.sh**. **NOTE** A path must be unique in the same cluster, but can be the same for different clusters.
Parameter	Bootstrap action script parameters
Execution Node	Select a type of the node where the bootstrap action script is executed.
Executed	Select the time when the bootstrap action script is executed. • Before initial component start • After initial component start **NOTE** You can only manually run the third-party component installation script on the node to install a running cluster component.
Action upon Failure	Whether to continue to execute subsequent scripts and create a cluster after the script fails to be executed. **NOTE** You are advised to set this parameter to **Continue** in the debugging phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.
Run as root	Whether to escalate the permission to user **root** If the bootstrap action requires root user operations, enable this function, or the bootstrap action may fail to execute. **NOTE** This parameter is available for MRS 3.1.5 clusters.

Step 5 Click **OK** to save the configuration.

Step 6 Click **Yes**.

----End

3.11.5 Modifying a Bootstrap Action

Scenario

Modify an existing bootstrap action on an MRS cluster.

Procedure

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of your desired cluster.

Step 3 On the page that is displayed, click the **Bootstrap Actions** tab.

Step 4 In the list, select the item to be modified and click **Edit**.

Step 5 Modify the parameters as needed.

Step 6 Click **OK** to save the modification.

Step 7 Click **Yes**.

----End

3.11.6 Deleting a Bootstrap Action

Scenario

Delete a bootstrap action on an MRS cluster.

Procedure

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of your desired cluster.

Step 3 On the page that is displayed, click the **Bootstrap Actions** tab.

Step 4 In the list, select the item to be deleted and click **Delete**.

Step 5 Click **OK**.

----End

4 Using an MRS Client

4.1 Installing a Client

4.1.1 Installing a Client (MRS 3.*x* or Later)

Scenario

Install clients of all services (excluding Flume) in an MRS cluster. For details about how to install the Flume client, see **Installing the Flume Client**.

You can install the clients on a node in or outside the cluster.

After modifying the server configuration of a cluster component, reinstall the component client to ensure the server and the client both run the same version to provide services properly.

Prerequisites

- If the node where the client is to be installed is outside the cluster, the node must be able to communicate with the nodes in the cluster. Otherwise, client installation will fail.
- The node where the client is to be installed must have the NTP service enabled and synchronized time with the server. Otherwise, client installation will fail.
- You install the client as user **root** or any OS user. The user must have the operation permission on the client file storage directory and installation directory. The permission on the two directories is **755**.

  This section uses the OS user **user_client** as an example to describe how to install the client in the **/opt/hadoopclient** directory.
- When you install the client as a user other than **omm** or **root**, and the **/var/tmp/patch** directory already exists, you have changed the permission for the directory to **777** and changed the permission for the logs in the directory to **666**.

Installing a Client on a Node Inside a Cluster

Step 1 Obtain the client software package.

Log in to FusionInsight Manager by referring to **Accessing FusionInsight Manager (MRS 3.x or Later)**. On the **Cluster** > **Dashboard** page, click the more sign (...) and select **Download Client**. In the **Download Cluster Client** dialog box displayed, configure parameters and click **OK**.

NOTE

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

Table 4-1 Client download parameters

Parameter	Description	Example Value
Select Client Type	• **Complete Client**: contains the complete client software package and configuration files, which applies to non-development task scenarios. • **Configuration Files Only**: downloads only client configuration files in the scenario where the administrator modifies the server configuration on FusionInsight Manager after the complete client is downloaded and installed in an application development task, and developers need to update client configuration files.	Complete Client
Select Platform Type	**The client type must match the architecture of the node where the client is to be installed. Otherwise, the installation fails. For clusters of the LTS version, only the client software package whose type is the same as that of FusionInsight Manager can be downloaded.** • **x86_64**: indicates the client software package that can be deployed on a x86 platform. • **aarch64**: indicates the client software package that can be deployed on a TaiShan server.	x86_64

Parameter	Description	Example Value
Save to Path	The path for storing the client software package on the active OMS node • Select **Save to Path**: Customize the path for storing the client software package on the active OMS node. User **omm** must have the read, write, and execute permissions on the path. If the path is not changed, the client file generated is saved in the **/tmp/FusionInsight-Client** directory on the active OMS node in the cluster by default. • Not to select **Save to Path**: The generated client file is automatically downloaded and saved to the local host. Before installing the client, you need to upload the file to a specified directory on the target node.	Select **Save to Path**

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

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

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

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

Step 3 Log in to the target node as a user, for example, **user_client**.

NOTE

You can install the client as user **root** or any other OS user. The user must have the operation permission on the client file storage directory and installation directory. The permission on the two directories is **755**.

Step 4 Decompress the client software package.

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

cd /tmp/clienttemp

Run the following commands to decompress the package and obtain **FusionInsight_Cluster_1_Services_ClientConfig.tar**:

tar -xvf FusionInsight_Cluster_1_Services_Client.tar

Run the following command to decompress **FusionInsight_Cluster_1_Services_ClientConfig.tar**:

tar -xvf FusionInsight_Cluster_1_Services_ClientConfig.tar

Step 5 Go to the package directory and run the following command to install the client to a specified directory:

cd FusionInsight_Cluster_1_Services_ClientConfig

./install.sh *Client installation directory*

For example, run the **./install.sh /opt/hadoopclient** command to install the client and wait until the installation is complete.

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

NOTE

- **If there is no specified client installation directory, it will be automatically created. If you specify an existing directory, it must be empty, and the path cannot contain spaces. The value can contain only uppercase letters, lowercase letters, digits, and underscores (_).**
- You need to manually delete the client installation directory when uninstalling a client.
- To ensure that the installed client can be used only by the installation user, add the **-o** parameter during the installation. For example, run the **./install.sh /opt/hadoopclient -o** command to install the client.

Step 6 Use the client by referring to "Using the client of Each Component".

----End

Installing a Client on a Node Outside a Cluster

Step 1 Create an ECS that meets the following requirements:

- A Linux ECS has been prepared. For details about the supported OS of the ECS, see **Table 4-2**.

  Table 4-2 Reference list

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

  In addition, sufficient disk space is allocated for the ECS, for example, 40 GB.
- The ECS and the MRS cluster are in the same VPC.

- The security group of the ECS must be the same as that of the master node in the MRS cluster.
- The NTP service has been installed on the ECS OS and is running properly.

  If the NTP service is not installed, run the **yum install ntp -y** command to install it when the **yum** source is configured.
- A user can log in to the Linux ECS using the password (in SSH mode).
- All ports in the inbound direction of the MRS cluster security group are open to the client node. For details, see **Adding a Security Group Rule**.

Step 2 Perform NTP time synchronization to synchronize the time of nodes outside the cluster with the time of the MRS cluster.

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 4-1 Adding the master node IP addresses

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

   driftfile /var/lib/ntp/drift

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

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

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

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

   # Enable public key cryptography.
   #crypto
   ```

2. Run the **service ntpd stop** command to stop the NTP service.
3. Run the following command to manually synchronize the time:

   /usr/sbin/ntpdate *192.168.10.8*

   NOTE

   192.168.10.8 indicates the IP address of the active Master node.

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.

Step 3 Obtain the client software package.

Log in to FusionInsight Manager by referring to **Accessing FusionInsight Manager (MRS 3.x or Later)**. On the **Cluster** > **Dashboard** page, click the more sign (...) and select **Download Client**. In the **Download Cluster Client** dialog box displayed, configure parameters and click **OK**.

NOTE

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

Table 4-3 Client download parameters

Parameter	Description	Example Value
Select Client Type	• **Complete Client**: contains the complete client software package and configuration files, which applies to non-development task scenarios. • **Configuration Files Only**: downloads only client configuration files in the scenario where the administrator modifies the server configuration on FusionInsight Manager after the complete client is downloaded and installed in an application development task, and developers need to update client configuration files.	Complete Client
Select Platform Type	**The client type must match the architecture of the node where the client is to be installed. Otherwise, the installation fails. For clusters of the LTS version, only the client software package whose type is the same as that of FusionInsight Manager can be downloaded.** • **x86_64**: indicates the client software package that can be deployed on a x86 platform. • **aarch64**: indicates the client software package that can be deployed on a TaiShan server.	x86_64

Parameter	Description	Example Value
Save to Path	The path for storing the client software package on the active OMS node • Select **Save to Path**: Customize the path for storing the client software package on the active OMS node. User **omm** must have the read, write, and execute permissions on the path. If the path is not changed, the client file generated is saved in the **/tmp/FusionInsight-Client** directory on the active OMS node in the cluster by default. • Not to select **Save to Path**: The generated client file is automatically downloaded and saved to the local host. Before installing the client, you need to upload the file to a specified directory on the target node.	Select **Save to Path**

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

The generated client software package is stored on the active OMS node of the cluster by default. You need to log in to the active OMS node as user **omm** and run the following command to copy the software package to a specified ECS:

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

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

Step 5 Log in to the target node as a user, for example, **user_client**.

NOTE

You can install the client as user **root** or any other OS user. The user must have the operation permission on the client file storage directory and installation directory. The permission on the two directories is **755**.

Step 6 Decompress the client software package.

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

cd /tmp/clienttemp

Run the following commands to decompress the package and obtain **FusionInsight_Cluster_1_Services_ClientConfig.tar**:

tar -xvf FusionInsight_Cluster_1_Services_Client.tar

Run the following command to decompress **FusionInsight_Cluster_1_Services_ClientConfig.tar**:

tar -xvf FusionInsight_Cluster_1_Services_ClientConfig.tar

Step 7 Check the network connection of the client.

1. Ensure that the host where the client is installed can communicate with the hosts listed in the **hosts** file stored in the directory containing the decompressed package, for example, **/tmp/FusionInsight_Cluster_1_Services_ClientConfig/hosts**.
2. If the host where the client is installed is not a host in the cluster, you need to set the mapping between the host name and the service plane IP address for each cluster node in **/etc/hosts**, as user **root**. Each host name uniquely maps an IP address. You can perform the following steps to import the domain name mapping of the cluster to the **hosts** file:
   a. Switch to user **root** or a user who has the permission to modify the **hosts** file.

      su - root

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

      cd /tmp/clienttemp/FusionInsight_Cluster_1_Services_ClientConfig

   c. Run the **cat realm.ini >> /etc/hosts** command to import the domain name mapping to the **hosts** file.

NOTE

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

Step 8 Log in to the node where the client is to be installed as user **user_client**, go to the client software package directory, and run the following command to install the client to a specified directory:

cd /tmp/clienttemp/FusionInsight_Cluster_1_Services_ClientConfig

./install.sh *Client installation directory*

For example, run the **./install.sh /opt/hadoopclient** command to install the client and wait until the installation is complete.

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

NOTE

- **If there is no specified client installation directory, it will be automatically created. If you specify an existing directory, it must be empty, and the path cannot contain spaces. The value can contain only uppercase letters, lowercase letters, digits, and underscores (_).**
- You need to manually delete the client installation directory when uninstalling a client.
- To ensure that the installed client can be used only by the installation user, add the **-o** parameter during the installation. For example, run the **./install.sh /opt/hadoopclient -o** command to install the client.

Step 9 Use the client by referring to "Using the client of Each Component".

----End

4.1.2 Installing a Client (Versions Earlier Than 3.x)

Scenario

An MRS client is required. The MRS cluster client can be installed on the Master or Core node in the cluster or on a node outside the cluster.

After a cluster of versions earlier than MRS 3.x is created, a client is installed on the active Master node by default. You can directly use the client. The installation directory is **/opt/client**.

For details about how to install a client of MRS 3.x or later, see **Installing a Client (MRS 3.x or Later)**.

NOTE

If a client has been installed on the node outside the MRS cluster and the client only needs to be updated, update the client using the user who installed the client, for example, user **root**.

Installing a Client on the Core Node

1. Log in to MRS Manager and choose **Services** > **Download Client** to download the client installation package to the active management node.

   NOTE

   If only the client configuration file needs to be updated, see method 2 in **Updating a Client (Versions Earlier Than 3.x)**.

2. Use the IP address to search for the active management node, and log in to the active management node using VNC.
3. Log in to the active management node, and run the following command to switch the user:

   sudo su - omm

4. On the MRS management console, view the IP address on the **Nodes** tab page of the specified cluster.

   Record the IP address of the Core node where the client is to be used.

5. On the active management node, run the following command to copy the client installation package to the Core node:

   scp -p /tmp/MRS-client/MRS_Services_Client.tar *IP address of the Core node*:**/opt/client**

6. Log in to the Core node as user **root**.

   Master nodes support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation.

7. Run the following commands to install the client:

   cd /opt/client

   tar -xvf MRS_Services_Client.tar

   tar -xvf MRS_Services_ClientConfig.tar

   cd /opt/client/MRS_Services_ClientConfig

   ./install.sh *Client installation directory*

For example, run the following command:

./install.sh /opt/client

8. For details about how to use the client, see **Using an MRS Client**.

Using an MRS Client

1. On the node where the client is installed, run the **sudo su - omm** command to switch the user. Run the following command to go to the client directory:

   cd /opt/client

2. Run the following command to configure environment variables:

   source bigdata_env

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

   kinit *MRS cluster user*

   Example: **kinit admin**

   NOTE

   User **admin** is created by default for MRS clusters with Kerberos authentication enabled and is used for administrators to maintain the clusters.

4. Run the client command of a component directly.

   For example, run the **hdfs dfs -ls /** command to view files in the HDFS root directory.

Installing a Client on a Node Outside the Cluster

Step 1 Create an ECS that meets the following requirements:

- For clusters of versions earlier than MRS 3.*x*, you need to confirm the CPU architecture of the current MRS cluster nodes. For clusters of versions earlier than MRS 3.*x*, the CPU architecture of the ECS must be the same as that of the MRS cluster node. For MRS 3.*x* or later, the MRS client is compatible with both of the following CPU architectures.
- An ECS has been prepared. For details about the OS and its version of the ECS, see **Table 4-4**.

Table 4-4 Reference list

CPU Architecture	OS	Supported Version
x86 computing	EulerOS	– Available: EulerOS 2.2 – Available: EulerOS 2.3 – Available: EulerOS 2.5

For example, a user can select an ECS running the EulerOS.

In addition, sufficient disk space is allocated for the ECS, for example, 40 GB.

- The ECS and the MRS cluster are in the same VPC.
- The security group of the ECS is the same as that of the master node in the MRS cluster.

  If this requirement is not met, modify the ECS security group or configure the inbound and outbound rules of the ECS security group to allow the ECS security group to be accessed by all security groups of MRS cluster nodes.
- To enable users to log in to a Linux ECS using a password (SSH), see *"Instances" > "Logging In to a Linux ECS" > "Login Using an SSH Password" in Elastic Cloud Server User Guide*.
- All ports in the inbound direction of the MRS cluster security group are open to the client node. For details, see **Adding a Security Group Rule**.

Step 2 Log in to MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**. Then, choose **Services**.

Step 3 Click **Download Client**.

Step 4 In **Client Type**, select **All client files**.

Step 5 In **Download To**, select **Remote host**.

Step 6 Set **Host IP Address** to the IP address of the ECS, **Host Port** to **22**, and **Save Path** to **/tmp**.

- If the default port **22** for logging in to an ECS using SSH has been changed, set **Host Port** to the new port.
- **Save Path** contains a maximum of 256 characters.

Step 7 Set **Login User** to **root**.

If other users are used, ensure that the users have read, write, and execute permission on the save path.

Step 8 Select **Password** or **SSH Private Key** for **Login Mode**.

- **Password**: Enter the password of user **root** set during cluster creation.
- **SSH Private Key**: Select and upload the key file used for creating the cluster.

Step 9 Click **OK** to generate a client file.

If the following information is displayed, the client package is saved. Click **Close**. Obtain the client file from the save path on the remote host that is set when the client is downloaded.

```
Client files downloaded to the remote host successfully.
```

If the following information is displayed, check the username, password, and security group configurations of the remote host. Ensure that the username and password are correct and an inbound rule of the SSH (22) port has been added to the security group of the remote host. And then, go to **Step 2** to download the client again.

```
Failed to connect to the server. Please check the network connection or parameter settings.
```

NOTE

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

Step 10 Log in to the ECS using VNC. For details, see **Instance** > **Logging In to a Linux** > **Logging In to a Linux** in the *Elastic Cloud Server User Guide*

All images support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation. It is recommended that you change the password upon the first login.

Step 11 Perform NTP time synchronization to synchronize the time of nodes outside the cluster with the time of the MRS cluster.

1. Check whether the NTP service is installed. If it is not installed, run the **yum install ntp -y** command to install it.
2. Run the **vim /etc/ntp.conf** command to edit the NTP client configuration file, add the IP address of the Master node in the MRS cluster, and comment out the IP addresses of other servers.

   server *master1_ip* prefer
   server *master2_ip*

   Figure 4-2 Adding the master node IP addresses

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

   driftfile /var/lib/ntp/drift

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

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

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

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

   # Enable public key cryptography.
   #crypto
   ```

3. Run the **service ntpd stop** command to stop the NTP service.
4. Run the following command to manually synchronize the time:

   /usr/sbin/ntpdate *192.168.10.8*

   NOTE

   192.168.10.8 indicates the IP address of the active Master node.

5. Run the **service ntpd start** or **systemctl restart ntpd** command to start the NTP service.

6. Run the **ntpstat** command to check the time synchronization result:

Step 12 On the ECS, switch to user **root** and copy the installation package in **Save Path** in **Step 6** to the **/opt** directory. For example, if **Save Path** is set to **/tmp**, run the following commands:

sudo su - root

cp /tmp/MRS_Services_Client.tar /opt

Step 13 Run the following command in the **/opt** directory to decompress the package and obtain the verification file and the configuration package of the client:

tar -xvf MRS_Services_Client.tar

Step 14 Run the following command to verify the configuration file package of the client:

sha256sum -c MRS_Services_ClientConfig.tar.sha256

The command output is as follows:

```
MRS_Services_ClientConfig.tar: OK
```

Step 15 Run the following command to decompress **MRS_Services_ClientConfig.tar**:

tar -xvf MRS_Services_ClientConfig.tar

Step 16 Run the following command to install the client to a new directory, for example, **/opt/Bigdata/client**. A directory is automatically generated during the client installation.

sh /opt/MRS_Services_ClientConfig/install.sh /opt/Bigdata/client

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

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

Step 17 Check whether the IP address of the ECS node is connected to the IP address of the cluster Master node.

For example, run the following command: **ping** *Master node IP address*.

- If yes, go to **Step 18**.
- If no, check whether the VPC and security group are correct and whether the ECS and the MRS cluster are in the same VPC and security group, and go to **Step 18**.

Step 18 Run the following command to configure environment variables:

source /opt/Bigdata/client/bigdata_env

Step 19 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step.

kinit *MRS cluster user*

Example: **kinit admin**

Step 20 Run the client command of a component.

For example, run the following command to query the HDFS directory:

hdfs dfs -ls /

----End

4.2 Updating a Client

4.2.1 Updating a Client (Version 3.x or Later)

A cluster provides a client for you to connect to a server, view task results, or manage data. If you modify service configuration parameters on Manager and restart the service, you need to download and install the client again or use the configuration file to update the client.

Updating the Client Configuration

Method 1:

Step 1 Log in to FusionInsight Manager. For details, see **Accessing FusionInsight Manager (MRS 3.x or Later)**. Click the name of the cluster to be operated in the **Cluster** drop-down list.

Step 2 Choose **More** > **Download Client** > **Configuration Files Only**.

The generated compressed file contains the configuration files of all services.

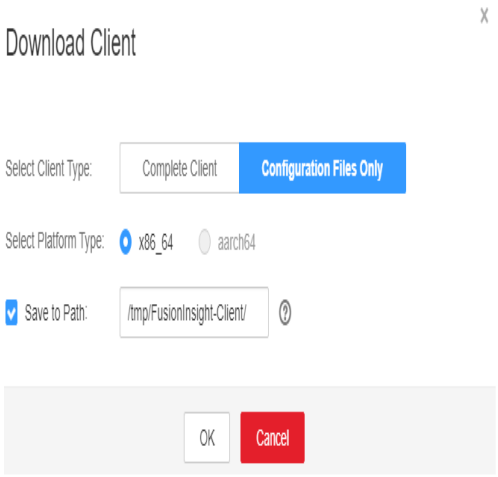

Step 3 Determine whether to generate a configuration file on the cluster node.

- If yes, select **Save to Path**, and click **OK** to generate the client file. By default, the client file is generated in **/tmp/FusionInsight-Client** on the active management node. You can also store the client file in other directories, and user **omm** has the read, write, and execute permissions on the directories. Then go to **Step 4**.
- If no, click **OK**, specify a local save path, and download the complete client. Wait until the download is complete and go to **Step 4**.

Step 4 Use WinSCP to save the compressed file to the client installation directory, for example, **/opt/hadoopclient**, as the client installation user.

Step 5 Decompress the software package.

Run the following commands to go to the directory where the client is installed, and decompress the file to a local directory. For example, the downloaded client file is **FusionInsight_Cluster_1_Services_Client.tar**.

cd /opt/hadoopclient

tar -xvf FusionInsight_Cluster_1_Services_Client.tar

Step 6 Verify the software package.

Run the following command to verify the decompressed file and check whether the command output is consistent with the information in the **sha256** file.

sha256sum -c
FusionInsight_Cluster_1_Services_ClientConfig_ConfigFiles.tar.sha256

```
FusionInsight_Cluster_1_Services_ClientConfig_ConfigFiles.tar: OK
```

Step 7 Decompress the package to obtain the configuration file.

tar -xvf FusionInsight_Cluster_1_Services_ClientConfig_ConfigFiles.tar

Step 8 Run the following command in the client installation directory to update the client using the configuration file:

sh refreshConfig.sh *Client installation directory Directory where the configuration file is located*

For example, run the following command:

sh refreshConfig.sh /opt/hadoopclient /opt/hadoopclient/
FusionInsight_Cluster_1_Services_ClientConfig_ConfigFiles

If the following information is displayed, the configurations have been updated successfully.

```
Succeed to refresh components client config.
```

----End

Method 2:

Step 1 Log in to the client installation node as user **root**.

Step 2 Go to the client installation directory, for example, **/opt/hadoopclient** and run the following commands to update the configuration file:

cd /opt/hadoopclient

sh autoRefreshConfig.sh

Step 3 Enter the username and password of the FusionInsight Manager administrator and the floating IP address of OMS.

NOTE

To obtain the floating IP address of OMS, log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of OMS. Record the value of **inet**. If the floating IP address of OMS cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the IP address on the Master node.

Step 4 Enter the names of the components whose configuration needs to be updated. Use commas (,) to separate the component names. Press **Enter** to update the configurations of all components if necessary.

If the following information is displayed, the configurations have been updated successfully.

```
Succeed to refresh components client config.
```

----End

4.2.2 Updating a Client (Versions Earlier Than 3.x)

NOTE

This section applies to clusters of versions earlier than MRS 3.x. For MRS 3.x or later, see **Updating a Client (Version 3.x or Later)**.

Updating a Client Configuration File

Scenario

An MRS cluster provides a client for you to connect to a server, view task results, or manage data. Before using an MRS client, you need to download and update the client configuration file if service configuration parameters are modified and a service is restarted on MRS Manager.

During cluster creation, the original client is stored in the **/opt/client** directory on all nodes in the cluster by default. After the cluster is created, only the client of a Master node can be directly used. To use the client of a Core node, you need to update the client configuration file first.

Procedure

Method 1: applicable to all versions

Step 1 Log in to MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**. Then, choose **Services**.

Step 2 Click **Download Client**.

Set **Client Type** to **Only configuration files**, **Download To** to **Server**, and click **OK** to generate the client configuration file. The generated file is saved in the **/tmp/MRS-client** directory on the active management node by default. You can customize the file path.

Figure 4-3 Downloading the client configuration file

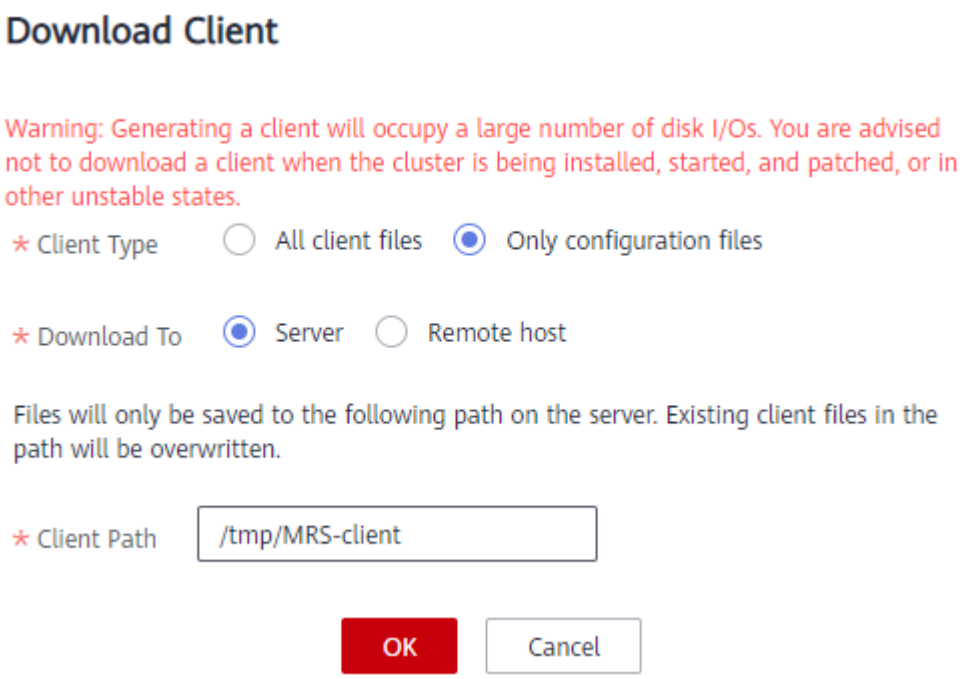

Step 3 Query and log in to the active Master node.

Step 4 If you use the client in the cluster, run the following command to switch to user **omm**. If you use the client outside the cluster, switch to user **root**.

sudo su - omm

Step 5 Run the following command to go to the client directory:

cd *Client installation directory*

Step 6 Run the following command to update client configurations:

sh refreshConfig.sh *Client installation directory Full path of the client configuration file package*

For example, run the following command:

sh refreshConfig.sh /opt/Bigdata/client /tmp/MRS-client/MRS_Services_Client.tar

If the following information is displayed, the configurations have been updated successfully.

```
ReFresh components client config is complete.
Succeed to refresh components client config.
```

----End

Method 2:

Step 1 After the cluster is installed, run the following command to switch to user **omm**. If you use the client outside the cluster, switch to user **root**.

sudo su - omm

Step 2 Run the following command to go to the client directory:

cd *Client installation directory*

Step 3 Run the following command and enter the name of an MRS Manager user with the download permission and its password (for example, the username is **admin** and the password is the one set during cluster creation) as prompted to update client configurations.

sh autoRefreshConfig.sh

Step 4 After the command is executed, the following information is displayed, where *XXX* indicates the name of the component installed in the cluster. To update client configurations of all components, press **Enter**. To update client configurations of some components, enter the component names and separate them with commas (,).

```
Components "xxx" have been installed in the cluster. Please input the comma-separated names of the components for which you want to update client configurations. If you press Enter without inputting any component name, the client configurations of all components will be updated:
```

If the following information is displayed, the configurations have been updated successfully.

```
Succeed to refresh components client config.
```

If the following information is displayed, the username or password is incorrect.

```
login manager failed,Incorrect username or password.
```

NOTE

- This script automatically connects to the cluster and invokes the **refreshConfig.sh** script to download and update the client configuration file.
- By default, the client uses the floating IP address specified by **wsom=xxx** in the **Version** file in the installation directory to update the client configurations. To update the configuration file of another cluster, modify the value of **wsom=xxx** in the **Version** file to the floating IP address of the corresponding cluster before performing this step.

----End

Fully Updating the Original Client of the Active Master Node

Scenario

During cluster creation, the original client is stored in the **/opt/client** directory on all nodes in the cluster by default. The following uses **/opt/Bigdata/client** as an example.

- For a normal MRS cluster, you will use the pre-installed client on a Master node to submit a job on the management console page.
- You can also use the pre-installed client on the Master node to connect to a server, view task results, and manage data.

After installing the patch on the cluster, you need to update the client on the Master node to ensure that the functions of the built-in client are available.

Procedure

Step 1 Log in to MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**. Then, choose **Services**.

Step 2 Click **Download Client**.

Set **Client Type** to **All client files**, **Download To** to **Server**, and click **OK** to generate the client configuration file. The generated file is saved in the **/tmp/MRS-client** directory on the active management node by default. You can customize the file path.

Step 3 Query and log in to the active Master node.

Step 4 On the ECS, switch to user **root** and copy the installation package to the **/opt** directory.

sudo su - root

cp /tmp/MRS-client/MRS_Services_Client.tar /opt

Step 5 Run the following command in the **/opt** directory to decompress the package and obtain the verification file and the configuration package of the client:

tar -xvf MRS_Services_Client.tar

Step 6 Run the following command to verify the configuration file package of the client:

sha256sum -c MRS_Services_ClientConfig.tar.sha256

The command output is as follows:

```
MRS_Services_ClientConfig.tar: OK
```

Step 7 Run the following command to decompress **MRS_Services_ClientConfig.tar**:

tar -xvf MRS_Services_ClientConfig.tar

Step 8 Run the following command to move the original client to the **/opt/Bigdata/client_bak** directory:

mv /opt/Bigdata/client /opt/Bigdata/client_bak

Step 9 Run the following command to install the client in a new directory. The client path must be **/opt/Bigdata/client**.

sh /opt/MRS_Services_ClientConfig/install.sh /opt/Bigdata/client

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

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

Step 10 Run the following command to modify the user and user group of the **/opt/Bigdata/client** directory:

chown omm:wheel /opt/Bigdata/client -R

Step 11 Run the following command to configure environment variables:

source /opt/Bigdata/client/bigdata_env

Step 12 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the user. If Kerberos authentication is disabled for the current cluster, skip this step.

kinit *MRS cluster user*

Example: **kinit admin**

Step 13 Run the client command of a component.

For example, run the following command to query the HDFS directory:

hdfs dfs -ls /

----End

Fully Updating the Original Client of the Standby Master Node

Step 1 Repeat **Step 1** to **Step 3** to log in to the standby Master node, and run the following command to switch to user **omm**:

sudo su - omm

Step 2 Run the following command on the standby master node to copy the downloaded client package from the active master node:

scp omm@*master1 nodeIP address***:/tmp/MRS-client/ MRS_Services_Client.tar /tmp/MRS-client/**

NOTE

- In this command, **master1** node is the active master node.
- **/tmp/MRS-client/** is an example target directory of the standby master node.

Step 3 Repeat **Step 4** to **Step 13** to update the client of the standby Master node.

----End

4.3 Using the Client of Each Component

4.3.1 Using a ClickHouse Client

ClickHouse is a column-based database oriented to online analysis and processing. It supports SQL query and provides good query performance. The aggregation analysis and query performance based on large and wide tables is excellent, which is one order of magnitude faster than other analytical databases.

Prerequisites

The client has been installed in a directory, for example, **/opt/client**. The client directory in the following operations is only an example. Change it to the actual installation directory. Before using the client, download and update the client configuration file, and ensure that the active management node of Manager is available.

Procedure

Step 1 Install the client. For details, see **Installing a Client**.

Step 2 Log in to the node where the client is installed as the client installation user.

Step 3 Run the following command to go to the client installation directory:

cd /opt/client

Step 4 Run the following command to configure environment variables:

source bigdata_env

Step 5 If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the user. The current user must have the permission to create ClickHouse tables. For details about how to configure the permission, see **ClickHouse User and Permission Management**, and bind roles to the user. If Kerberos authentication is disabled for the current cluster, skip this step.

For an MRS 3.1.0 cluster, run the **export CLICKHOUSE_SECURITY_ENABLED=true** command first.

kinit *Component service user*

Example: **kinit clickhouseuser**

Step 6 Run the client command of the ClickHouse component.

Run the **clickhouse -h** command to view the command help of ClickHouse.

The command output is as follows:

```
Use one of the following commands:
clickhouse local [args]
clickhouse client [args]
clickhouse benchmark [args]
clickhouse server [args]
clickhouse performance-test [args]
clickhouse extract-from-config [args]
clickhouse compressor [args]
clickhouse format [args]
clickhouse copier [args]
clickhouse obfuscator [args]
...
```

For MRS 3.1.0, run the **clickhouse client** command to connect to the ClickHouse server.

- Command for using a non-SSL mode to log in to a ClickHouse cluster with Kerberos authentication disabled

  clickhouse client --host *IP address of the ClickHouse instance* **--port 9000 --user** *Username* **--password**

  Enter the user password.

- Using SSL for login when Kerberos authentication is enabled for the current cluster:

  There are no default users in clusters with Kerberos authentication enabled. You must create a user on FusionInsight Manager. For details about how to create a user, see **ClickHouse User and Permission Management**.

  After the user authentication is successful, you do not need to carry the **--user** and **--password** parameters when logging in to the client as the authenticated user.

  clickhouse client --host *IP address of the ClickHouse instance* **--port 9440 --secure**

For MRS 3.1.2 or later, run the **clickhouse client** command to connect to the ClickHouse server.

- Command for using a non-SSL mode to log in to a ClickHouse cluster with Kerberos authentication disabled

  clickhouse client --host *IP address of the ClickHouse instance* **--port 9000 --user** *Username* **--password**

  Enter the user password.

- Using SSL for login when Kerberos authentication is enabled for the current cluster:

  There are no default users in clusters with Kerberos authentication enabled. You must create a user on FusionInsight Manager.

  clickhouse client --host *IP address of the ClickHouse instance* **--port 9440 --user** *Username* **--password --secure**

  Enter the user password.

Run the **quit;** command to exit the ClickHouse server connection.

Table 4-5 describes related parameters.

Table 4-5 Parameters of the **clickhouse client** command

Parameter	Description
--host	Host name of the server. The default value is **localhost**. You can use the host name or IP address of the node where the ClickHouse instance is located. **NOTE** You can log in to FusionInsight Manager and choose **Cluster** > **Services** > **ClickHouse** > **Instance** to obtain the service IP address of the ClickHouseServer instance.
--port	Port for connection. • If the SSL security connection is used, the default port number is **9440**, the parameter **--secure** must be carried. For details about the port number, search for the **tcp_port_secure** parameter in the ClickHouseServer instance configuration. • If non-SSL security connection is used, the default port number is **9000**, the parameter **--secure** does not need to be carried. For details about the port number, search for the **tcp_port** parameter in the ClickHouseServer instance configuration.

Parameter	Description
--user	Username. You can create a user on FusionInsight Manager and bind roles to it. For details about how to create a user, see **ClickHouse User and Permission Management**. • If Kerberos authentication has been enabled for the current cluster (the cluster is in security mode) and the user authentication is successful, you do not need to carry the **--user** and **--password** parameters during your login to the client as the authenticated user. You must create a user with this name on Manager because there is no default user in the Kerberos cluster scenario. • If Kerberos authentication has not been enabled for the current cluster (the cluster is in normal mode), you cannot use the ClickHouse user created on FusionInsight Manager if you need to specify the username and password when you log in to the client. You need to execute the **create user SQL** statement on the client to create a ClickHouse user. If you do not need to specify the username and password during your login to the client, the default user is used by default.
--password	Password. The default password is an empty string. This parameter is used together with the **--user** parameter. You can set a password when creating a user on Manager.
--query	Query to process when using non-interactive mode.
--database	Current default database. The default value is **default**, which is the default configuration on the server.
--multiline	If this parameter is specified, multiline queries are allowed. (**Enter** only indicates line feed and does not indicate that the query statement is complete.)
--multiquery	If this parameter is specified, multiple queries separated with semicolons (;) can be processed. This parameter is valid only in non-interactive mode.
--format	Specified default format used to output the result.
--vertical	If this parameter is specified, the result is output in vertical format by default. In this format, each value is printed on a separate line, which helps to display a wide table.
--time	If this parameter is specified, the query execution time is printed to **stderr** in non-interactive mode.
--stacktrace	If this parameter is specified, stack trace information will be printed when an exception occurs.
--config-file	Name of the configuration file.
--secure	If this parameter is specified, the server will be connected in SSL mode.

Parameter	Description
--history_file	Path of files that record command history.
--param_<name>	Query with parameters. Pass values from the client to the server. For details, see **https://clickhouse.tech/docs/en/interfaces/cli/#cli-queries-with-parameters**.

----End

4.3.2 Using a Flink Client

This section describes how to use Flink to run wordcount jobs.

Prerequisites

- Flink has been installed in an MRS cluster.
- The cluster runs properly and the client has been correctly installed, for example, in the **/opt/hadoopclient** directory. The client directory in the following operations is only an example. Change it to the actual installation directory.

Using the Flink Client (MRS 3.*x* or Later)

Step 1 Install a client.

The following uses installing a Flink client on a node in the cluster as an example:

1. Log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, click **More**, and select **Download Client**.
2. In the **Download Cluster Client** dialog box that is displayed, select **Complete Client** for **Select Client Type**, select the platform type that matches the architecture of the node where the client is to install, select **Save to Path**, and click **OK**.
   - The generated file is stored in the **/tmp/FusionInsight-Client** directory on the active management node by default.
   - The name of the client software package is in the follow format: **FusionInsight_Cluster_***<Cluster ID>***_Services_Client.tar**. In this section, the cluster ID **1** is used as an example. Replace it with the actual cluster ID.
3. Log in to the server where the client is to be installed as the client installation user.
4. Go to the directory where the installation package is stored and run the following commands to decompress the package:

   cd /tmp/FusionInsight-Client

   tar -xvf FusionInsight_Cluster_1_Services_Client.tar
5. Run the following command to verify the decompressed file and check whether the command output is consistent with the information in the **sha256** file:

sha256sum -c FusionInsight_Cluster_1_Services_ClientConfig.tar.sha256

```
FusionInsight_Cluster_1_Services_ClientConfig.tar: OK
```

6. Decompress the obtained installation file.

   tar -xvf FusionInsight_Cluster_1_Services_ClientConfig.tar

7. Go to the directory where the installation package is stored, and run the following command to install the client to a specified directory (absolute path), for example, **/opt/hadoopclient**:

   cd /tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_ClientConfig

   ./install.sh /opt/hadoopclient

   The client is installed if information similar to the following is displayed:

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

Step 2 Log in to the node where the client is installed as the client installation user.

Step 3 Run the following command to go to the client installation directory:

cd /opt/hadoopclient

Step 4 Run the following command to initialize environment variables:

source /opt/hadoopclient/bigdata_env

Step 5 Perform the following operations if Kerberos authentication is enabled for the cluster. Otherwise, skip these operations.

1. Prepare a user for submitting Flink jobs.

   Log in to FusionInsight Manager and choose **System** > **Permission** > **Role**. Click **Create Role** and configure **Role Name** and **Description**. In **Configure Resource Permission**, choose *Name of the desired cluster* > **Flink** and select **FlinkServer Admin Privilege**. Then click **OK**.

   Choose **System** > **Permission** > **User** and click **Create User**. Configure **Username**, set **User Type** to **Human-Machine**, configure **Password** and **Confirm Passowrd**, click **Add** next to **User Group** to add the **hadoop**, **yarnviewgroup**, and **hadooppmanager** user groups as needed, click **Add** next to **Role** to add the **System_administrator**, **default**, and the created role, and click **OK**. (If you create a Flink job user for the first time, log in to FusionInsight Manager as the user and change the password.)

2. Log in to Manager and download the authentication credential.

   Log in to FusionInsight Manager. For details, see **Accessing FusionInsight Manager (MRS 3.x or Later)**. Choose **System** > **Permission** > **User**. In the **Operation** column of the created user, click **More** and select **Download Authentication Credential**.

   Figure 4-4 Downloading the authentication credential

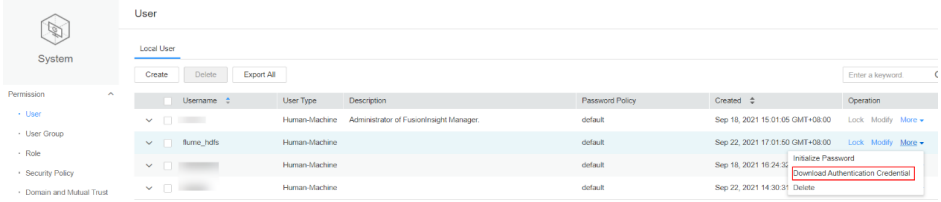

3. Decompress the downloaded authentication credential package and copy the obtained file to a directory on the client node, for example, **/opt/hadoopclient/Flink/flink/conf**. If the client is installed on a node outside the cluster, copy the obtained file to the **/etc/** directory on this node.
4. Add the service IP address of the node where the client is installed and IP addresses of all master nodes to the **jobmanager.web.access-control-allow-origin** and **jobmanager.web.allow-access-address** configuration items in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** file. Use commas (,) to separate the IP addresses.

   jobmanager.web.access-control-allow-origin: *xx.xx.xxx.xxx,xx.xx.xxx.xxx,xx.xx.xxx.xxx*
   jobmanager.web.allow-access-address: *xx.xx.xxx.xxx,xx.xx.xxx.xxx,xx.xx.xxx.xxx*

   NOTE

   To obtain the service IP address of the node where the client is installed, perform the following operations:

   - Node inside the cluster:

     In the navigation tree of the MRS management console, choose **Clusters > Active Clusters**, select a cluster, and click its name to switch to the cluster details page.

     On the **Nodes** tab page, view the IP address of the node where the client is installed.
   - Node outside the cluster: IP address of the ECS where the client is installed.
5. Configure security authentication by adding the **keytab** path and username in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** configuration file.

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

   Example:

   security.kerberos.login.keytab: /opt/hadoopclient/Flink/flink/conf/user.keytab
   security.kerberos.login.principal: test
6. In the **bin** directory of the Flink client, run the following command to perform security hardening and configure a password used for submitting jobs:

   cd /opt/hadoopclient/Flink/flink/bin

   sh generate_keystore.sh

   The script automatically replaces the SSL value in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** file.

   NOTE

   After authentication and encryption, the **flink.keystore** and **flink.truststore** files are generated in the **conf** directory on the Flink client and the following configuration items are set to the default values in the **flink-conf.yaml** file:

   - Set **security.ssl.keystore** to the absolute path of the **flink.keystore** file.
   - Set **security.ssl.truststore** to the absolute path of the **flink.truststore** file.
   - Set **security.cookie** to a random password automatically generated by the **generate_keystore.sh** script.
   - By default, **security.ssl.encrypt.enabled** is set to **false** in the **flink-conf.yaml** file by default. The **generate_keystore.sh** script sets **security.ssl.key-password**, **security.ssl.keystore-password**, and **security.ssl.truststore-password** to the password entered when the **generate_keystore.sh** script is called. There can be security risks if a configuration file contains the authentication password. You are advised to delete the configuration file or use other secure methods to keep the password.

7. Configure paths for the client to access the **flink.keystore** and **flink.truststore** files.
   - Relative path (recommended):

     Perform the following steps to set the file path of **flink.keystore** and **flink.truststore** to the relative path and ensure that the directory where the Flink client command is executed can directly access the relative paths.

     i. Create a directory, for example, **ssl**, in **/opt/hadoopclient/Flink/flink/conf/**.

     cd /opt/hadoopclient/Flink/flink/conf/

     mkdir ssl

     ii. Move the **flink.keystore** and **flink.truststore** files to the **/opt/hadoopclient/Flink/flink/conf/ssl/** directory.

     mv flink.keystore ssl/

     mv flink.truststore ssl/

     iii. Change the values of the following parameters to relative paths in the **flink-conf.yaml** file:

     ```
     security.ssl.keystore: ssl/flink.keystore
     security.ssl.truststore: ssl/flink.truststore
     ```

   - Absolute path:

     After the **generate_keystore.sh** script is executed, the file path of **flink.keystore** and **flink.truststore** is automatically set to the absolute path **/opt/hadoopclient/Flink/flink/conf/** in the **flink-conf.yaml** file. In this case, you need to move the **flink.keystore** and **flink.truststore** files from the **conf** directory to this absolute path on the Flink client and YARN nodes.

Step 6 Run a wordcount job.

NOTICE

When a user submits or runs a job in Flink, the user must have the following permissions based on whether Ranger authentication is enabled for related services (such as HDFS and Kafka):

- If Ranger authentication is enabled, the current user must belong to the **hadoop** group or the user has been granted the **/flink** read and write permissions in Ranger.
- If Ranger authentication is disabled, the current user must belong to the **hadoop** group.

- For a normal cluster (Kerberos authentication disabled), you can submit jobs in either of the following ways:
  - Run the following commands to start a session and submit a job in the session:

    yarn-session.sh -nm "*session-name***" -d**

    flink run /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

  - Run the following command to submit a single job on Yarn:

    flink run -m yarn-cluster /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

- For a security cluster (Kerberos authentication enabled), you can submit jobs in either of the following ways based on the paths of the **flink.keystore** and **flink.truststore** files:
  - If the **flink.keystore** and **flink.truststore** files are stored in the relative path:
    - Run the following command in the directory at the same level as **ssl** to start the session and submit the job in the session:

      ssl is a relative path. For example, if **ssl** is in **opt/hadoopclient/Flink/flink/conf/**, run the command in the **opt/hadoopclient/Flink/flink/conf/** directory.

      cd /opt/hadoopclient/Flink/flink/conf

      yarn-session.sh -t ssl/ -nm "*session-name***" -d**

      flink run /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

    - Run the following command to submit a single job on Yarn:

      cd /opt/hadoopclient/Flink/flink/conf

      flink run -m yarn-cluster -yt ssl/ /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
  - If the **flink.keystore** and **flink.truststore** files are stored in the absolute path:
    - Run the following commands to start a session and submit a job in the session:

      cd /opt/hadoopclient/Flink/flink/conf

      yarn-session.sh -nm "*session-name***" -d**

      flink run /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

    - Run the following command to submit a single job on Yarn:

      flink run -m yarn-cluster /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

Step 7 After the job has been successfully submitted, the following information is displayed on the client:

Figure 4-5 Job submitted successfully on Yarn

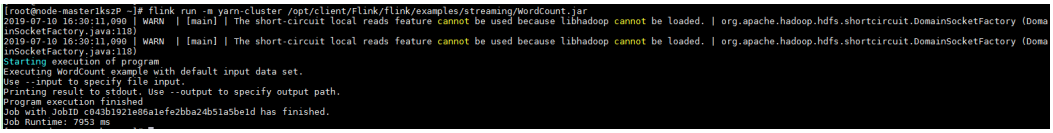

Figure 4-6 Session started successfully

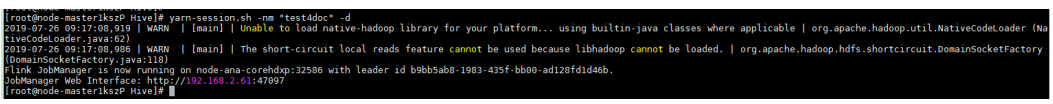

Figure 4-7 Job submitted successfully in the session

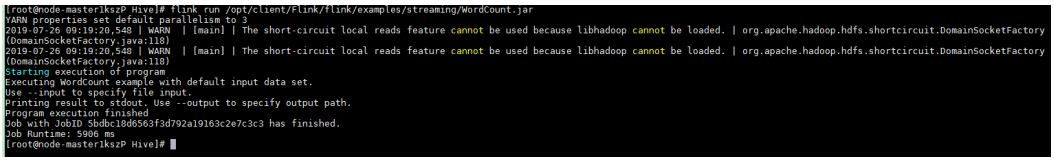

Step 8 Go to the native Yarn service page, find the application of the job, and click the application name to go to the job details page.

- If the job is not completed, click **Tracking URL** to go to the native Flink page and view the job running information.
- If the job submitted in a session has been completed, you can click **Tracking URL** to log in to the native Flink service page to view job information.

Figure 4-8 application

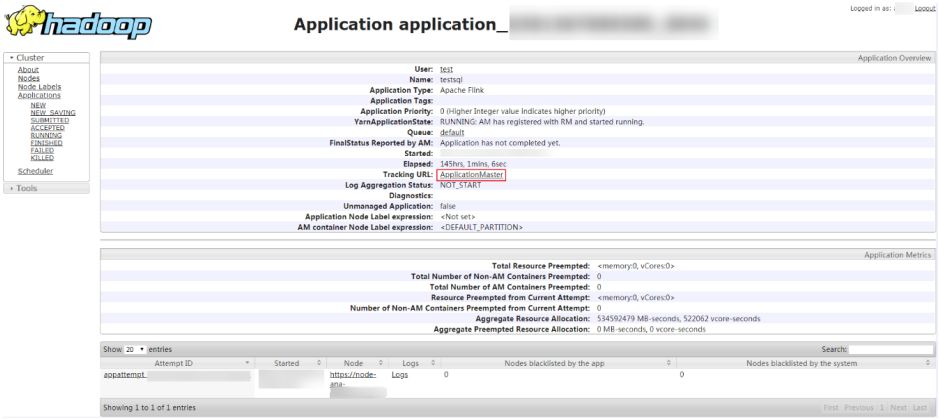

----End

Using the Flink Client (Versions Earlier Than MRS 3.x)

Step 1 Install a client.

The following uses installing a Flink client on a core node as an example:

1. Log in to MRS Manager and choose **Services** > **Download Client** to download the client installation package to the active management node.
2. Use the IP address to search for the active management node, and log in to the active management node over VNC.
3. Log in to the active management node, and run the following command to switch the user:

   sudo su - omm
4. On the MRS management console, view the IP address on the **Nodes** tab page of the specified cluster.

   Record the IP address of the core node where the client is to be used.
5. On the active management node, run the following command to copy the client installation package to the core node:

   scp -p /tmp/MRS-client/MRS_Services_Client.tar *IP address of the core node*:**/opt/client**

6. Log in to the core node as user **root**.

   Master nodes support Cloud-Init. The preset username for Cloud-Init is **root**, and the password is the one set during cluster creation.

7. Run the following commands to install the client:

   cd /opt/client

   tar -xvf MRS_Services_Client.tar

   tar -xvf MRS_Services_ClientConfig.tar

   cd /opt/client/MRS_Services_ClientConfig

   ./install.sh *Client installation directory*

   For example, run the following command:

   ./install.sh /opt/hadoopclient

Step 2 Log in to the node where the client is installed as the client installation user.

Step 3 Run the following command to go to the client installation directory:

cd /opt/hadoopclient

Step 4 Run the following command to initialize environment variables:

source /opt/hadoopclient/bigdata_env

Step 5 If Kerberos authentication is enabled for the cluster, perform the following steps. If not, skip this whole step.

1. Prepare a user for submitting Flink jobs.

   Log in to MRS Manager and choose **System** > **Manage Role**. Click **Add Role** to add a role, for example, **flinkrole**. In **Permission**, choose **HDFS** > **File System** > **hdfs://hacluster/** and select **Read**, **Write**, and **Execute**. After you finish configuring this service, click **Service** in the **Permission** area. Choose **Yarn** > **Scheduler Queue** > **root**, select **Submit** for **default**, and click **OK**.

   Choose **System** > **Manage User Group**. Click **Create User Group** to create a user group for the sample project, for example, **flinkgroup**. Choose **System** > **Manage User**. Click **Create User** to create a user for the sample project. Enter a username, for example, **flinkuser**. Set **User Type** to **Human-machine**, add the user to user groups **flinkgroup** and **hadoop**, bind the role **flinkrole** to the user to obtain permissions, and click **OK**. (If you create a user for the first time, log in to MRS Manager as this user to change the password.)

2. Log in to Manager and download the authentication credential.

   Log in to Manager of the cluster. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**. Choose **System Settings** > **User Management**. In the **Operation** column of the row that contains the added user, choose **More** > **Download Authentication Credential**.

   Figure 4-9 Downloading the authentication credential

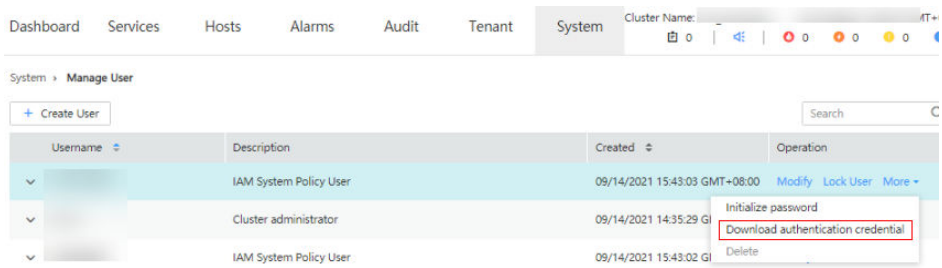

3. Decompress the downloaded authentication credential package and copy the obtained file to a directory on the client node, for example, **/opt/hadoopclient/Flink/flink/conf**. If the client is installed on a node outside the cluster, copy the obtained file to the **/etc/** directory on this node.
4. Configure security authentication by adding the **keytab** path and username in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** configuration file.
   security.kerberos.login.keytab: *<user.keytab file path>*
   security.kerberos.login.principal: *<Username>*

   Example:

   ```
   security.kerberos.login.keytab: /opt/hadoopclient/Flink/flink/conf/user.keytab
   security.kerberos.login.principal: test
   ```

5. Generate the **generate_keystore.sh** script by referring to **Example of Issuing a Certificate** and place it in the **bin** directory on the Flink client. In the **bin** directory of the Flink client, run the following command to perform security hardening and configure a password used for submitting jobs:

   cd /opt/hadoopclient/Flink/flink/bin

   sh generate_keystore.sh

   This script automatically replaces the SSL value in the **/opt/hadoopclient/Flink/flink/conf/flink-conf.yaml** file. For MRS 2.*x* or earlier, external SSL is disabled for security clusters by default. To enable external SSL, configure it by referring to "Authentication and Encryption" and run this script again.

   NOTE

   - You do not need to manually generate the **generate_keystore.sh** script.
   - After authentication and encryption, the generated **flink.keystore**, **flink.truststore**, and **security.cookie** items are automatically filled in the corresponding configuration items in **flink-conf.yaml**.

6. Configure paths for the client to access the **flink.keystore** and **flink.truststore** files.
   - Relative path (recommended):

     Perform the following steps to set the file path of **flink.keystore** and **flink.truststore** to the relative path and ensure that the directory where the Flink client command is executed can directly access the relative paths.

     i. Create a directory, for example, **ssl**, in **/opt/hadoopclient/Flink/flink/conf/**.

        cd /opt/hadoopclient/Flink/flink/conf/

        mkdir ssl

     ii. Move the **flink.keystore** and **flink.truststore** files to the **/opt/hadoopclient/Flink/flink/conf/ssl/** directory.

        mv flink.keystore ssl/

        mv flink.truststore ssl/

     iii. Change the values of the following parameters to relative paths in the **flink-conf.yaml** file:

        ```
        security.ssl.internal.keystore: ssl/flink.keystore
        security.ssl.internal.truststore: ssl/flink.truststore
        ```

   - Absolute path:

     After the **generate_keystore.sh** script is executed, the file path of **flink.keystore** and **flink.truststore** is automatically set to the absolute

path **/opt/hadoopclient/Flink/flink/conf/** in the **flink-conf.yaml** file. In this case, you need to move the **flink.keystore** and **flink.truststore** files from the **conf** directory to this absolute path on the Flink client and YARN nodes.

Step 6 Run a wordcount job.

NOTICE

When a user submits or runs a job in Flink, the user must have the following permissions based on whether Ranger authentication is enabled for related services (such as HDFS and Kafka):

- If Ranger authentication is enabled, the current user must belong to the **hadoop** group or the user has been granted the **/flink** read and write permissions in Ranger.
- If Ranger authentication is disabled, the current user must belong to the **hadoop** group.

- For a normal cluster (Kerberos authentication disabled), you can submit jobs in either of the following ways:
  - Run the following commands to start a session and submit a job in the session:

    yarn-session.sh -nm "*session-name***" -d**

    flink run /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
  - Run the following command to submit a single job on Yarn:

    flink run -m yarn-cluster /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
- For a security cluster (Kerberos authentication enabled), you can submit jobs in either of the following ways based on the paths of the **flink.keystore** and **flink.truststore** files:
  - If the **flink.keystore** and **flink.truststore** files are stored in the relative path:
    - Run the following command in the directory at the same level as **ssl** to start the session and submit the job in the session:

      ssl is a relative path. For example, if **ssl** is in **opt/hadoopclient/Flink/flink/conf/**, run the command in the **opt/hadoopclient/Flink/flink/conf/** directory.

      cd /opt/hadoopclient/Flink/flink/conf

      yarn-session.sh -t ssl/ -nm "*session-name***" -d**

      flink run /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar
    - Run the following command to submit a single job on Yarn:

      cd /opt/hadoopclient/Flink/flink/conf

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

  - If the **flink.keystore** and **flink.truststore** files are stored in the absolute path:
    - Run the following commands to start a session and submit a job in the session:

      cd /opt/hadoopclient/Flink/flink/conf

      yarn-session.sh -nm "*session-name***" -d**

      flink run /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

    - Run the following command to submit a single job on Yarn:

      flink run -m yarn-cluster /opt/hadoopclient/Flink/flink/examples/streaming/WordCount.jar

Step 7 After the job has been successfully submitted, the following information is displayed on the client:

Figure 4-10 Job submitted successfully on Yarn

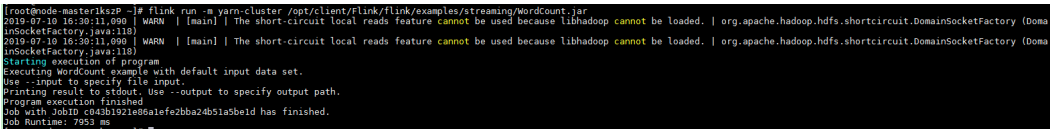

Figure 4-11 Session started successfully

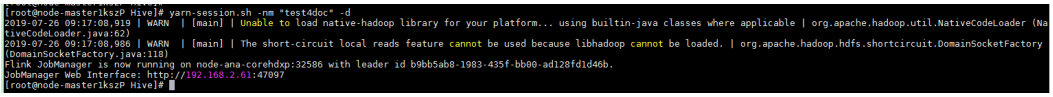

Figure 4-12 Job submitted successfully in the session

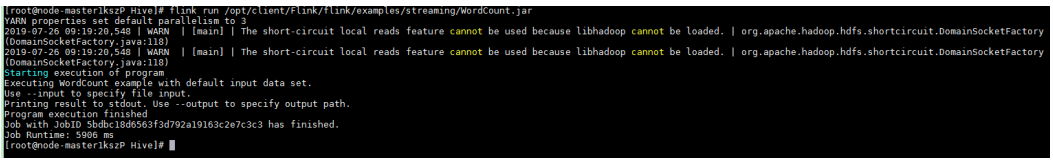

Step 8 Go to the native Yarn service page, find the application of the job, and click the application name to go to the job details page.

- If the job is not completed, click **Tracking URL** to go to the native Flink page and view the job running information.
- If the job submitted in a session has been completed, you can click **Tracking URL** to log in to the native Flink service page to view job information.

Figure 4-13 Application

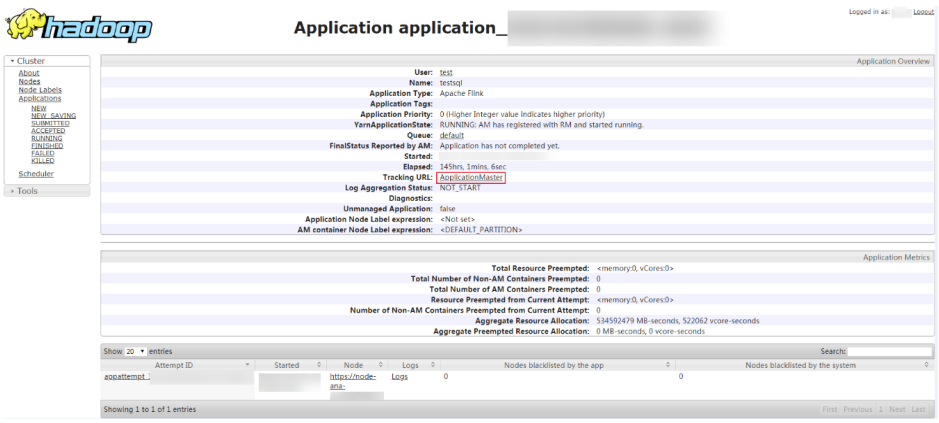

----End

4.3.3 Using a Flume Client

Scenario

You can use Flume to import collected log information to Kafka.

Prerequisites

- A streaming cluster that contains components such as Flume and Kafka and has Kerberos authentication enabled has been created. For details, see **Buying a Custom Cluster**.
- The streaming cluster can properly communicate with the node where logs are generated.

Using the Flume Client (Versions Earlier Than MRS 3.x)

NOTE

You do not need to perform **Step 2** to **Step 6** for a normal cluster.

Step 1 Install the Flume client.

Install the Flume client in a directory, for example, **/opt/Flumeclient**, on the node where logs are generated by referring to **Installing the Flume Client on Clusters of Versions Earlier Than MRS 3.x**. The Flume client installation directories in the following steps are only examples. Change them to the actual installation directories.

Step 2 Copy the configuration file of the authentication server from the Master1 node to the *Flume client installation directory*/**fusioninsight-flume-***Flume component version number*/**conf** directory on the node where the Flume client is installed.

${BIGDATA_HOME}/MRS_Current/1_*X*_**KerberosClient/etc/kdc.conf** is used as the full file path.

In the preceding paths, **X** indicates a random number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 3 Check the service IP address of any node where the Flume role is deployed.

Log in to the cluster details page, choose *Name of the desired cluster* > **Components** > **Flume** > **Instances**, and check the service IP address of any node where the Flume role is deployed.

NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

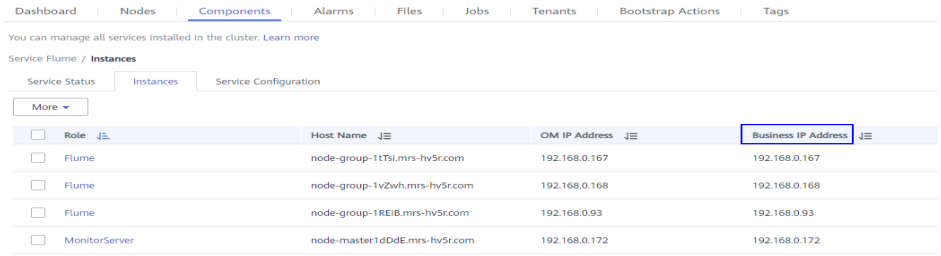

Step 4 Copy the user authentication file from this node to the *Flume client installation directory*/**fusioninsight-flume-Flume component version number/conf** directory on the Flume client node.

${BIGDATA_HOME}/MRS_*XXX***/install/FusionInsight-Flume-***Flume component version number***/flume/conf/flume.keytab** is used as the full file path.

In the preceding paths, **XXX** indicates the product version number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 5 Copy the **jaas.conf** file from this node to the **conf** directory on the Flume client node.

${BIGDATA_HOME}/MRS_Current/1_*X***_Flume/etc/jaas.conf** is used as the full file path.

In the preceding path, **X** indicates a random number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 6 Log in to the Flume client node and go to the client installation directory. Run the following command to modify the file:

vi conf/jaas.conf

Change the full path of the user authentication file defined by **keyTab** to the **Flume client installation directory/fusioninsight-flume-*Flume component version number*/conf** saved in **Step 4**, and save the modification and exit.

Step 7 Run the following command to modify the **flume-env.sh** configuration file of the Flume client:

vi *Flume client installation directory*/**fusioninsight-flume-***Flume component version number***/conf/flume-env.sh**

Add the following information after **-XX:+UseCMSCompactAtFullCollection**:


```
-Djava.security.krb5.conf=Flume client installation directory/fusioninsight-flume-1.9.0/conf/kdc.conf -
Djava.security.auth.login.config=Flume client installation directory/fusioninsight-flume-1.9.0/conf/jaas.conf -
Dzookeeper.request.timeout=120000
```

Example: **"-XX:+UseCMSCompactAtFullCollection -Djava.security.krb5.conf=/opt/FlumeClient/fusioninsight-flume-***Flume component version number***/conf/kdc.conf -Djava.security.auth.login.config=/opt/FlumeClient/fusioninsight-flume-***Flume component version number***/conf/jaas.conf -Dzookeeper.request.timeout=120000"**

Change *Flume client installation directory* to the actual installation directory. Then save and exit.

Step 8 Run the following command to restart the Flume client:

cd *Flume client installation directory***/fusioninsight-flume-***Flume component version number***/bin**

./flume-manage.sh restart

Example:

cd /opt/FlumeClient/fusioninsight-flume-*Flume component version number***/bin**

./flume-manage.sh restart

NOTE

The Flume client will be automatically restarted after being stopped. If you do not need automatic restart, run the following command:

./flume-manage.sh stop force

If you want to restart the Flume client, run the following command:

./flume-manage.sh start force

Step 9 Run the following command to configure and save jobs in the Flume client configuration file **properties.properties** based on service requirements.

vi *Flume client installation directory***/fusioninsight-flume-***Flume component version number***/conf/properties.properties**

The following uses SpoolDir Source+File Channel+Kafka Sink as an example:

```
################################################################################################
client.sources = static_log_source
client.channels = static_log_channel
client.sinks = kafka_sink
################################################################################################
#LOG_TO_HDFS_ONLINE_1

client.sources.static_log_source.type = spooldir
client.sources.static_log_source.spoolDir = Monitoring directory
client.sources.static_log_source.fileSuffix = .COMPLETED
client.sources.static_log_source.ignorePattern = ^$
client.sources.static_log_source.trackerDir = Metadata storage path during transmission
client.sources.static_log_source.maxBlobLength = 16384
client.sources.static_log_source.batchSize = 51200
client.sources.static_log_source.inputCharset = UTF-8
client.sources.static_log_source.deserializer = LINE
client.sources.static_log_source.selector.type = replicating
client.sources.static_log_source.fileHeaderKey = file
```



```
client.sources.static_log_source.fileHeader = false
client.sources.static_log_source.basenameHeader = true
client.sources.static_log_source.basenameHeaderKey = basename
client.sources.static_log_source.deletePolicy = never

client.channels.static_log_channel.type = file
client.channels.static_log_channel.dataDirs = Data cache path. Multiple paths, separated by commas (,), can be configured to improve performance.
client.channels.static_log_channel.checkpointDir = Checkpoint storage path
client.channels.static_log_channel.maxFileSize = 2146435071
client.channels.static_log_channel.capacity = 1000000
client.channels.static_log_channel.transactionCapacity = 612000
client.channels.static_log_channel.minimumRequiredSpace = 524288000

client.sinks.kafka_sink.type = org.apache.flume.sink.kafka.KafkaSink
client.sinks.kafka_sink.kafka.topic = Topic to which data is written, for example, flume_test
client.sinks.kafka_sink.kafka.bootstrap.servers = XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number
client.sinks.kafka_sink.flumeBatchSize = 1000
client.sinks.kafka_sink.kafka.producer.type = sync
client.sinks.kafka_sink.kafka.security.protocol = SASL_PLAINTEXT
client.sinks.kafka_sink.kafka.kerberos.domain.name = Kafka domain name. This parameter is mandatory for a security cluster, for example, hadoop.xxx.com.
client.sinks.kafka_sink.requiredAcks = 0

client.sources.static_log_source.channels = static_log_channel
client.sinks.kafka_sink.channel = static_log_channel
```

NOTE

- **client.sinks.kafka_sink.kafka.topic**: Topic to which data is written. If the topic does not exist in Kafka, it is automatically created by default.
- **client.sinks.kafka_sink.kafka.bootstrap.servers**: List of Kafka Brokers, which are separated by commas (,). By default, the port is **21007** for a security cluster and **9092** for a normal cluster.
- **client.sinks.kafka_sink.kafka.security.protocol**: The value is **SASL_PLAINTEXT** for a security cluster and **PLAINTEXT** for a normal cluster.
- **client.sinks.kafka_sink.kafka.kerberos.domain.name**:

  You do not need to set this parameter for a normal cluster. For a security cluster, the value of this parameter is the value of **kerberos.domain.name** in the Kafka cluster.

  Obtain the value by checking **${BIGDATA_HOME}/MRS_Current/1_*X*_Broker/etc/server.properties** on the node where the broker instance resides.

  In the preceding paths, **X** indicates a random number. Change it based on site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 10 After the parameters are set and saved, the Flume client automatically loads the content configured in **properties.properties**. When new log files are generated by spoolDir, the files are sent to Kafka producers and can be consumed by Kafka consumers.

----End

Using the Flume Client (MRS 3.*x* or Later)

NOTE

You do not need to perform **Step 2** to **Step 6** for a normal cluster.

Step 1 Install the Flume client.

Install the Flume client in a directory, for example, **/opt/Flumeclient**, on the node where logs are generated by referring to **Installing the Flume Client on MRS 3.x**

or Later Clusters. The Flume client installation directories in the following steps are only examples. Change them to the actual installation directories.

Step 2 Copy the configuration file of the authentication server from the Master1 node to the *Flume client installation directory*/**fusioninsight-flume-***Flume component version number*/**conf** directory on the node where the Flume client is installed.

The full file path is **${BIGDATA_HOME}/FusionInsight_BASE_***XXX*/**1_***X*_KerberosClient/etc/kdc.conf**. In the preceding path, **XXX** indicates the product version number. **X** indicates a random number. Replace them based on site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 3 Check the service IP address of any node where the Flume role is deployed.

Log in to FusionInsight Manager, click **Cluster**, click **Services**, and click Flume. On the displayed page, click **Instance**. Check the service IP address of any node where the Flume role is deployed.

NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 4 Copy the user authentication file from this node to the *Flume client installation directory*/**fusioninsight-flume-Flume component version number/conf** directory on the Flume client node.

The full file path is **${BIGDATA_HOME}/FusionInsight_Porter_***XXX*/**install/FusionInsight-Flume-***Flume component version number*/**flume/conf/flume.keytab**.

In the preceding paths, **XXX** indicates the product version number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 5 Copy the **jaas.conf** file from this node to the **conf** directory on the Flume client node.

The full file path is **${BIGDATA_HOME}/FusionInsight_Current/1_***X*_Flume/etc/jaas.conf**.

In the preceding path, **X** indicates a random number. Change it based on the site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 6 Log in to the Flume client node and go to the client installation directory. Run the following command to modify the file:

vi conf/jaas.conf

Change the full path of the user authentication file defined by **keyTab** to the **Flume client installation directory/fusioninsight-flume-*Flume component version number*/conf** saved in **Step 4**, and save the modification and exit.

Step 7 Run the following command to modify the **flume-env.sh** configuration file of the Flume client:

vi *Flume client installation directory*/**fusioninsight-flume-**_Flume component version number_/**conf/flume-env.sh**

Add the following information after **-XX:+UseCMSCompactAtFullCollection**:

```
-Djava.security.krb5.conf=Flume client installation directory/fusioninsight-flume-1.9.0/conf/kdc.conf -Djava.security.auth.login.config=Flume client installation directory/fusioninsight-flume-1.9.0/conf/jaas.conf -Dzookeeper.request.timeout=120000
```

Example: **"-XX:+UseCMSCompactAtFullCollection -Djava.security.krb5.conf=/opt/FlumeClient/fusioninsight-flume-***Flume component version number*/**conf/kdc.conf -Djava.security.auth.login.config=/opt/FlumeClient/fusioninsight-flume-***Flume component version number*/**conf/jaas.conf -Dzookeeper.request.timeout=120000"**

Change *Flume client installation directory* to the actual installation directory. Then save and exit.

Step 8 Run the following command to restart the Flume client:

cd *Flume client installation directory*/**fusioninsight-flume-***Flume component version number*/**bin**

./flume-manage.sh restart

Example:

cd /opt/FlumeClient/fusioninsight-flume-*Flume component version number*/**bin**

./flume-manage.sh restart

Step 9 Configure jobs based on actual service scenarios.

- Some parameters, for MRS 3.*x* or later, can be configured on Manager.
- Set the parameters in the **properties.properties** file. The following uses SpoolDir Source+File Channel+Kafka Sink as an example.

  Run the following command on the node where the Flume client is installed. Configure and save jobs in the Flume client configuration file **properties.properties** based on actual service requirements.

  vi *Flume client installation directory*/**fusioninsight-flume-***Flume component version number*/**conf/properties.properties**

  ```
  ################################################################################################################
  client.sources = static_log_source
  client.channels = static_log_channel
  client.sinks = kafka_sink
  ################################################################################################################
  #LOG_TO_HDFS_ONLINE_1

  client.sources.static_log_source.type = spooldir
  client.sources.static_log_source.spoolDir = Monitoring directory
  client.sources.static_log_source.fileSuffix = .COMPLETED
  client.sources.static_log_source.ignorePattern = ^$
  client.sources.static_log_source.trackerDir = Metadata storage path during transmission
  client.sources.static_log_source.maxBlobLength = 16384
  client.sources.static_log_source.batchSize = 51200
  client.sources.static_log_source.inputCharset = UTF-8
  client.sources.static_log_source.deserializer = LINE
  client.sources.static_log_source.selector.type = replicating
  client.sources.static_log_source.fileHeaderKey = file
  ```



```
client.sources.static_log_source.fileHeader = false
client.sources.static_log_source.basenameHeader = true
client.sources.static_log_source.basenameHeaderKey = basename
client.sources.static_log_source.deletePolicy = never

client.channels.static_log_channel.type = file
client.channels.static_log_channel.dataDirs = Data cache path. Multiple paths, separated by commas (,), can be configured to improve performance.
client.channels.static_log_channel.checkpointDir = Checkpoint storage path
client.channels.static_log_channel.maxFileSize = 2146435071
client.channels.static_log_channel.capacity = 1000000
client.channels.static_log_channel.transactionCapacity = 612000
client.channels.static_log_channel.minimumRequiredSpace = 524288000

client.sinks.kafka_sink.type = org.apache.flume.sink.kafka.KafkaSink
client.sinks.kafka_sink.kafka.topic = Topic to which data is written, for example, flume_test
client.sinks.kafka_sink.kafka.bootstrap.servers = XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number,XXX.XXX.XXX.XXX:Kafka port number
client.sinks.kafka_sink.flumeBatchSize = 1000
client.sinks.kafka_sink.kafka.producer.type = sync
client.sinks.kafka_sink.kafka.security.protocol = SASL_PLAINTEXT
client.sinks.kafka_sink.kafka.kerberos.domain.name = Kafka domain name. This parameter is mandatory for a security cluster, for example, hadoop.xxx.com.
client.sinks.kafka_sink.requiredAcks = 0

client.sources.static_log_source.channels = static_log_channel
client.sinks.kafka_sink.channel = static_log_channel
```

NOTE

- **client.sinks.kafka_sink.kafka.topic**: Topic to which data is written. If the topic does not exist in Kafka, it is automatically created by default.
- **client.sinks.kafka_sink.kafka.bootstrap.servers**: List of Kafka Brokers, which are separated by commas (,). By default, the port is **21007** for a security cluster and **9092** for a normal cluster.
- **client.sinks.kafka_sink.kafka.security.protocol**: The value is **SASL_PLAINTEXT** for a security cluster and **PLAINTEXT** for a normal cluster.
- **client.sinks.kafka_sink.kafka.kerberos.domain.name**:

  You do not need to set this parameter for a normal cluster. For a security cluster, the value of this parameter is the value of **kerberos.domain.name** in the Kafka cluster.

  Obtain the value by checking **${BIGDATA_HOME}/MRS_Current/1_*X*_Broker/etc/server.properties** on the node where the broker instance resides.

  In the preceding paths, **X** indicates a random number. Change it based on site requirements. The file must be saved by the user who installs the Flume client, for example, user **root**.

Step 10 After the parameters are set and saved, the Flume client automatically loads the content configured in **properties.properties**. When new log files are generated by spoolDir, the files are sent to Kafka producers and can be consumed by Kafka consumers.

----End

4.3.4 Using an HBase Client

Scenario

This section describes how to use the HBase client in an O&M scenario or a service scenario.

Prerequisites

- The client has been installed. For example, the installation directory is **/opt/hadoopclient**. The client directory in the following operations is only an example. Change it to the actual installation directory.
- Service component users have been created by the MRS cluster administrator.

  A machine-machine user needs to download the **keytab** file and a human-machine user needs to change the password upon the first login.
- If a non-**root** user uses the HBase client, ensure that the owner of the HBase client directory is this user. Otherwise, run the following command to change the owner.

  chown user:group -R *Client installation directory*/**HBase**

Using the HBase Client (Versions Earlier Than MRS 3.x)

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

Step 2 Run the following command to go to the client directory:

cd /opt/hadoopclient

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create HBase tables. If Kerberos authentication is disabled for the current cluster, skip this step.

kinit *Component service user*

For example, **kinit hbaseuser**.

Step 5 Run the following HBase client command:

hbase shell

----End

Using the HBase Client (MRS 3.*x* or Later)

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

Step 2 Run the following command to go to the client directory:

cd /opt/hadoopclient

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If Kerberos authentication has been enabled for the current cluster, run the following command to authenticate the current user. The current user must have the permission to create HBase tables. If Kerberos authentication is disabled for the current cluster, skip this step.

kinit *Component service user*

For example, **kinit hbaseuser**.

Step 5 Run the following HBase client command:

hbase shell

----End

Common HBase client commands

The following table lists common HBase client commands. For more commands, see **http://hbase.apache.org/2.2/book.html**.

Table 4-6 HBase client commands

Command	Description
create	Used to create a table, for example, **create 'test', 'f1', 'f2', 'f3'**.
disable	Used to disable a specified table, for example, **disable 'test'**.
enable	Used to enable a specified table, for example, **enable 'test'**.
alter	Used to alter the table structure. You can run the **alter** command to add, modify, or delete column family information and table-related parameter values, for example, **alter 'test', {NAME => 'f3', METHOD => 'delete'}**.
describe	Used to obtain the table description, for example, **describe 'test'**.
drop	Used to delete a specified table, for example, **drop 'test'**. Before deleting a table, you must stop it.
put	Used to write the value of a specified cell, for example, **put 'test','r1','f1:c1','myvalue1'**. The cell location is unique and determined by the table, row, and column.
get	Used to get the value of a row or the value of a specified cell in a row, for example, **get 'test','r1'**.
scan	Used to query table data, for example, **scan 'test'**. The table name and scanner must be specified in the command.

4.3.5 Using an HDFS Client

Scenario

This section describes how to use the HDFS client in an O&M scenario or service scenario.

Prerequisites

- The client has been installed.

For example, the installation directory is **/opt/client**. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite.

- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user needs to change the password upon the first login. (This operation is not required in normal mode.)

Using the HDFS Client

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

Step 2 Run the following command to go to the client installation directory:

cd /opt/client

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If the cluster is in security mode, run the following command to authenticate the user. In normal mode, user authentication is not required.

kinit *Component service user*

Step 5 Run the HDFS Shell command. Example:

hdfs dfs -ls /

----End

Common HDFS Client Commands

The following table lists common HDFS client commands.

Table 4-7 Common HDFS client commands

Command	Description	Example
hdfs dfs -mkdir *Folder name*	Used to create a folder.	**hdfs dfs -mkdir /tmp/mydir**
hdfs dfs -ls *Folder name*	Used to view a folder.	**hdfs dfs -ls /tmp**
hdfs dfs -put *Local file on the client node Specified HDFS path*	Used to upload a local file to a specified HDFS path.	**hdfs dfs -put /opt/test.txt /tmp** Upload the **/opt/test.txt** file on the client node to the **/tmp** directory of HDFS.
hdfs dfs -get *Specified file on HDFS Specified path on the client node*	Used to download the HDFS file to the specified local path.	**hdfs dfs -get /tmp/test.txt /opt/** Download the **/tmp/test.txt** file on HDFS to the **/opt** path on the client node.

Command	Description	Example
hdfs dfs -rm -r -f *Specified folder on HDFS*	Used to delete a folder.	**hdfs dfs -rm -r -f /tmp/mydir**
hdfs dfs -chmod *Permission parameter File directory*	Used to configure the HDFS directory permission for a user.	**hdfs dfs -chmod 700 /tmp/test**

Client-related FAQs

1. What do I do when the HDFS client exits abnormally and error message "java.lang.OutOfMemoryError" is displayed after the HDFS client command is running?

   This problem occurs because the memory required for running the HDFS client exceeds the preset upper limit (128 MB by default). You can change the memory upper limit of the client by modifying **CLIENT_GC_OPTS** in *<Client installation path>***/HDFS/component_env**. For example, if you want to set the upper limit to 1 GB, run the following command:

   ```
   CLIENT_GC_OPTS="-Xmx1G"
   ```

   After the modification, run the following command to make the modification take effect:

   source <*Client installation path*>/**/bigdata_env**

2. How do I set the log level when the HDFS client is running?

   By default, the logs generated during the running of the HDFS client are printed to the console. The default log level is INFO. To enable the DEBUG log level for fault locating, run the following command to export an environment variable:

   export HADOOP_ROOT_LOGGER=DEBUG,console

   Then run the HDFS Shell command to generate the DEBUG logs.

   If you want to print INFO logs again, run the following command:

   export HADOOP_ROOT_LOGGER=INFO,console

3. How do I delete HDFS files permanently?

   HDFS provides a recycle bin mechanism. Typically, after an HDFS file is deleted, the file is moved to the recycle bin of HDFS. If the file is no longer needed and the storage space needs to be released, clear the corresponding recycle bin directory, for example, **hdfs://hacluster/user/xxx/.Trash/Current/***xxx*.

4.3.6 Using a Hive Client

Scenario

This section guides users to use a Hive client in an O&M or service scenario.

Prerequisites

- The client has been installed. For example, the client is installed in the **/opt/hadoopclient** directory. The client directory in the following operations is only an example. Change it to the actual installation directory.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login.

Using the Hive Client (Versions Earlier Than MRS 3.x)

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

Step 2 Run the following command to go to the client installation directory:

cd /opt/hadoopclient

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 Log in to the Hive client based on the cluster authentication mode.

- In security mode, run the following command to complete user authentication and log in to the Hive client:

  kinit *Component service user*

  beeline

- In common mode, run the following command to log in to the Hive client. If no component service user is specified, the current OS user is used to log in to the Hive client.

  beeline -n *component service user*

  NOTE

  After a beeline connection is established, you can compile and submit HQL statements to execute related tasks. To run the Catalog client command, you need to run the **!q** command first to exit the beeline environment.

Step 5 Run the following command to execute the HCatalog client command:

hcat -e *"cmd"*

cmd must be a Hive DDL statement, for example, **hcat -e "show tables"**.

NOTE

- To use the HCatalog client, choose **More** > **Download Client** on the service page to download the clients of all services. This restriction does not apply to the beeline client.
- Due to permission model incompatibility, tables created using the HCatalog client cannot be accessed on the HiveServer client. However, the tables can be accessed on the WebHCat client.
- If you use the HCatalog client in Normal mode, the system performs DDL commands using the current user who has logged in to the operating system.
- Exit the beeline client by running the **!q** command instead of by pressing **Ctrl + c**. Otherwise, the temporary files generated by the connection cannot be deleted and a large number of junk files will be generated as a result.
- If multiple statements need to be entered during the use of beeline clients, separate the statements from each other using semicolons (**;**) and set the value of **entireLineAsCommand** to **false**.

  Setting method: If beeline has not been started, run the **beeline --entireLineAsCommand=false** command. If the beeline has been started, run the **!set entireLineAsCommand false** command.

  After the setting, if a statement contains semicolons (**;**) that do not indicate the end of the statement, escape characters must be added, for example, **select concat_ws('\;', collect_set(col1)) from tbl**.

----End

Using the Hive Client (MRS 3.*x* or Later)

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

Step 2 Run the following command to go to the client installation directory:

cd /opt/hadoopclient

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 Log in to the Hive client based on the cluster authentication mode.

- In security mode, run the following command to complete user authentication and log in to the Hive client:

  kinit *Component service user*

  beeline

- In common mode, run the following command to log in to the Hive client. If no component service user is specified, the current OS user is used to log in to the Hive client.

  beeline -n *component service user*

Step 5 Run the following command to execute the HCatalog client command:

hcat -e *"cmd"*

cmd must be a Hive DDL statement, for example, **hcat -e "show tables"**.

NOTE

- To use the HCatalog client, choose **More** > **Download Client** on the service page to download the clients of all services. This restriction does not apply to the beeline client.
- Due to permission model incompatibility, tables created using the HCatalog client cannot be accessed on the HiveServer client. However, the tables can be accessed on the WebHCat client.
- If you use the HCatalog client in Normal mode, the system performs DDL commands using the current user who has logged in to the operating system.
- Exit the beeline client by running the **!q** command instead of by pressing **Ctrl + C**. Otherwise, the temporary files generated by the connection cannot be deleted and a large number of junk files will be generated as a result.
- If multiple statements need to be entered during the use of beeline clients, separate the statements from each other using semicolons (**;**) and set the value of **entireLineAsCommand** to **false**.

  Setting method: If beeline has not been started, run the **beeline --entireLineAsCommand=false** command. If the beeline has been started, run the **!set entireLineAsCommand false** command.

  After the setting, if a statement contains semicolons (**;**) that do not indicate the end of the statement, escape characters must be added, for example, **select concat_ws('\;', collect_set(col1)) from tbl**.

----End

Common Hive Client Commands

The following table lists common Hive Beeline commands.

For more commands, see **https://cwiki.apache.org/confluence/display/Hive/HiveServer2+Clients#HiveServer2Clients-BeelineCommands**.

Table 4-8 Common Hive Beeline commands

Command	Description
set <key>=<value>	Sets the value of a specific configuration variable (key). **NOTE** If the variable name is incorrectly spelled, the Beeline does not display an error.
set	Prints the list of configuration variables overwritten by users or Hive.
set -v	Prints all configuration variables of Hadoop and Hive.
add FILE[S] <filepath> <filepath>* add JAR[S] <filepath> <filepath>* add ARCHIVE[S] <filepath> <filepath>*	Adds one or more files, JAR files, or ARCHIVE files to the resource list of the distributed cache.

Command	Description
add FILE[S] <ivyurl> <ivyurl>* add JAR[S] <ivyurl> <ivyurl>* add ARCHIVE[S] <ivyurl> <ivyurl>*	Adds one or more files, JAR files, or ARCHIVE files to the resource list of the distributed cache using the lvy URL in the **ivy://goup:module:version?query_string** format.
list FILE[S] list JAR[S] list ARCHIVE[S]	Lists the resources that have been added to the distributed cache.
list FILE[S] <filepath>* list JAR[S] <filepath>* list ARCHIVE[S] <filepath>*	Checks whether given resources have been added to the distributed cache.
delete FILE[S] <filepath>* delete JAR[S] <filepath>* delete ARCHIVE[S] <filepath>*	Deletes resources from the distributed cache.
delete FILE[S] <ivyurl> <ivyurl>* delete JAR[S] <ivyurl> <ivyurl>* delete ARCHIVE[S] <ivyurl> <ivyurl>*	Delete the resource added using **<ivyurl>** from the distributed cache.
reload	Enable HiveServer2 to discover the change of the JAR file **hive.reloadable.aux.jars.path** in the specified path. (You do not need to restart HiveServer2.) Change actions include adding, deleting, or updating JAR files.
dfs <dfs command>	Runs the **dfs** command.
<query string>	Executes the Hive query and prints the result to the standard output.

4.3.7 Using an Impala Client

Impala is a massively parallel processing (MPP) SQL query engine for processing vast amounts of data stored in Hadoop clusters. It is an open source software written in C++ and Java. It provides high performance and low latency compared with other SQL engines for Hadoop.

Background

Suppose a user develops an application to manage users who use service A in an enterprise. The procedure of operating service A on the Impala client is as follows:

Operations on common tables:

- Create the **user_info** table.
- Add users' educational backgrounds and titles to the table.
- Query user names and addresses by user ID.
- Delete the user information table after service A ends.

Table 4-9 User information

No.	Name	Gender	Age	Address
12005000201	A	Male	19	City A
12005000202	B	Female	23	City B
12005000203	C	Male	26	City C
12005000204	D	Male	18	City D
12005000205	E	Female	21	City E
12005000206	F	Male	32	City F
12005000207	G	Female	29	City G
12005000208	H	Female	30	City H
12005000209	I	Male	26	City I
12005000210	J	Female	25	City J

Prerequisites

The client has been installed. For example, the client is installed in the **/opt/hadoopclient** directory. The client directory in the following operations is only an example. Change it to the actual installation directory.

Procedure

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

Step 2 Run the following command to go to the client installation directory:

cd /opt/hadoopclient

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 Run the Impala client command to implement service A.

- **Operations on internal tables**:

  Run the Impala client command **impala-shell**.

  NOTE

  By default, **impala-shell** attempts to connect to the Impala daemon on port 21000 of **localhost**. To connect to another host, use the **-i < host:port >** option, for example, **impala-shell -i** *xxx.xxx.xxx.xxx***:21000**. To automatically connect to a specific Impala database, use the **-d <database>** option. For example, if all your Kudu tables are in the **impala_kudu** database, **-d impala_kudu** can use this database. To exit Impala Shell, run the **quit** command.

  a. Create the **user_info** user information table according to **Table 4-9** and add data to it.

  ```
  create table user_info(id string,name string,gender string,age int,addr string);
  insert into table user_info(id,name,gender,age,addr) values("12005000201", "A", "Male", 19, "City A");
  ```

  ... (Other statements are the same.)

  b. Add users' educational backgrounds and titles to the **user_info** table.

  For example, to add educational background and title information about user 12005000201, run the following commands.

  ```
  alter table user_info add columns(education string,technical string);
  ```

  c. Query user names and addresses by user ID.

  For example, to query the name and address of user 12005000201, run the following command:

  ```
  select name,addr from user_info where id='12005000201';
  ```

  d. Delete the user information table:

  ```
  drop table user_info;
  ```

- **Operations on external partition tables**:

  Create an external partition table and import data.

  a. Create a path for storing external table data.

  kinit *hive* (Run this command only in security mode.)

  NOTE

  The user must have the hive administrator permissions.

  hdfs dfs -mkdir /hive

  hdfs dfs -mkdir /hive/user_info

  b. Create a table.

  impala-shell

  NOTE

  By default, **impala-shell** attempts to connect to the Impala daemon on port 21000 of **localhost**. To connect to another host, use the **-i < host:port >** option, for example, **impala-shell -i** *xxx.xxx.xxx.xxx***:21000**. To automatically connect to a specific Impala database, use the **-d <database>** option. For example, if all your Kudu tables are in the **impala_kudu** database, **-d impala_kudu** can use this database. To exit Impala Shell, run the **quit** command.

  ```
  create external table user_info(id string,name string,gender string,age int,addr string)
  partitioned by(year string) row format delimited fields terminated by ' ' lines terminated by '\n'
  stored as textfile location '/hive/user_info';
  ```


NOTE

- **fields terminated** indicates delimiters, for example, spaces.
- **lines terminated** indicates line breaks, for example, **\n**.
- **/hive/user_info** indicates the path of the data file.

c. Import data.

i. Execute the **insert** statement to insert data.

```
insert into user_info partition(year="2018") values ("12005000201", "A", "Male", 19, "City A");
```

ii. Run the **load data** command to import file data.

1) Create a file based on the data in **Table 4-9**. For example, the file name is **txt.log**. Fields are separated by space, and the line feed characters are used as the line breaks.

2) Upload the file to HDFS.

hdfs dfs -put txt.log /tmp

3) Load data to the table.

load data inpath '/tmp/txt.log' into table user_info partition (year='2018');

d. Query the imported data:

```
select * from user_info;
```

e. Delete the user information table:

```
drop table user_info;
```

----End

4.3.8 Using a Kafka Client

Scenario

You can create, query, and delete topics on a cluster client. For details, see **Managing Kafka User Permissions**.

Prerequisites

The client has been installed in a directory, for example, **/opt/client**. The client directory in the following operations is only an example. Change it based on site requirements.

Using the Kafka Client (Versions Earlier Than MRS 3.x)

Step 1 Install the client. For details, see **Installing a Client**.

Step 2 Access the ZooKeeper instance page.

Click the cluster name to go to the cluster details page and choose **Components** > **ZooKeeper** > **Instances**.

> **NOTE**
>
> If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 3 View the IP addresses of the ZooKeeper role instance.

Record any IP address of the ZooKeeper instance.

Step 4 Log in to the node where the client is installed.

Step 5 Run the following command to switch to the client installation directory, for example, **/opt/client/Kafka/kafka/bin**.

cd /opt/client/Kafka/kafka/bin

Step 6 Run the following command to configure environment variables:

source /opt/client/bigdata_env

Step 7 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step.

kinit *Kafka user*

Step 8 Create a topic.

sh kafka-topics.sh --create --topic *Topic name* **--partitions** *Number of partitions occupied by the topic* **--replication-factor** *Number of replicas of the topic* **--zookeeper** *IP address of the node where the ZooKeeper instance resides:clientPort*/**kafka**

Example: **sh kafka-topics.sh --create --topic TopicTest --partitions 3 --replication-factor 3 --zookeeper 10.10.10.100:2181/kafka**

Step 9 Run the following command to view the topic information in the cluster:

sh kafka-topics.sh --list --zookeeper *IP address of the node where the ZooKeeper instance resides:clientPort*/**kafka**

Example: **sh kafka-topics.sh --list --zookeeper 10.10.10.100:2181/kafka**

Step 10 Delete the topic created in **Step 8**.

sh kafka-topics.sh --delete --topic *Topic name* **--zookeeper** *IP address of the node where the ZooKeeper instance resides*:*clientPort*/**kafka**

Example: **sh kafka-topics.sh --delete --topic TopicTest --zookeeper 10.10.10.100:2181/kafka**

Type **y** and press **Enter**.

----End

Using the Kafka Client (MRS 3.x or Later)

Step 1 Install the client. For details, see **Installing a Client**.

Step 2 Access the ZooKeeper instance page.

Log in to FusionInsight Manager. For details, see **Accessing FusionInsight Manager (MRS 3.x or Later)**. Choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Instance**.

Step 3 View the IP addresses of the ZooKeeper role instance.

Record any IP address of the ZooKeeper instance.

Step 4 Log in to the node where the client is installed.

Step 5 Run the following command to switch to the client installation directory, for example, **/opt/client/Kafka/kafka/bin**.

cd /opt/client/Kafka/kafka/bin

Step 6 Run the following command to configure environment variables:

source /opt/client/bigdata_env

Step 7 If Kerberos authentication is enabled for the current cluster, run the following command to authenticate the current user. If Kerberos authentication is disabled for the current cluster, skip this step.

kinit *Kafka user*

Step 8 Log in to FusionInsight Manager, choose **Cluster** > **Name of the desired cluster** > **Services** > **ZooKeeper**, and click the **Configurations** tab and then **All Configurations**. On the displayed page, search for the **clientPort** parameter and record its value.

Step 9 Create a topic.

sh kafka-topics.sh --create --topic *Topic name* **--partitions** *Number of partitions occupied by the topic* **--replication-factor** *Number of replicas of the topic* **--zookeeper** *IP address of the node where the ZooKeeper instance resides:clientPort*/**kafka**

Example: **sh kafka-topics.sh --create --topic TopicTest --partitions 3 --replication-factor 3 --zookeeper 10.10.10.100:2181/kafka**

Step 10 Run the following command to view the topic information in the cluster:

sh kafka-topics.sh --list --zookeeper *IP address of the node where the ZooKeeper instance resides:clientPort*/**kafka**

Example: **sh kafka-topics.sh --list --zookeeper 10.10.10.100:2181/kafka**

Step 11 Delete the topic created in **Step 9**.

sh kafka-topics.sh --delete --topic *Topic name* **--zookeeper** *IP address of the node where the ZooKeeper instance resides*:*clientPort*/**kafka**

Example: **sh kafka-topics.sh --delete --topic TopicTest --zookeeper 10.10.10.100:2181/kafka**

----End

4.3.9 Using a Kudu Client

Kudu is a columnar storage manager developed for the Apache Hadoop platform. Kudu shares the common technical properties of Hadoop ecosystem applications. It is horizontally scalable and supports highly available operations.

Prerequisites

The cluster client has been installed. For example, the client is installed in the **/opt/hadoopclient** directory. The client directory in the following operations is only an example. Change it to the actual installation directory.

Procedure

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

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

Step 2 Run the following command to go to the client installation directory:

cd /opt/hadoopclient

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 Run the Kudu command line tool.

Run the command line tool of the Kudu component to view help information.

kudu -h

The command output is as follows:

```
Usage: kudu <command> [<args>]

<command> can be one of the following:
        cluster   Operate on a Kudu cluster
       diagnose   Diagnostic tools for Kudu servers and clusters
             fs   Operate on a local Kudu filesystem
            hms   Operate on remote Hive Metastores
  local_replica   Operate on local tablet replicas via the local filesystem
         master   Operate on a Kudu Master
            pbc   Operate on PBC (protobuf container) files
           perf   Measure the performance of a Kudu cluster
 remote_replica   Operate on remote tablet replicas on a Kudu Tablet Server
          table   Operate on Kudu tables
         tablet   Operate on remote Kudu tablets
           test   Various test actions
        tserver   Operate on a Kudu Tablet Server
            wal   Operate on WAL (write-ahead log) files
```

NOTE

The Kudu command line tool does not support DDL and DML operations, but provides the refined query function for the **cluster**, **master**, **tserver**, **fs**, and **table** parameters.

Common operations:

- Check the tables in the current cluster.

  kudu table list *KuduMaster instance IP1:7051, KuduMaster instance IP2:7051, KuduMaster instance IP3:7051*

- Query the configurations of the KuduMaster instance of the Kudu service.

  kudu master get_flags *KuduMaster instance IP:7051*

- Query the schema of a table.

  kudu table describe *KuduMaster instance IP1:7051, KuduMaster instance IP2:7051, KuduMaster instance IP3:7051 Table name*

- Delete a table.

  kudu table delete *KuduMaster instance IP1:7051, KuduMaster instance IP2:7051, KuduMaster instance IP3:7051 Table name*

  NOTE

  To obtain the IP address of the KuduMaster instance, choose **Components** > **Kudu** > **Instances** on the cluster details page.

----End

4.3.10 Using the Oozie Client

Scenario

This section describes how to use the Oozie client in an O&M scenario or service scenario.

Prerequisites

- The client has been installed in a directory, for example, **/opt/client**. The client directory in the following operations is only an example. Change it based on site requirements.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login.

Using the Oozie Client

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

Step 2 Run the following command to switch to the client installation directory (change it to the actual installation directory):

cd /opt/client

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 Check the cluster authentication mode.

- If the cluster is in security mode, run the following command to authenticate the user: *exampleUser* indicates the name of the user who submits tasks.

  kinit *exampleUser*

- If the cluster is in normal mode, go to **Step 5**.

Step 5 Perform the following operations to configure Hue:

1. Configure the Spark2x environment (skip this step if the Spark2x task is not involved):

hdfs dfs -put /opt/client/Spark2x/spark/jars/*.jar /user/oozie/share/lib/spark2x/

When the JAR package in the HDFS directory **/user/oozie/share** changes, you need to restart the Oozie service.

2. Upload the Oozie configuration file and JAR package to HDFS.

   hdfs dfs -mkdir /user/*exampleUser*

   hdfs dfs -put -f /opt/client/Oozie/oozie-client-*/examples /user/*exampleUser*/

   NOTE

   - *exampleUser* indicates the name of the user who submits tasks.
   - If the user who submits the task and other files except **job.properties** are not changed, client installation directory **Oozie/oozie-client-*/examples** can be repeatedly used after being uploaded to HDFS.
   - Resolve the JAR file conflict between Spark and Yarn about Jetty.

     hdfs dfs -rm -f /user/oozie/share/lib/spark/jetty-all-9.2.22.v20170606.jar
   - In normal mode, if **Permission denied** is displayed during the upload, run the following commands:

     su - omm

     source /opt/client/bigdata_env

     hdfs dfs -chmod -R 777 /user/oozie

     exit

----End

4.3.11 Using a Storm Client

Scenario

This section describes how to use the Storm client in an O&M scenario or service scenario.

Prerequisites

- You have installed the client. For example, the installation directory is **/opt/hadoopclient**.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users have downloaded the keytab file. A human-machine user must change the password upon the first login. (Not involved in normal mode)

Procedure

Step 1 Prepare the client based on service requirements. Log in to the node where the client is installed.

Log in to the node where the client is installed based on the client location. For details, see **Installing a Client**.

Step 2 Run the following command to go to the client installation directory:

cd /opt/hadoopclient

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If multiple Storm instances are installed, run the following command to load the environment variables of a specific instance when running the Storm command to submit the topology. Otherwise, skip this step. The following command uses the instance Storm-2 as an example.

source Storm-2/component_env

Step 5 Run the following command to perform user authentication (skip this step in normal mode):

kinit *Component service user*

Step 6 Run the following command to perform operations on the client:

For example, run the following command:

- **cql**
- **storm**

NOTE

A Storm client cannot be connected to secure and non-secure ZooKeepers at the same time.

----End

4.3.12 Using a Yarn Client

Scenario

This section guides users to use a Yarn client in an O&M or service scenario.

Prerequisites

- The client has been installed.

  For example, the installation directory is **/opt/client**. The client directory in the following operations is only an example. Change it based on the actual installation directory onsite.
- Service component users have been created by the MRS cluster administrator. In security mode, machine-machine users need to download the keytab file. A human-machine user must change the password upon the first login. In common mode, you do not need to download the keytab file or change the password.

Using the Yarn Client

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

Step 2 Run the following command to go to the client installation directory:

cd /opt/client

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If the cluster is in security mode, run the following command to authenticate the user. In normal mode, user authentication is not required.

kinit *Component service user*

Step 5 Run the Yarn command. The following provides an example:

yarn application -list

----End

Client-related FAQs

1. What Do I Do When the Yarn Client Exits Abnormally and Error Message "java.lang.OutOfMemoryError" Is Displayed After the Yarn Client Command Is Run?

   This problem occurs because the memory required for running the Yarn client exceeds the upper limit (128 MB by default) set on the Yarn client. For clusters of MRS 3.x or later: You can modify **CLIENT_GC_OPTS** in *<Client installation path>***/HDFS/component_env** to change the memory upper limit of the Yarn client. For example, if you want to set the maximum memory to 1 GB, run the following command:

   ```
   export CLIENT_GC_OPTS="-Xmx1G"
   ```

   For clusters earlier than MRS 3.x: You can modify **GC_OPTS_YARN** in *< Client installation path >***/HDFS/component_env** to change the memory upper limit of the Yarn client. For example, if you want to set the maximum memory to 1 GB, run the following command:

   ```
   export GC_OPTS_YARN="-Xmx1G"
   ```

   After the modification, run the following command to make the modification take effect:

   source <*Client installation path*>/**/bigdata_env**

2. How Can I Set the Log Level When the Yarn Client Is Running?

   By default, the logs generated during the running of the Yarn client are printed to the console. The default log level is INFO. To enable the DEBUG log level for fault locating, run the following command to export an environment variable:

   export YARN_ROOT_LOGGER=DEBUG,console

   Then run the Yarn Shell command to print DEBUG logs.

   If you want to print INFO logs again, run the following command:

   export YARN_ROOT_LOGGER=INFO,console

5 Configuring a Cluster with Decoupled Storage and Compute

5.1 MRS Storage-Compute Decoupling

In scenarios that require large storage capacity and elastic compute resources, MRS enables you to store data in OBS and use an MRS cluster for data computing only. In this way, storage and compute are separated.

NOTE

In the big data storage-compute decoupling scenario, the OBS parallel file system must be used to configure a cluster. Using common object buckets will greatly affect the cluster performance.

Perform the following steps to use the storage-compute decoupling function:

1. Configure a cluster with decoupled storage and compute.

   Select one of the following configurations (Using an agency is recommended.):

   - Bind an agency of the ECS type to an MRS cluster to access OBS, preventing the AK/SK from being exposed in the configuration file. For details, see **Configuring a Storage-Compute Decoupled Cluster (Agency)**.
   - Configure the AK/SK in an MRS cluster. The AK/SK will be exposed in the configuration file in plaintext. Exercise caution when performing this operation. For details, see **Configuring a Storage-Compute Decoupled Cluster (AK/SK)**.
   - MRS uses the Guardian component to connect to OBS, providing other components with the capabilities of obtaining temporary authentication credentials and fine-grained permission control for accessing OBS. For details, see **Interconnecting the Guardian Service with OBS**.

NOTE

- Currently, using Guardian components to connect to OBS is supported only in MRS 3.3.0-LTS or later versions. For details about configurations in clusters of other versions, see **Interconnecting MRS with OBS Using an Agency**.
- Job submission based on the Guardian storage and compute decoupling management plane depends on JobGateway instead of Executor.

2. Use the cluster.

   After the required permissions for accessing OBS are obtained, components in the MRS cluster can access the corresponding files through the client.

   For details about how to configure components to access OBS, see the following content:

   - **Interconnecting MRS with OBS Using an Agency**
   - **Interconnecting Components with OBS Using Guardian**

5.2 Interconnecting with OBS Using the Cluster Agency Mechanism

5.2.1 Configuring a Storage-Compute Decoupled Cluster (Agency)

MRS allows you to store data in OBS and use an MRS cluster for data computing only. In this way, storage and compute are separated. You can create an IAM agency, which enables ECS to automatically obtain the temporary AK/SK to access OBS. This prevents the AK/SK from being exposed in the configuration file.

By binding an agency, ECSs or BMSs can manage some of your resources. Determine whether to configure an agency based on the actual service scenario.

MRS provides the following configuration modes for accessing OBS. You can select one of them. The agency mode is recommended.

- Bind an agency of the ECS type to an MRS cluster to access OBS, preventing the AK/SK from being exposed in the configuration file. For details, see the following part in this section.
- Configure the AK/SK in an MRS cluster. The AK/SK will be exposed in the configuration file in plaintext. Exercise caution when performing this operation. For details, see **Configuring a Storage-Compute Decoupled Cluster (AK/SK)**.

  This function is available for components Hadoop, Hive, Spark, Presto, and Flink in clusters of .

(Optional) Step 1: Create an ECS Agency with OBS Access Permissions

NOTE

- MRS presets **MRS_ECS_DEFAULT_AGENCY** in the agency list of IAM so that you can select this agency when creating a cluster. This agency has the **OBSOperateAccess** permission and the **CESFullAccess** (only available for users who have enabled fine-grained policies), **CES Administrator**, and **KMS Administrator** permissions in the region where the cluster is located. Do not modify **MRS_ECS_DEFAULT_AGENCY** on IAM.
- If you want to use the preset agency, skip the step for creating an agency. If you want to use a custom agency, perform the following steps to create an agency. (To create or modify an agency, you must have the Security Administrator permission.) If you need fine-grained permission control on specified paths in the OBS file system, see **Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS** and create custom role policies.

1. Log in to the Huawei Cloud management console.
2. Choose **Service List** > **Management & Governance** > **Identity and Access Management**.
3. Choose **Agencies**. On the displayed page, click **Create Agency**.
4. Enter an agency name, for example, **mrs_ecs_obs**.
5. Set **Agency Type** to **Cloud service** and select **ECS BMS** to authorize ECS or BMS to invoke OBS. See **Figure 5-1**.
6. Set **Validity Period** to **Unlimited** and click **Next**.

Figure 5-1 Creating an agency

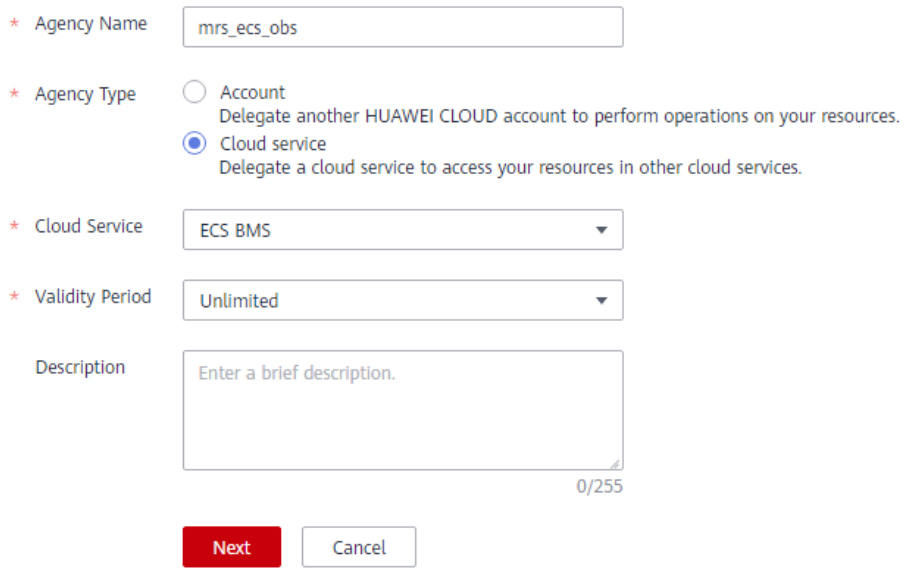

7. On the displayed page, search for the **OBS OperateAccess** and select it, as shown in .

NOTE

If KMS encryption is configured for an OBS bucket, the **KMS Administrator** policy must be selected.

Figure 5-2 Configuring permissions

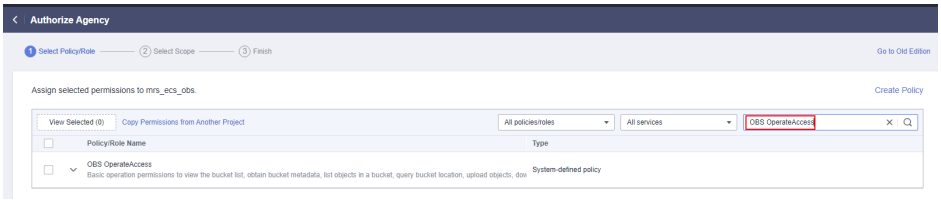

8. Click **Next**. On the page that is displayed, select the desired scope for the permissions you selected. By default, **All resources** is selected. Click **Show More**, select **Global resources**, and click **OK**.
9. In the dialog box that is displayed, click **OK** to start authorization. After the message "**Authorization successful.**" is displayed, click **Finish**. The agency is successfully created.

Step 2: Create a Cluster with Storage and Compute Separated

You can configure an agency when creating a cluster or bind an agency to an existing cluster to separate storage and compute. This section uses a cluster with Kerberos authentication enabled as an example.

Configuring an agency when creating a cluster:

1. Go to the **Buy Cluster** page.
2. Click the **Custom Config** tab.

Figure 5-3 Custom purchase of a cluster

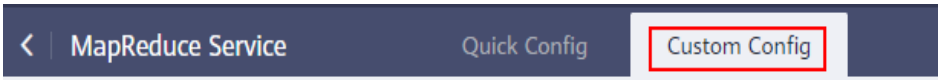

3. On the **Custom Config** tab page, set software parameters.
   - **Region**: Select a region as required.
   - **Cluster Name**: You can use the default name. However, you are advised to include a project name abbreviation or date for consolidated memory and easy distinguishing.
   - Cluster Version: Select a cluster version.
   - **Cluster Type**: Select **Analysis cluster** or **Hybrid cluster** and select all components.
   - **Metadata**: Select **Local**.
4. Click **Next** and set hardware parameters.
   - **Billing Mode**: Select **Pay-per-use**.
   - **AZ**: Use the default value.
   - **VPC**: Use the default value.
   - **Subnet**: Use the default value.
   - **Security Group**: Use the default value.
   - **EIP**: Use the default value.
   - **Enterprise Project**: Use the default value.
   - **CPU Architecture**: Use the default value. This parameter is supported in the versions earlier than MRS 3.x .

- **Cluster Node**: Select the number of cluster nodes and node specifications based on site requirements.

5. Click **Next** and set related parameters.
   - **Kerberos Authentication**: This function is enabled by default. You can enable or disable it.
   - **Username**: The default username is **admin**, which is used to log in to MRS Manager.
   - **Password**: Set a password for user **admin**.
   - **Confirm Password**: Enter the password of user **admin** again.
   - **Login Mode**: Select a method for logging in to ECSs. In this example, select **Password**.
   - **Username**: The default username is **root**, which is used to remotely log in to ECSs.
   - **Password**: Set a password for user **root**.
   - **Confirm Password**: Enter the password of user **root** again.
6. In this example, configure an agency and leave other parameters blank. For details about how to configure other parameters, see **Advanced Options**.

   Agency: Select the agency created in **(Optional) Step 1: Create an ECS Agency with OBS Access Permissions** or **MRS_ECS_DEFAULT_AGENCY** preset in IAM.

   Figure 5-4 Configuring an agency

   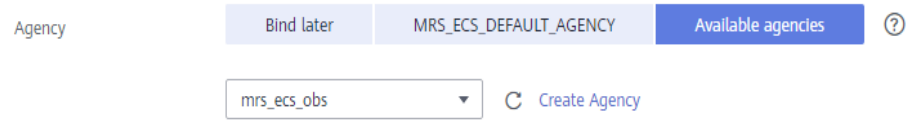

7. Select the check box for secure communications. For details, see **Communication Security Authorization**.
8. Click **Buy Now** and wait until the cluster is created.

   If Kerberos authentication is enabled for a cluster, check whether Kerberos authentication is required. If yes, click **Continue**. If no, click **Back** to disable Kerberos authentication and then create a cluster.

Configuring an agency for an existing cluster:

1. Log in to the MRS management console. In the left navigation pane, choose **Clusters** > **Active Clusters**.
2. Click the name of the cluster to enter its details page.
3. On the **Dashboard** page, click **Synchronize** on the right of **IAM User Sync** to synchronize IAM users.
4. On the **Dashboard** tab page, click **Manage Agency** on the right side of **Agency** to select an agency and click **OK** to bind it. Alternatively, click **Create Agency** to go to the IAM console to create an agency and select it.

Figure 5-5 Binding an agency

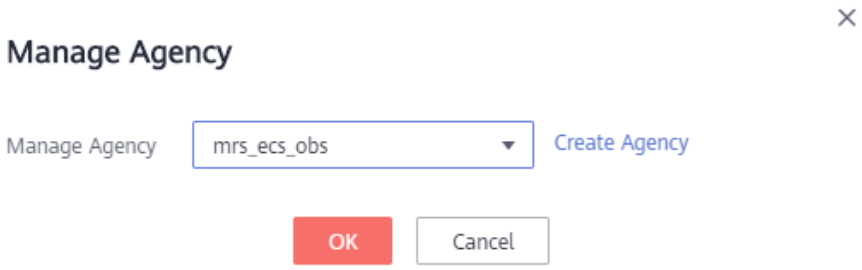

Step 3: Create an OBS File System for Storing Data

NOTE

In storage-compute decoupled scenarios, the OBS parallel file system must be used to store data. The cluster performance will be significantly affected if common object buckets are used.

1. Log in to the OBS Console.
2. Choose **Parallel File System** > **Create Parallel File System**.
3. Enter the file system name, for example, **mrs-word001**.

   Set other parameters as required.

Figure 5-6 Creating an OBS parallel file system

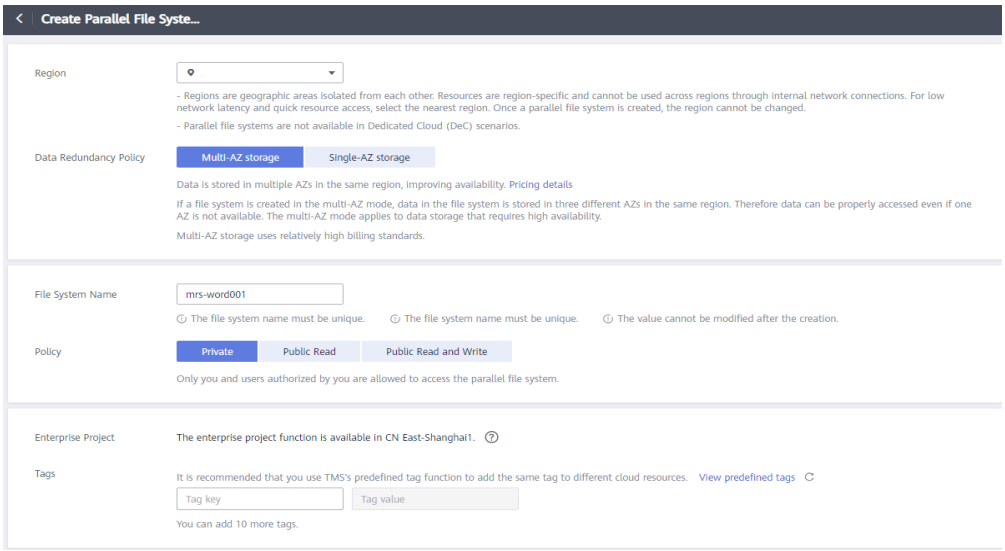

4. Click **Create Now**.
5. In the parallel file system list on the OBS console, click the file system name to go to the details page.
6. In the navigation pane, choose **Files** and create the **program** and **input** folders.
   - **program**: Upload the program package to this folder.
   - **input**: Upload the input data to this folder.

Step 4: Access the OBS File System

1. Log in to a Master node as user **root**. For details, see **Logging In to an ECS**.
2. Run the following command to set the environment variables:

   For versions earlier than MRS 3.x, run the **source /opt/client/bigdata_env** command.

   For MRS 3.x or later, run the **source /opt/Bigdata/client/bigdata_env** command.
3. Verify that Hadoop can access OBS.

   a. View the list of files in the file system **mrs-word001**.

   hadoop fs -ls obs://mrs-word001/

   b. Check whether the file list is returned. If it is returned, OBS access is successful.

   Figure 5-7 Returned file list

   ```
   Found 2 items
   drwxrwxrwx   - root root          0 2019-12-21 11:04 obs://mrs-word001/input
   drwxrwxrwx   - root root          0 2019-12-21 11:04 obs://mrs-word001/program
   ```
4. Verify that Hive can access OBS.

   a. If Kerberos authentication has been enabled for the cluster, run the following command to authenticate the current user. The current user must have a permission to create Hive tables. For details about how to configure a role with a permission to create Hive tables, see **Creating a Role**. For details about how to create a user and bind a role to the user, see **Creating a User**. If Kerberos authentication is disabled for the current cluster, skip this step.

   kinit *MRS cluster user*

   Example: **kinit hiveuser**

   b. Run the client command of the Hive component.

   beeline

   c. Access the OBS directory in the beeline. For example, run the following command to create a Hive table and specify that data is stored in the **test_obs** directory of the file system **mrs-word001**:

   create table test_obs(a int, b string) row format delimited fields terminated by ',' stored as textfile location "obs://mrs-word001/test_obs";

   d. Run the following command to query all tables. If table **test_obs** is displayed in the command output, OBS access is successful.

   show tables;

   Figure 5-8 Returned table name

   ```
   +-----------+
   | tab_name  |
   +-----------+
   | test_obs  |
   +-----------+
   1 row selected (0.352 seconds)
   ```


e. Press **Ctrl+C** to exit the Hive beeline.

5. Verify that Spark can access OBS.

a. Run the client command of the Spark component.

spark-beeline

b. Access OBS in spark-beeline. For example, create table **test** in the **obs://mrs-word001/table/** directory.

create table test(id int) location 'obs://mrs-word001/table/';

c. Run the following command to query all tables. If table **test** is displayed in the command output, OBS access is successful.

show tables;

Figure 5-9 Returned table name

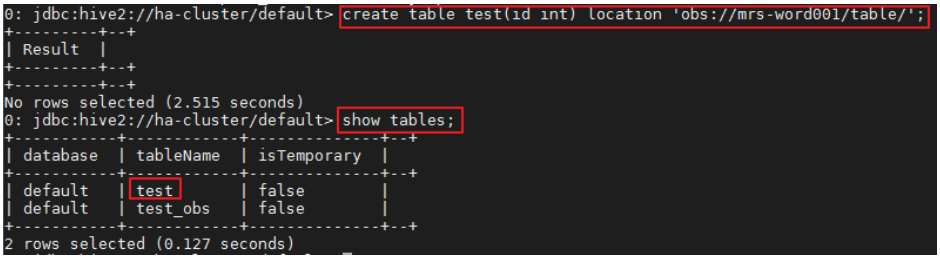

d. Press **Ctrl+C** to exit the Spark beeline.

6. Verify that Presto can access OBS.

– For normal clusters with Kerberos authentication disabled

i. Run the following command to connect to the client:

presto_cli.sh

ii. On the Presto client, run the following statement to create a schema and set **location** to an OBS path:

CREATE SCHEMA hive.demo WITH (location = 'obs://mrs-word001/presto-demo002/');

iii. Create a table in the schema. The table data is stored in the OBS file system. The following is an example.

CREATE TABLE hive.demo.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;

Figure 5-10 Return result

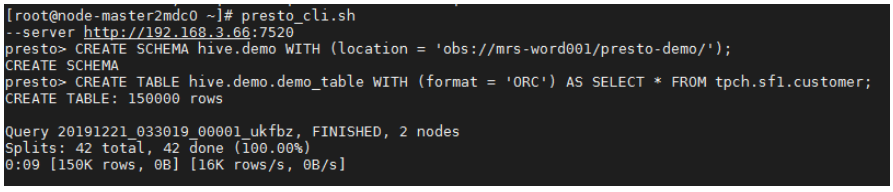

iv. Run **exit** to exit the client.

– For security clusters with Kerberos authentication enabled

i. Log in to MRS Manager and create a role with the Hive Admin Privilege permissions, for example, **prestorole**. For details about how to create a role, see **Creating a Role**.

ii. Create a user that belongs to the Presto and Hive groups and bind the role created in **6.i** to the user, for example, **presto001**. For details about how to create a user, see **Creating a User**.

iii. Authenticate the current user.

kinit presto001

iv. Download the user credential.

1) For MRS 3.x earlier, on MRS Manager, choose **System** > **Manage User**. In the row of the new user, choose **More** > **Download Authentication Credential**.

Figure 5-11 Downloading the Presto user authentication credential

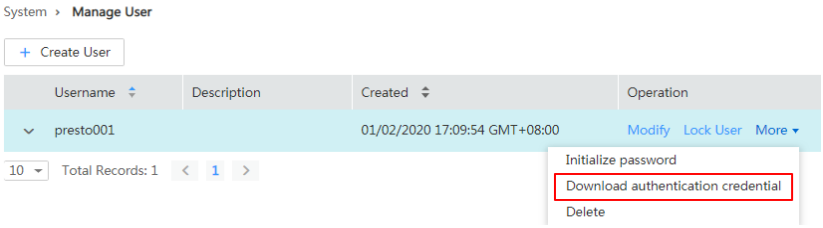

2) On FusionInsight Manager for MRS 3.x or later,, choose **System > Permission > User**. In the row that contains the newly added user, click **More > Download Authentication Credential**.

v. Decompress the downloaded user credential file, and save the obtained **krb5.conf** and **user.keytab** files to the client directory, for example, **/opt/Bigdata/client/Presto/**.

vi. Run the following command to obtain a user principal:

klist -kt /opt/Bigdata/client/Presto/user.keytab

vii. For clusters with Kerberos authentication enabled, run the following command to connect to the Presto Server of the cluster:

presto_cli.sh --krb5-config-path {krb5.conf file path} --krb5-principal {user principal} --krb5-keytab-path {user.keytab file path} --user {presto username}

- **krb5.conf** file path: Replace it with the file path set in **6.v**, for example, **/opt/Bigdata/client/Presto/krb5.conf**.
- **user.keytab** file path: Replace it with the file path set in **6.v**, for example, **/opt/Bigdata/client/Presto/user.keytab**.
- **user principal**: Replace it with the result returned in **6.vi**.
- **presto username**: Replace it with the name of the user created in **6.ii**, for example, **presto001**.

Example: presto_cli.sh --krb5-config-path /opt/Bigdata/client/Presto/krb5.conf --krb5-principal prest001@xxx_xxx_xxx_xxx.COM --krb5-keytab-path /opt/Bigdata/client/Presto/user.keytab --user presto001

viii. On the Presto client, run the following statement to create a schema and set **location** to an OBS path:

CREATE SCHEMA hive.demo01 WITH (location = 'obs://mrs-word001/presto-demo002/');

ix. Create a table in the schema. The table data is stored in the OBS file system. The following is an example.

CREATE TABLE hive.demo01.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;

Figure 5-12 Return result

x. Run **exit** to exit the client.

7. Verify that Flink can access OBS.
   a. On the **Dashboard** page, click **Synchronize** on the right of **IAM User Sync** to synchronize IAM users.
   b. After user synchronization is complete, choose **Jobs** > **Create** on the cluster details page to create a Flink job. In **Parameters**, enter parameters in **--input <Job input path> --output <Job output path>** format. You can click **OBS** to select a job input path, and enter a job output path that does not exist, for example, **obs://mrs-word001/output/**.
   c. On OBS Console, go to the output path specified during job creation. If the output directory is automatically created and contains the job execution results, OBS access is successful.

Figure 5-13 Flink job execution result

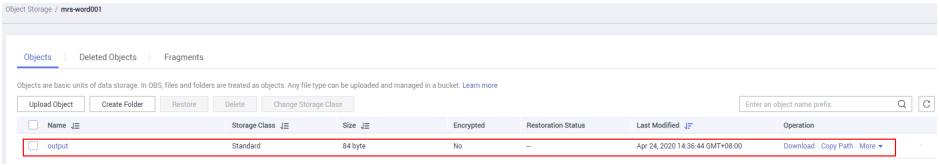

Step 5: Configure a Lifecycle Rule

In MRS 3.2.0-LTS.1 and later versions, components prevent mis-deletion by default. That is, file data deleted by component users is not directly deleted but stored in the recycle bin directory in the OBS file system.

To save OBS space, you need to enable periodical deletion of file data from the OBS recycle bin by referring to **Configuring the Policy for Clearing Component Data in the Recycle Bin**.

Reference

For details about how to control permissions to access OBS, see **Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS**.

5.2.2 Configuring a Storage-Compute Decoupled Cluster (AK/SK)

In MRS 1.9.2 or later, OBS can be interconnected with MRS using **obs://**. Currently, Hadoop, Hive, Spark, Presto, and Flink are supported. HBase cannot use **obs://** to interconnect with OBS.

MRS provides the following configuration modes for accessing OBS. You can select one of them. The agency mode is recommended.

- Bind an agency of the ECS type to an MRS cluster to access OBS, preventing the AK/SK from being exposed in the configuration file. For details, see **Configuring a Storage-Compute Decoupled Cluster (Agency)**.
- Configure the AK/SK in an MRS cluster. The AK/SK will be exposed in the configuration file in plaintext. Exercise caution when performing this operation. For details, see the following part in this section.

NOTICE

- To improve data write performance, log in to the Manager and choose **Cluster** > **Services** > *Name of the service to be modified* > **Configurations**. Change the value of **fs.obs.buffer.dir** to the data disk directory.
- In the storage-compute decoupled scenario, the OBS parallel file system must be used to configure a cluster. For details, see **Parallel File System**. Using common object buckets will greatly affect the cluster performance.
- **In MRS 3.2.0-LTS.1 and later versions, components prevent mis-deletion by default. That is, file data deleted by component users is not directly deleted but stored in the recycle bin directory in the OBS file system.**

  To save OBS space, you need to enable periodical deletion of file data from the OBS recycle bin by referring to **Configuring the Policy for Clearing Component Data in the Recycle Bin**.
- Configuration files containing authentication passwords pose security risks. Delete such files after configuration or store them securely.
- Commands carrying authentication passwords pose security risks. Disable historical command recording before running such commands to prevent information leakage.

Using Hadoop to Access OBS

- Add the following content to the **core-site.xml** file in the *Client installation directory*/**HDFS/hadoop/etc/hadoop** directory on the HDFS client.

```
<property>
   <name>fs.obs.access.key</name>
   <value>ak</value>
</property>
<property>
   <name>fs.obs.secret.key</name>
   <value>sk</value>
</property>
<property>
   <name>fs.obs.endpoint</name>
   <value>obs endpoint</value>
</property>
```


If you use commands that need to submit jobs to Yarn, such as **distcp**, you need to add the preceding content to the **core-site.xml** file in the Yarn directory (**$client_home/Yarn/config**) on the MRS client.

NOTICE

AK and SK will be displayed as plaintext in the configuration file. Exercise caution when setting AK and SK in the file.

After the configuration is added, you can directly access data on OBS without manually adding the AK/SK and endpoint. For example, run the following command to view the file list of the **test_obs_orc** directory in the **obs-test** file system:

hadoop fs -ls "obs://obs-test/test_obs_orc"

- Add AK/SK and endpoint to the command line to access data on OBS.

  hadoop fs -Dfs.obs.endpoint=xxx -Dfs.obs.access.key=xx -Dfs.obs.secret.key=xx -ls "obs://obs-test/ test_obs_orc"

Using Hive to Access OBS

Step 1 Log in to the service configuration page.

- For versions earlier than MRS 3.x, log in to the cluster details page and choose **Components** > **Hive** > **Service Configuration**.
- For MRS 3.*x* or later, log in to FusionInsight Manager. For details, see **Accessing FusionInsight Manager (MRS 3.x or Later)**. Choose **Cluster** > **Services** > **Hive** > **Configurations**.

Step 2 In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.

Step 3 Configure the AK and SK of OBS.

- For versions earlier than MRS 3.*x*, click **Hive**, select **Customization**, and add the following two configurations to **core.site.customized.configs**: **Name**: **fs.obs.access.key**; **Value**: AK for accessing OBS; **Name**: **fs.obs.secret.key**; **Value**: SK for accessing OBS.

  Figure 5-14 Setting the AK/SK for accessing OBS

  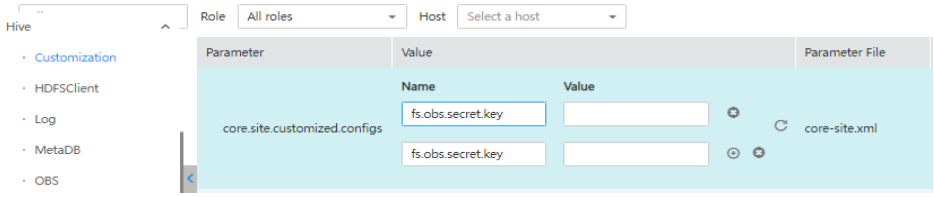

- For MRS 3.*x* or later, click **Hive(Service)**, select **Customization**, and add the following two configurations to **core.site.customized.configs**: **Name**: **fs.obs.access.key**; **Value**: AK for accessing OBS; **Name**: **fs.obs.secret.key**; **Value**: SK for accessing OBS.

Figure 5-15 Configuring the AK/SK for accessing OBS

Step 4 Save the configurations and restart Hive.

Step 5 Access the OBS directory in Beeline. For example, run the following command to create a Hive table and specify that data is stored in the **test_obs** directory in the **test-bucket** file system:

create table test_obs(a int, b string) row format delimited fields terminated by "," stored as textfile location "obs://test-bucket/test_obs";

----End

Using Spark to Access OBS

NOTE

- SparkSQL depends on Hive. Therefore, when configuring OBS on Spark, you need to modify the OBS configuration used in **Using Hive to Access OBS**.
- In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

- spark-beeline and spark-sql

  You can use spark-beeline or spark-sql to log in to the Spark client and run the following commands to configure AK and SK information for accessing OBS:

  set fs.obs.access.key=*AK*

  set fs.obs.secret.key=*SK*

  set fs.obs.endpoint=*OBS Endpoint*

- spark-beeline

  The spark-beeline can access OBS by configuring service parameters on Manager. The procedure is as follows:

  a. Log in to the service configuration page.

     - For versions earlier than MRS 3.x, log in to the cluster details page and choose **Components** > **Spark** > **Service Configuration**.
     - For MRS 3.*x* or later, log in to FusionInsight Manager. For details, see **Accessing FusionInsight Manager (MRS 3.x or Later)**. Choose **Cluster** > **Services** > **Spark2x** > **Configurations**.

  b. In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.

  c. Choose **JDBCServer** > **OBS**, and set values for **fs.obs.access.key** and **fs.obs.secret.key**.

     If the preceding two parameters cannot be found in the current cluster, choose **JDBCServer** > **Customization** in the navigation tree on the left and add the two parameters to the customized parameter **spark.core-site.customized.configs**.

Figure 5-16 Parameters for adding an OBS

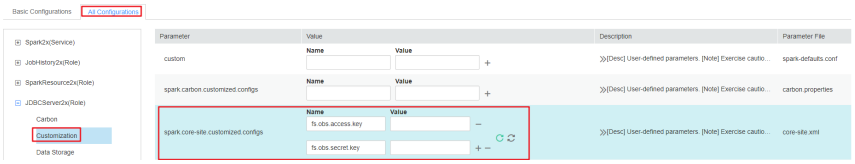

d. Save the configurations and restart Spark.

e. Access OBS in **spark-beeline**. For example, access the **obs://obs-demo-input/table/** directory.

create table test(id int) location 'obs://obs-demo-input/table/';

- spark-sql and spark-submit

  Both spark-sql and spark-submit can access OBS if you add the following content to the **core-site.xml** configuration file in the *Client installation directory*/**Spark/spark/conf** directory:

```
<property>
    <name>fs.obs.access.key</name>
    <value>ak</value>
</property>
<property>
    <name>fs.obs.secret.key</name>
    <value>sk</value>
</property>
<property>
    <name>fs.obs.endpoint</name>
    <value>obs endpoint</value>
</property>
```

Using Presto to Access OBS

Step 1 Go to the cluster details page and choose **Components** > **Presto** > **Service Configuration**.

Step 2 In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.

Step 3 Search for and configure the following parameters:

- Set **fs.obs.access.key** to **AK**.
- Set **fs.obs.secret.key** to **SK**.

If the preceding two parameters cannot be found in the current cluster, choose **Presto > Hive** in the navigation tree on the left and add the two parameters to the customized parameter **core.site.customized.configs**.

Step 4 Save the configurations and restart Presto.

Step 5 Choose **Components** > **Hive** > **Service Configuration**.

Step 6 In the configuration type drop-down box, switch **Basic Configurations** to **All Configurations**.

Step 7 Search for and configure the following parameters:

- Set **fs.obs.access.key** to **AK**.
- Set **fs.obs.secret.key** to **SK**.

Step 8 Save the configurations and restart Hive.

Step 9 On the Presto client, run the following statement to create a schema and set **location** to an OBS path:

presto_cli.sh

CREATE SCHEMA hive.demo WITH (location = 'obs://obs-demo/presto-demo/');

Step 10 Create a table in the schema. The table data is stored in the OBS file system. The following is an example.

USE hive.demo;

CREATE TABLE *Table name* **(id int);**

INSERT INTO *Table name* **VALUES (*2*);** In this command, **2** is only used an example. Replace it with the real value.

CREATE TABLE hive.demo.demo_table WITH (format = 'ORC') AS SELECT * FROM tpch.sf1.customer;

----End

Using Flink to Access OBS

Add the following configuration to the Flink configuration file of the MRS client in ***Client installation path*/Flink/flink/conf/flink-conf.yaml**:

```
fs.obs.access.key: ak
fs.obs.secret.key: sk
fs.obs.endpoint: OBS Endpoint
```

NOTICE

AK and SK will be displayed as plaintext in the configuration file. Exercise caution when setting AK and SK in the file.

After the configuration is added, you can directly access data on OBS without manually adding the AK/SK and endpoint.

5.2.3 Configuring the Policy for Clearing Component Data in the Recycle Bin

Scenario

By default, components in an MRS 3.2.0-LTS.1 or later cluster support prevention against accidental data deletion. Native HDFS garbage collection can be used in the Hadoop big data systems that use OBS.

The file data deleted by a component user is not directly deleted, but is stored in the recycle bin of the OBS file system instead. This section describes how to set a lifecycle rule for the recycle bin directory to periodically clear related data.

CAUTION

- **For clusters that use decoupled storage and compute, configure a lifecycle policy for the related directories by referring to this chapter. Otherwise, the storage space may be used up and storage fees may increase.**
- The recycle bin directory is created per user. When a user is created in the MRS cluster and the user has the permission to delete component data, you need to configure the recycle bin clearing rule for this new user.
- For HBase components that use decoupled storage and compute in MRS 3.1.2 or later versions, refer to this topic to set a policy for clearing component data in the recycle bin.

You need to configure lifecycle policies for the recycle bin directories of preset users in the MRS cluster and the recycle bin directories of new users who need accidental deletion prevention. If a low privileged agency is used or only the permission for MRS users to access OBS file system directories is configured by referring to **Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS**, you will need the operation permission for the recycle bin directory.

Table 5-1 Directories for which a lifecycle policy needs to be configured

Cluster Version	Directory Type	Component	Directory	How to Create
Versions earlier than MRS 3.3.0-LTS	Recycle bin directories that must be configured by default for each component in an MRS cluster	Hive	• user/omm/.Trash • user/hive/.Trash	If the **.Trash** folder does not exist, create it on the cluster client as user **omm**. Run the following command: **hdfs dfs -mkdir -p obs://***Name of the OBS parallel file system where the table is stored***/***Folder path*
		Spark	• user/omm/.Trash • user/root/.Trash • user/spark2x/.Trash	
		HetuEngine	• user/omm/.Trash • user/hetuserver/.Trash	
		HBase	• user/hbase/.Trash • user/omm/.Trash	

Cluster Version	Directory Type	Component	Directory	How to Create
	Recycle bin directories of users who need accidental deletion prevention	Hive/ Spark/ HetuEngine	user/*<New service user>*/.Trash	
MRS 3.3.0-LTS or later	Default recycle bin directories configured for each component in an MRS cluster	Hive/ Spark/ HetuEngine	/user/.Trash	

For example, if a new user in the cluster has the following permissions, you need to create a recycle bin directory clearing rule for the user in the parallel file system:

- Permissions to delete the HDFS files
- **DROP**, **INSERT OVERWRITE**, and **TRUNCATE** permissions on Hive tables
- **DROP**, **TRUNCATE**, **DELETE**, **INSERT OVERWRITE**, and **LOAD OVERWRITE** permissions on HetuEngine

Configuring the Lifecycle Rule of an OBS Directory

Step 1 Log in to the OBS console.

Step 2 Click **Parallel File Systems** and click the name of the file system used by the current MRS cluster.

Step 3 In the navigation pane on the left, choose **Basic Configurations** > **Lifecycle Rules**. Click **Create** to create a lifecycle rule for a specified directory. For details about the parameters, see **Configuring a Lifecycle Rule**.

Table 5-2 Parameters for creating a lifecycle rule

Name	Description	Example Value
Status	Whether to enable the lifecycle rule.	Enable
Rule Name	Rule name that identifies different lifecycle configurations.	rule-test

Name	Description	Example Value
Prefix	Prefix of the objects to which the lifecycle rule applies. Objects that have the specified prefix will be managed by the lifecycle rule. The prefix cannot start with a slash (/), have consecutive slashes (/), or contain the following special characters: **\:*?"<>|** If this parameter is not specified, the rule will take effect for the entire file system. **NOTE** To prevent other service data from being deleted by mistake, you are not advised to use the lifecycle rule configured for the entire file system or high-level directories. Generally, the recycle bin directory of MRS components is in the following format. If the folder does not exist, create it. user/*<Username>*/.Trash	user/omm/.Trash
Delete Files After (Days)	The object within the rule configuration scope expires and is automatically deleted by OBS if the number of days since its last update reaches this parameter value.	30 days

Step 4 Click **OK** to complete the lifecycle rule configuration.

You can click **Edit** in the **Operation** column of a lifecycle rule to edit it. You can also click **Disable** or **Enable** to disable or enable it.

Step 5 Repeat the preceding steps to create recycle bin directory clearing rules for all users who have the data deletion permission in the current MRS cluster one by one until all recycle bin directories in the OBS file system are configured.

----End

5.2.4 Interconnecting MRS with OBS Using an Agency

5.2.4.1 Interconnecting Flink with OBS

Before performing the following operations, ensure that you have configured a storage-compute decoupled cluster by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)**.

Step 1 Log in to the Flink client installation node as the client installation user.

Step 2 Run the following command to initialize environment variables:

source *Client installation directory*/**bigdata_env**

Step 3 Configure the Flink client.

Step 4 Start a session.

- Normal cluster (Kerberos authentication disabled)

  yarn-session.sh -nm "*session-name***" -d**

- Security cluster (Kerberos authentication enabled)
  - If the **flink.keystore** and **flink.truststore** file paths are relative paths:

    Run the following command in the directory at the same level as **ssl** to start the session. **ssl/** is a relative path.

    cd /opt/hadoopclient/Flink/flink/conf/

    yarn-session.sh -t ssl/ -nm "*session-name***" -d**

    ```
    ...
    Cluster started: Yarn cluster with application id application_1624937999496_0017
    JobManager Web Interface: http://192.168.1.150:32261
    ```

  - If the **flink.keystore** and **flink.truststore** file paths are absolute paths:

    Run the following command to start a session:

    cd /opt/hadoopclient/Flink/flink/conf/

    yarn-session.sh -nm "*session-name***" -d**

Step 5 For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

kinit *Username*

Step 6 Explicitly add the OBS file system to be accessed in the Flink command line.

echo -e 'test' >/tmp/test

hdfs dfs -mkdir -p obs://*Parallel file system name***/tmp/flinkjob**

hdfs dfs -put /tmp/test/ obs://*Parallel file system name***/tmp/flinkjob/**

flink run Client installation directory**/Flink/flink/examples/batch/WordCount.jar -input obs://***Parallel file system name***/tmp/flinkjob/test -output obs://***Parallel file system name***/tmp/flinkjob/output**

----End

NOTE

Flink jobs are running on Yarn. Before configuring Flink to interconnect with the OBS file system, ensure that the interconnection between Yarn and the OBS file system is normal.

5.2.4.2 Interconnecting Flume with OBS

This section applies to MRS 3.x or later.

Before performing the following operations, ensure that you have configured a storage-compute decoupled cluster by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)** or **Configuring a Storage-Compute Decoupled Cluster (AK/SK)**.

Step 1 Configure an agency.

1. Log in to the MRS console. In the navigation pane on the left, choose **Clusters** > **Active Clusters**.

2. Click the name of a cluster to go to the cluster details page.
3. On the **Dashboard** page, click **Synchronize** on the right of **IAM User Sync** to synchronize IAM users.
4. Click **Manage Agency** on the right of **Agency**, select the target agency, and click **OK**.

Step 2 Create an OBS file system for storing data.

1. Log in to the OBS console.
2. In the navigation pane on the left, choose **Parallel File Systems**. On the displayed page, click **Create Parallel File System**.
3. Enter the file system name, for example, **esdk-c-test-pfs1**, and set other parameters as required. Click **Create Now**.
4. In the parallel file system list on the OBS console, click the created file system name to go to its details page.
5. In the navigation pane on the left, choose **Files** and click **Create Folder** to create the **testFlumeOutput** folder.

Step 3 Prepare the **properties.properties** file and upload it to the **/opt/flumeInput** directory.

1. Prepare the **properties.properties** file on the local host. Its content is as follows:

```
# source
server.sources = r1
# channels
server.channels = c1
# sink
server.sinks = obs_sink
# ----- define net source -----
server.sources.r1.type = seq
server.sources.r1.spooldir = /opt/flumeInput
# ---- define OBS sink ----
server.sinks.obs_sink.type = hdfs
server.sinks.obs_sink.hdfs.path = obs://esdk-c-test-pfs1/testFlumeOutput
server.sinks.obs_sink.hdfs.filePrefix = %[localhost]
server.sinks.obs_sink.hdfs.useLocalTimeStamp = true
# set file size to trigger roll
server.sinks.obs_sink.hdfs.rollSize = 0
server.sinks.obs_sink.hdfs.rollCount = 0
server.sinks.obs_sink.hdfs.rollInterval = 5
#server.sinks.obs_sink.hdfs.threadsPoolSize = 30
server.sinks.obs_sink.hdfs.fileType = DataStream
server.sinks.obs_sink.hdfs.writeFormat = Text
server.sinks.obs_sink.hdfs.fileCloseByEndEvent = false

# define channel
server.channels.c1.type = memory
server.channels.c1.capacity = 1000
# transaction size
server.channels.c1.transactionCapacity = 1000
server.channels.c1.byteCapacity = 800000
server.channels.c1.byteCapacityBufferPercentage = 20
server.channels.c1.keep-alive = 60
server.sources.r1.channels = c1
server.sinks.obs_sink.channel = c1
```

NOTE

The value of **server.sinks.obs_sink.hdfs.path** is the OBS file system created in **Step 2**.

2. Log in to the node where the Flume client is installed as user **root**.

3. Create the **/opt/flumeInput** directory and create a customized **.txt** file in this directory.
4. Log in to FusionInsight Manager.
5. Choose **Cluster** > *Name of the target cluster* > **Services** > **Flume**. On the displayed page, click **Configurations** and then **Upload File** in the **Value** column corresponding to the **flume.config.file** parameter, upload the **properties.properties** file prepared in **Step 3.1**, and click **Save**.

Step 4 View the result in the OBS system.

1. Log in to the OBS console.
2. Click **Parallel File Systems** and go to the folder created in **Step 2** to view the result.

----End

5.2.4.3 Interconnecting HDFS with OBS

Before performing the following operations, ensure that you have configured a storage-compute decoupled cluster by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)** or **Configuring a Storage-Compute Decoupled Cluster (AK/SK)**.

Step 1 Log in to the node on which the HDFS client is installed as a client installation user.

Step 2 Run the following command to switch to the client installation directory.

cd *Client installation directory*

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If the cluster is in security mode, authenticate the user. In normal mode, skip user authentication.

kinit *Component service user*

Step 5 Explicitly add the OBS file system to be accessed in the HDFS command line.

For example:

- Run the following command to access the OBS file system:

  hdfs dfs -ls obs://*OBS parallel file system name***/***Path*
- Run the following command to upload the **/opt/test.txt** file from the client node to the OBS file system path:

  hdfs dfs -put /opt/test.txt obs://*OBS parallel file system name/Path*

----End

NOTE

If a large number of logs are printed in the OBS file system, the read and write performance may be affected. You can adjust the log level of the OBS client as follows:

cd *Client installation directory*/**HDFS/hadoop/etc/hadoop**

vi log4j.properties

Add the OBS log level configuration to the file as follows:

log4j.logger.org.apache.hadoop.fs.obs=WARN

log4j.logger.com.obs=WARN

Figure 5-17 Adding an OBS log level

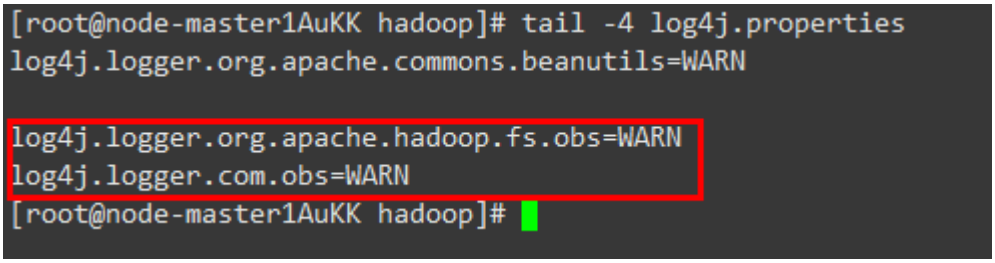

5.2.4.4 Interconnecting Hive with OBS

Before performing the following operations, ensure that you have configured a storage-compute decoupled cluster by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)** or **Configuring a Storage-Compute Decoupled Cluster (AK/SK)**.

Setting the Location to an OBS Path When Creating a Table

Step 1 Log in to the client installation node as the client installation user.

Step 2 Run the following command to initialize environment variables:

source *Client installation directory*/**bigdata_env**

Step 3 For a security cluster, run the following command to perform user authentication (the user must have the permission to perform Hive operations). If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

kinit *User performing Hive operations*

Step 4 Log in to FusionInsight Manager of a cluster earlier than MRS 3.2.0, choose **Cluster** > **Services** > **Hive**, and click **Configurations** > **All Configurations**.

In the navigation pane on the left, choose **Hive** > **Customization**. In custom configuration items, add **dfs.namenode.acls.enabled** to **hdfs.site.customized.configs** and set its value to **false**.

Figure 5-18 Adding custom parameters

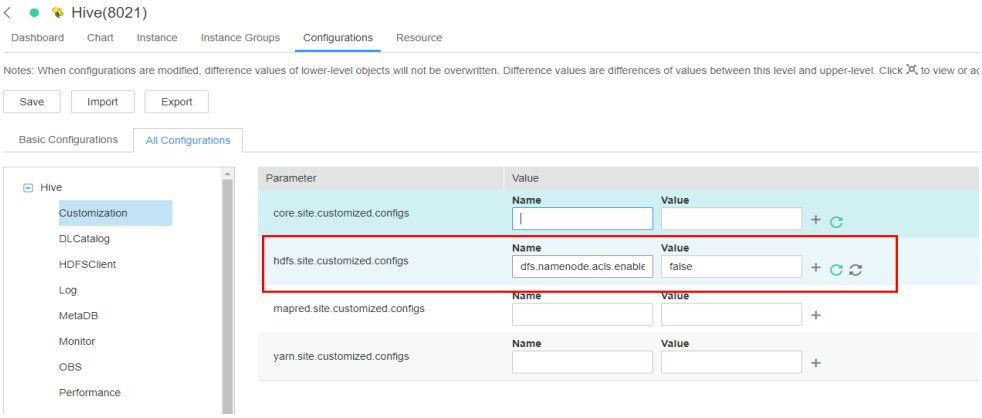

Step 5 Click **Save** to save the configuration for versions earlier than MRS 3.2.0. On the **Dashboard** page, click **More** and select **Restart Service**. Enter the password of the current user, click **OK**, and select **Restart upper-layer services**. Click **OK** to restart Hive.

Step 6 Log in to the beeline client and set **Location** to the OBS file system path when creating a table.

beeline

For example, run the following command to create the table **test** in **obs://***OBS parallel file system name***/user/hive/warehouse/***Database name***/***Table name*:

create table test(name string) location "obs://*OBS parallel file system name***/user/hive/warehouse/***Database name***/***Table name***";**

NOTE

You need to add the component operator to the URL policy in the Ranger policy. Set the URL to the complete path of the object on OBS. Select the Read and Write permissions.

----End

Interconnecting Hive with OBS Through MetaStore

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Hive** > **Configurations** > **All Configurations**.

- For versions earlier than MRS 3.2.0:
  - In the navigation pane on the left, choose **MetaStore (role)** > **Customization**. Add the configuration item **hive.metastore.warehouse.dir** to the custom parameter **hive.metastore.customized.configs** and set the value to an OBS path. For example, set it to **obs://hivetest/user/hive/warehouse/**, where **hivetest** is the name of the OBS parallel file system.

Figure 5-19 Configuring **hive.metastore.warehouse.dir**

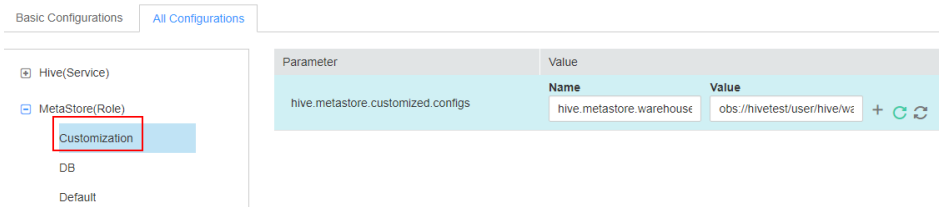

– In the navigation pane on the left, choose **HiveServer (role)** > **Customization**. Add the configuration item **hive.metastore.warehouse.dir** to **hive.metastore.customized.configs** and **hive.server.customized.configs** and set the value to an OBS path. For example, set it to **obs://hivetest/user/hive/warehouse/**, where **hivetest** is the name of the OBS parallel file system.

Figure 5-20 Configuring **hive.metastore.warehouse.dir**

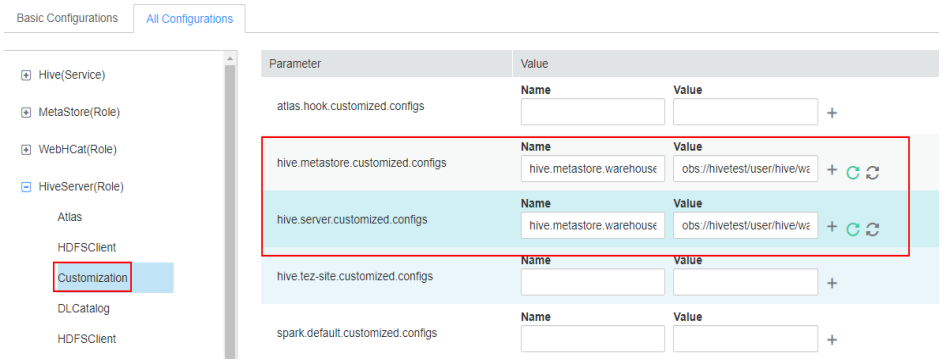

- For MRS 3.2.0 and later versions:

  Search for **hive.metastore.warehouse.dir** in the search box and change the parameter value to an OBS path, for example, **obs://hivetest/user/hive/warehouse/**. **hivetest** indicates the OBS file system name.

Figure 5-21 Configuring **hive.metastore.warehouse.dir**

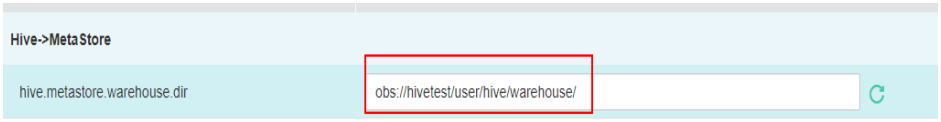

Step 2 Save the change and restart Hive.

Step 3 (Optional) Install the client by referring to **Installing a Client**. If the client has been installed in the cluster, go to **Step 4**.

Step 4 Update the client configuration file.

1. Run the following command to open **hivemetastore-site.xml** in the Hive configuration file directory on the client:

   vim *Client installation directory*/**Hive/config/hivemetastore-site.xml**

2. Change the value of **hive.metastore.warehouse.dir** to the corresponding OBS path, for example, **obs://hivetest/user/hive/warehouse/**, where **hivetest** is the OBS bucket name.

Figure 5-22 Configuring the OBS Path

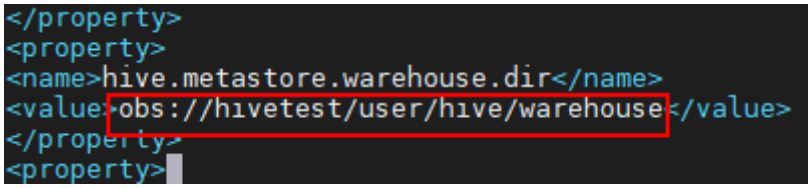

3. For MRS 3.2.0 and later versions, change the value of **hive.metastore.warehouse.dir** in **hivemetastore-site.xml** to the corresponding OBS path, for example, **obs://hivetest/user/hive/warehouse/**. The XML file is stored in the HCatalog client configuration file directory.

   vi *Client installation directory*/**Hive/HCatalog/conf/hivemetastore-site.xml**

Step 5 Log in to the beeline client, create a table, and check whether the location is the OBS path.

beeline

create table test(name string);

desc formatted test;

Location of the table is the OBS path.

Figure 5-23 Location of the Hive table

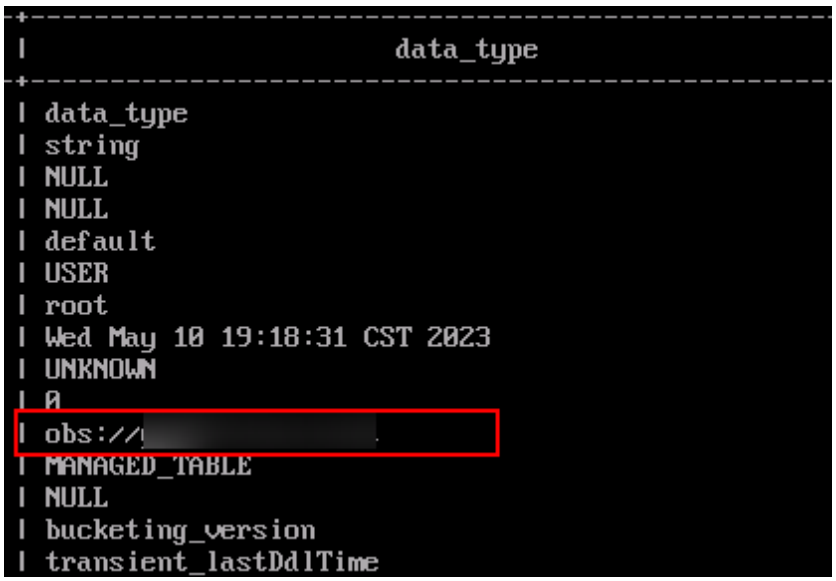

NOTE

If the location of the current database points to HDFS, tables created in the database also point to HDFS by default. You do not need to specify the location. To modify the default table creation policy, modify the location of the database to point to OBS. Perform the following steps to modify the parameters:

1. Run the following command to query the location of the database:

   show create database *obs_test*;

   Figure 5-24 Viewing the location of the Hive Table

   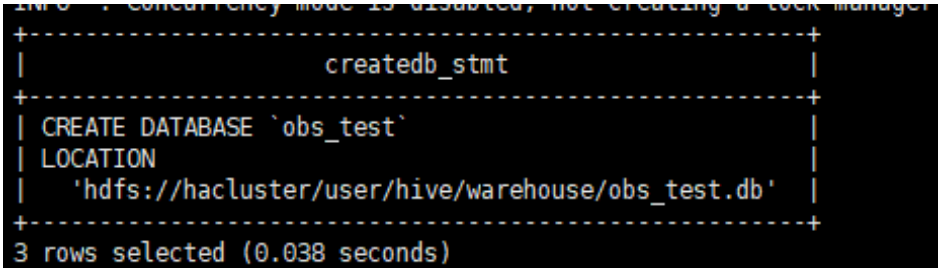

2. Run the following command to change the database location:

   alter database *obs_test* **set location** '**obs://***OBS parallel file system name*/**user/hive/warehouse/***Database name*'

   Run the **show create database** *obs_test* command to check whether the database location points to OBS.

   Figure 5-25 Check the location of the modified Hive table.

   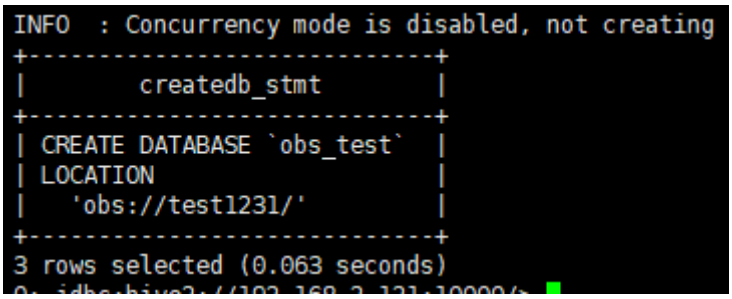

3. Run the following command to modify the table location:

   alter table *user_info* **set location** '**obs://***OBS parallel file system name*/**user/hive/warehouse/***Database name*/*Table name*'

   If the table contains data, migrate the original data file to the new location.

----End

5.2.4.5 Interconnecting MapReduce with OBS

Before performing the following operations, ensure that you have configured a storage-compute decoupled cluster by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)** or **Configuring a Storage-Compute Decoupled Cluster (AK/SK)**.

Step 1 Log in to the MRS management console and click the cluster name to go to the cluster details page.

Step 2 Choose **Components > MapReduce**. The **All Configurations** page is displayed. In the navigation tree on the left, choose **MapReduce > Customization**. In the

customized configuration items, add the configuration item **mapreduce.jobhistory.always-scan-user-dir** to **core-site.xml** and set its value to **true**.

Figure 5-26 Adding a custom parameter

Step 3 Save the configurations and restart the MapReduce service.

----End

5.2.4.6 Interconnecting Spark2x with OBS

The OBS file system can be interconnected with Spark2x after an MRS cluster is installed.

Before performing the following operations, ensure that you have configured a storage-compute decoupled cluster by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)** or **Configuring a Storage-Compute Decoupled Cluster (AK/SK)**.

Verifying OBS Access with Spark Beeline

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Spark2x** > **Configurations** > **All Configurations**.

In the left navigation tree, choose **JDBCServer2x** > **Customization**. Add **dfs.namenode.acls.enabled** to the **spark.hdfs-site.customized.configs** parameter and set its value to **false**.

Figure 5-27 Adding Spark custom parameters

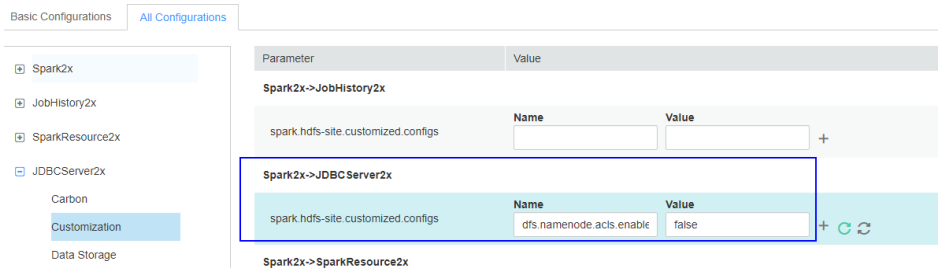

Step 2 Search for the **spark.sql.statistics.fallBackToHdfs** parameter and set its value to **false**.

Figure 5-28 Setting **spark.sql.statistics.fallBackToHdfs**

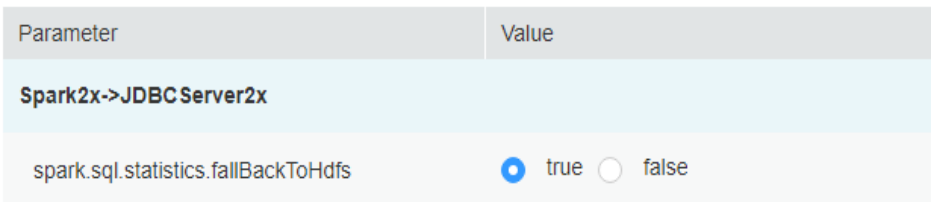

Step 3 Save the configurations and restart the JDBCServer2x instance.

Step 4 Log in to the client installation node as the client installation user.

Step 5 Run the following commands to configure environment variables:

source *Client installation directory*/**bigdata_env**

Step 6 For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

kinit *Username*

Step 7 Access OBS using Spark beeline. The following example creates a table named **test** in the **obs://mrs-word001/table/** directory.

create table test(id int) location '*obs://mrs-word001/table/*';

Step 8 Run the following command to query all tables. If table **test** is returned, OBS access is successful.

show tables;

Figure 5-29 Returned table names

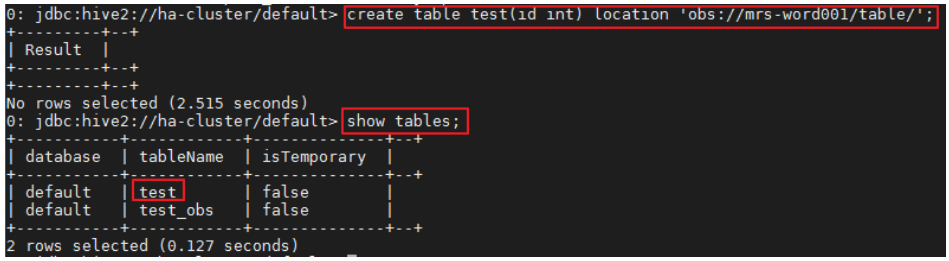

Step 9 Press **Ctrl+C** to exit Spark beeline.

----End

Verifying OBS Access with Spark SQL

Step 1 Log in to the client installation node as the client installation user.

Step 2 Run the following commands to configure environment variables:

source *Client installation directory*/**bigdata_env**

Step 3 Modify the configuration file:

vim *Client installation directory*/**Spark2x/spark/conf/hdfs-site.xml**

```
<property>
<name>dfs.namenode.acls.enabled</name>
<value>false</value>
</property>
```

Step 4 For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

kinit *Username*

Step 5 Access OBS using Spark SQL CLI. For example, create a table named **test** in the **obs://mrs-word001/table/** directory.

1. Go to the **cd** *Client installation directory*/**Spark2x/spark/bin** directory and run the **./spark-sql** command to log in to the Spark SQL CLI.
2. Run the following command in the Spark SQL CLI:

   create table test(id int) location '*obs://mrs-word001/table/*';

Step 6 Run the **show tables;** command to confirm that the table is created successfully.

Step 7 Run **exit;** to exit the Spark SQL CLI.

NOTE

If a large number of logs are printed in the OBS file system, read and write performance may be affected. You can adjust the log level of the OBS client as follows:

cd *Client installation directory*/**Spark2x/spark/conf**

vi log4j.properties

Add the OBS log level configuration to the file as follows:

log4j.logger.org.apache.hadoop.fs.obs=WARN

log4j.logger.com.obs=WARN

Figure 5-30 Adding an OBS log level

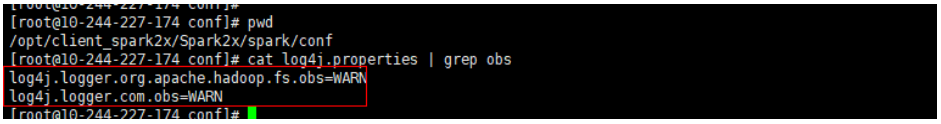

----End

Using Spark Shell to Read OBS Files

Step 1 Log in to the client installation node as the client installation user.

Step 2 Run the following commands to configure environment variables:

source *Client installation directory*/**bigdata_env**

Step 3 Modify the configuration file:

vim *Client installation directory*/**Spark2x/spark/conf/hdfs-site.xml**

```
<property>
<name>dfs.namenode.acls.enabled</name>
<value>false</value>
</property>
```

Step 4 For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

kinit *Username*

Step 5 Create an OBS file.

1. Run the following commands to log in to the Spark SQL CLI:

   cd *Client installation directory*/**Spark2x/spark/conf**

   ./spark-sql

2. Run the following commands to create a table and import data to the table:
   create database *test* **location "obs://***Parallel file system path/test***";**
   use *test*;
   create table *test1*(a int,b int) **using parquet;**
   insert into *test1* **values(1,2);**
   desc formatted *test1*;

   Figure 5-31 Checking the location of the table

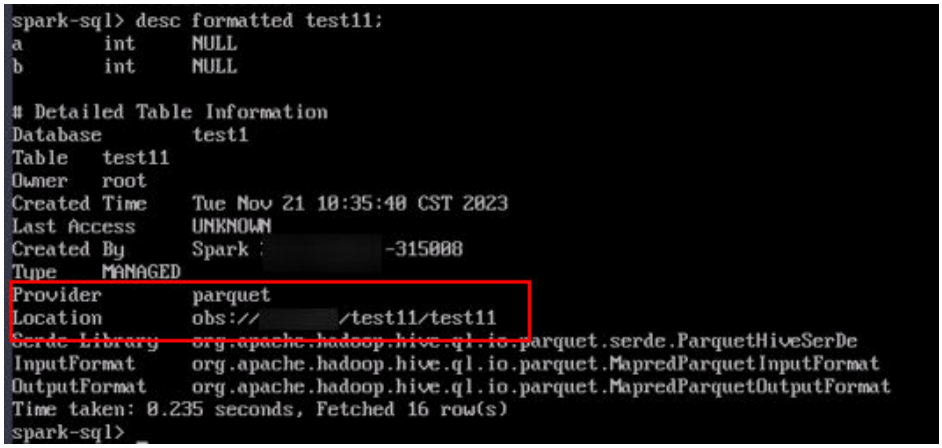

Step 6 Run the following command to go to the Spark **bin** directory:

cd *Client installation directory*/**Spark2x/spark/conf**

Run **./spark-sql** to log in to the Spark SQL CLI.

Step 7 In the Spark Shell CLI, run the following command to query the table created in **Step 5.2**:

spark.read.format("parquet").load ("obs://*Parallel file system path/test1***").show();**

Figure 5-32 Viewing table data

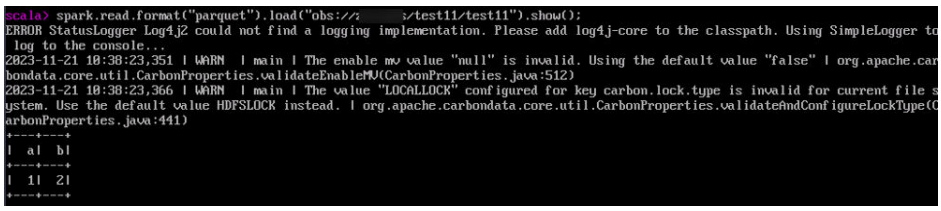

Step 8 Run the **:quit** command to exit the Spark Shell CLI.

----End

5.2.4.7 Interconnecting Sqoop with External Storage Systems

Before you start operations in this section, you have interconnected the HDFS client with OBS by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)**. You also need to download the MySQL driver package of the required version from the MySQL official website **https://downloads.mysql.com/**

archives/c-j/, decompress the package, and upload it to *Client installation directory*/**Sqoop/sqoop/lib** directory on the node where the Sqoop client is installed.

Exporting Data from HDFS to MySQL

Step 1 Log in to the node where the client is located.

Step 2 Run the following command to initialize environment variables:

source /opt/client/bigdata_env

Step 3 Run the following command to operate the Sqoop client:

sqoop export --connect jdbc:mysql://10.100.*xxx.xxx***:***3306/test* **--username root --password** *xxx* **--table** *component13* **-export-dir hdfs://hacluster/***user/hive/warehouse/component_test3* **--fields-terminated-by ',' -m 1**

Table 5-3 Parameter description

Parameter	Description
--connect	Specifies the URL for connecting to JDBC. The value is in **jdbc:mysql://***IP address of the MySQL database*:*MySQL Port*/*Database name* format.
--username	Specifies the username for connecting to the MySQL database.
-password	Specifies the password for connecting to the MySQL database. There can be security risks if a command contains the authentication password. You are advised to disable the command recording function (history) before running the command.
-table <table-name>	Specifies the name of the MySQL table used to store exported data.
-export-dir <dir>	Specifies the HDFS path of the Sqoop table to be exported.
--fields-terminated-by	Specifies the delimiter of the exported data, which must be the same as that in the HDFS data table to be exported.
-m or -num-mappers <n>	Starts *n* (4 by default) maps to import data concurrently. The value cannot be greater than the maximum number of maps in a cluster.
-direct	Imports data to a relational database using a database import tool, for example, mysqlimport of MySQL, more efficient than the JDBC connection mode.
-update-key <col-name>	Specifies the column used for updating the existing data in a relational database.

Parameter	Description
-update-mode <mode>	Specifies how updates are performed. The value can be **updateonly** or **allowinsert**. This parameter is used only when the relational data table does not contain the data record to be imported. For example, if the HDFS data to be imported to the destination table contains a data record **id=1** and the table contains an existing data record **id=2**, the update will fail.
-input-null-string <null-string>	This parameter is optional. If it is not specified, **null** will be used.
-input-null-non-string <null-string>	This parameter is optional. If it is not specified, **null** will be used.
-staging-table <staging-table-name>	Creates a table with the same data structure as the destination table for storing data before it is imported to the destination table. This parameter ensures the transaction security when data is imported to a relational database table. Due to multiple transactions during an import, this parameter can prevent other transactions from being affected when one transaction fails. For example, the imported data is incorrect or duplicate records exist.
-clear-staging-table	Clears data in the staging table before data is imported if the staging-table is not empty.

----End

Importing Data from MySQL to Hive

Step 1 Log in to the node where the client is located.

Step 2 Run the following command to initialize environment variables:

source /opt/client/bigdata_env

Step 3 Run the following command to operate the Sqoop client:

sqoop import --connect jdbc:mysql://10.100.*xxx.xxx***:3306/***test* **--username root --password** *xxx* **--table component --hive-import --hive-table** *component_test2* **--delete-target-dir --fields-terminated-by "," -m 1 --as-textfile**

Table 5-4 Parameter description

Parameter	Description
--hive-import	Imports data from a relational database to MRS Hive.
--delete-target-dir	Deletes the existing target file (if any) from Hive and imports again.

Parameter	Description
-append	Appends data to an existing dataset in the HDFS. Once this parameter is used, Sqoop imports data to a temporary directory, renames the temporary file where the data is stored, and moves the file to a formal directory to avoid duplicate file names in the directory.
-as-avrodatafile	Imports data to a data file in the Avro format.
-as-sequencefile	Imports data to a sequence file.
-as-textfile	Import data to a text file. After the text file is generated, you can run SQL statements in Hive to query the result.
-boundary-query <statement>	Specifies the SQL statement for performing boundary query. Before importing data, use a SQL statement to obtain a result set and import the data in the result set. The data format can be **-boundary-query 'select id,creationdate from person where id = 3'** (indicating a data record whose ID is 3) or **select min(<split-by>), max(<split-by>) from <table name>**. The fields to be queried cannot contain fields whose data type is string. Otherwise, the error message "java.sql.SQLException: Invalid value for getLong()" is displayed.
-columns<col,col,col...>	Specifies the fields to be imported. The format is ***-Column id,Username***.
-direct	Imports data to a relational database using a database import tool, for example, mysqlimport of MySQL, more efficient than the JDBC connection mode.
-direct-split-size	Splits the imported streams by byte. Especially when data is imported from PostgreSQL using the direct mode, a file that reaches the specified size can be divided into several independent files.
-inline-lob-limit	Sets the maximum value of an inline LOB.
-m or -num-mappers	Starts *n* (4 by default) maps to import data concurrently. The value cannot be greater than the maximum number of maps in a cluster.
-query, -e<statement>	Imports data from the query result. To use this parameter, you must specify the **-target-dir** and **-hive-table** parameters and use the query statement containing the WHERE clause as well as $CONDITIONS. Example: **-query'select * from person where $CONDITIONS' -target-dir /user/hive/warehouse/person -hive-table person**

Parameter	Description
-split-by<column-name>	Specifies the column of a table used to split work units. Generally, the column name is followed by the primary key ID.
-table <table-name>	Specifies the relational database table from which data is obtained.
-target-dir <dir>	Specifies the HDFS path.
-warehouse-dir <dir>	Specifies the directory for storing data to be imported. This parameter is applicable when data is imported to HDFS but cannot be used when you import data to Hive directories. This parameter cannot be used together with **-target-dir**.
-where	Specifies the WHERE clause when data is imported from a relational database, for example, **-where 'id = 2'**.
-z,-compress	Compresses sequence, text, and Avro data files using the GZIP compression algorithm. Data is not compressed by default.
–compression-codec	Specifies the Hadoop compression codec. GZIP is used by default.
–null-string <null-string>	Specifies the string to be interpreted as **NULL** for string columns.
–null-non-string<null-string>	Specifies the string to be interpreted as null for non-string columns. If this parameter is not specified, **NULL** will be used.
-check-column (col)	Specifies the column for checking incremental data import, for example, **id**.
-incremental (mode) append or last modified	Incrementally imports data. **append**: appends records, for example, appending records that are greater than the value specified by **last-value**. **lastmodified**: appends data that is modified after the date specified by **last-value**.
-last-value (value)	Specifies the maximum value (greater than the specified value) of the column after the last import. This parameter can be set as required.

----End

Sqoop Usage Example

- Importing data from MySQL to HDFS using the **sqoop import** command

sqoop import --connect jdbc:mysql://10.100.231.134:3306/test --username root --password *xxx* --query 'SELECT * FROM component where $CONDITIONS and component_id ="MRS 1.0_002"' --target-dir /tmp/component_test --delete-target-dir --fields-terminated-by "," -m 1 --as-textfile

- Exporting data from OBS to MySQL using the **sqoop export** command

  sqoop export --connect jdbc:mysql://10.100.231.134:3306/test --username root --password *xxx* --table component14 -export-dir obs://obs-file-bucket/xx/part-m-00000 --fields-terminated-by ',' -m 1

- Importing data from MySQL to OBS using the **sqoop import** command

  sqoop import --connect jdbc:mysql://10.100.231.134:3306/test --username root --password *xxx* --table component --target-dir obs://obs-file-bucket/xx --delete-target-dir --fields-terminated-by "," -m 1 --as-textfile

- Importing data from MySQL to OBS tables outside Hive

  sqoop import --connect jdbc:mysql://10.100.231.134:3306/test --username root --password *xxx* --table component --hive-import --hive-table component_test01 --fields-terminated-by "," -m 1 --as-textfile

MySQL Driver Package Is Missing During Data Import or Export

If the error "Could not load db driver class: com.mysql.jdbc.Driver" is reported when you run the **sqoop import** or **sqoop export** command, the MySQL driver package is missing. Download the MySQL driver package from the MySQL official website, decompress it, upload it to the *Client installation directory*/**Sqoop/sqoop/lib**, and run the command again.

Figure 5-33 An error indicating that the MySQL driver package is missing

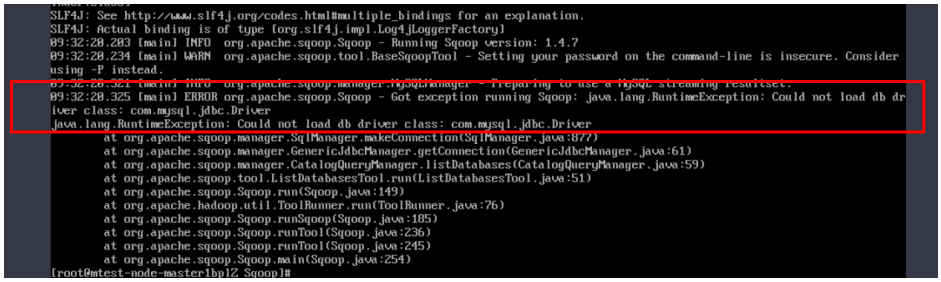

5.2.4.8 Interconnecting Hudi with OBS

Step 1 Log in to the client installation node as the client installation user.

Step 2 Run the following commands to configure environment variables:

source *Client installation directory*/**bigdata_env**

source *Client installation directory*/**Hudi/component_env**

Step 3 Modify the configuration file:

vim *Client installation directory*/**Hudi/hudi/conf/hdfs-site.xml**

```
<property>
<name>dfs.namenode.acls.enabled</name>
<value>false</value>
</property>
```


Step 4 For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

kinit *Username*

Step 5 Start spark-shell and run the following commands to create a COW table and save it in OBS:

```
import org.apache.hudi.QuickstartUtils._
import scala.collection.JavaConversions._
import org.apache.spark.sql.SaveMode._
import org.apache.hudi.DataSourceReadOptions._
import org.apache.hudi.DataSourceWriteOptions._
import org.apache.hudi.config.HoodieWriteConfig._
val tableName = "hudi_cow_table"
val basePath = "obs://testhudi/cow_table/"
val dataGen = new DataGenerator
val inserts = convertToStringList(dataGen.generateInserts(10))
val df = spark.read.json(spark.sparkContext.parallelize(inserts, 2))
df.write.format("org.apache.hudi").
options(getQuickstartWriteConfigs).
option(PRECOMBINE_FIELD_OPT_KEY, "ts").
option(RECORDKEY_FIELD_OPT_KEY, "uuid").
option(PARTITIONPATH_FIELD_OPT_KEY, "partitionpath").
option(TABLE_NAME, tableName).
mode(Overwrite).
save(basePath);
```

NOTE

obs://testhudi/cow_table/ is the OBS path, and **testhudi** is the bucket name. Change them based on site requirements.

Step 6 Use DataSource to check whether the table is successfully created and whether the data is normal.

```
val roViewDF = spark.
read.
format("org.apache.hudi").
load(basePath + "/*/*/*/*")
roViewDF.createOrReplaceTempView("hudi_ro_table")
```


spark.sql("select * from hudi_ro_table").show()

Step 7 Run the **:q** command to exit the spark-shell CLI.

----End

5.2.5 Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS

When fine-grained permission control is enabled, you can configure OBS access permissions to implement access control on directories in OBS file systems.

NOTE

This section does not apply to MRS 1.9.2.

This function enables you to control MRS users' access to OBS resources. For example, if you allow user group A to only access log files in a specified OBS file system, perform the following operations:

1. Configure an agency with OBS access permissions for an MRS cluster so that OBS can be accessed using the temporary AK/SK automatically obtained by the ECS. This prevents the AK/SK from being exposed in the configuration file.
2. Create a policy on the IAM console to allow access to log files in a specified OBS file system, and create an agency bound to the policy permission.
3. In the MRS cluster, bind the new agency to user group A so that user group A only has the permission to access log files in the specified OBS file system.

In the following scenarios, the username used for submitting jobs is an internal username so that MRS multi-user access to OBS is not supported.

- For spark-beeline, the internal username used for submitting jobs is **spark** in a security cluster and **omm** in a normal cluster.
- For the HBase shell, the internal username used for submitting jobs is **hbase** in a security cluster and **omm** in a normal cluster.
- For Presto, the internal username used for submitting jobs in the security cluster is **omm** or **hive**, and that in the normal cluster is **omm**. (Choose **Components** > **Presto** > **Service Configuration**. Change **Basic** to **All** in the parameter type drop-down box.) Then, search for and change the value of **hive.hdfs.impersonation.enabled** to **true** to enable MRS multi-user to access OBS with fine-grained permissions.

Prerequisites

- Fine-grained permission control has been enabled. For details about permissions management, see **Creating an MRS User**.
- You have a basic knowledge of **Cloud Service Delegation** and OBS fine-grained policies.

Step 1: Configuring an Agency with OBS Access Permission for a Cluster

Follow instructions in **Configuring a Storage-Compute Decoupled Cluster (Agency)** to configure an agency with OBS access permissions.

The agency takes effect for all users (including internal users) and user groups in the cluster. To control the permissions of users and user groups in the cluster to access OBS, perform the following operations.

NOTE

When you configure permissions on an OBS path, if the write permission is configured, you need to configure the corresponding recycle bin path.

The default recycle bin path is **/user/${current.user}/.Trash/**, in which **${current.user}** indicates the current user.

Step 2: Creating a Policy and an Agency on IAM

Create policies with different access permissions and bind the policies to the agency. For details, see **Creating a Policy and an Agency on IAM**.

Step 3: Configuring OBS Permission Control Mappings on the MRS Cluster Details Page

Step 1 On the MRS management console, choose **Clusters** > **Active Clusters** and click the cluster name.

Step 2 In the **Basic Information** area on the **Dashboard** tab page, click **Manage** next to **OBS Permission Control**.

Step 3 Click **Add Mapping** and set parameters according to **Table 5-5**.

Table 5-5 OBS permission control parameters

Parameter	Description
IAM Agency	Select the agency created in **Step 2**.
Type	• **User**: User-level mapping • **Group**: User group-level mapping **NOTE** • User-level mapping takes priority over user group-level mapping. If you select **Group**, you are advised to enter the primary group name in **MRS User (User Group)**. • Do not use the same username (user group) for multiple mapping records.

Parameter	Description
MRS User (User Group)	Use commas (,) to separate multiple names of users or user groups. **NOTE** • If OBS permission control is not configured for a user and no AK and SK are configured, the **OBS OperateAccess** permission in **MRS_ECS_DEFAULT_AGENCY** will be used for accessing OBS. You are advised not to bind the internal user of a component to an agency. • If you need to configure an agency for the internal user of a component when submitting a job in the following scenarios, the requirements are as follows: – To control permissions on spark-beeline operations, set the username to **spark** for a security cluster and **omm** for a normal cluster. – To control permissions on HBase shell operations, set the username to **hbase** for a security cluster and **omm** for a normal cluster. – To control permissions on Presto, set the username to **omm**, **hive**, and the username used for logging in to the client for a security cluster and **omm** and the username used for logging in to the client for a normal cluster. – If you want to use Hive to create tables in beeline mode, set the username to the internal user **hive**.

Step 4 Click **OK**.

Step 5 Select **I agree to authorize the trust relationships between MRS Users (Groups) and IAM agencies**, and click **OK**. The mapping between the MRS user and OBS permission is added.

If appears next to **OBS Permission Control** on the **Dashboard** tab page or the mapping table has been updated for OBS permission control, the mapping takes effect. It takes about 1 minute to for the mapping to take effect.

In the **Operation** column of the mapping list, you can edit or delete the added mapping.

NOTE

- If OBS permission control is not configured for a user and no AK and SK are configured, the permissions owned by the agency configured for the cluster in the **Object Storage Service (OBS)** project will be used to access OBS.
- Regardless of whether OBS permission control is configured, AK/SK permission is used for accessing OBS once it is configured.
- Security Administrator permission is required to modify, create, or delete a mapping.
- To apply the mapping changes in spark-beeline, hive beeline, and Presto, you need to restart Spark, exit beeline and enter again, and restart Presto, respectively.

----End

Component Access to OBS When OBS Permission Control Is Enabled

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

Step 2 Set environment variables (The default client installation path for MRS 3.x or later is /opt/Bigdata/client, and for versions earlier than MRS 3.x is /opt/client. Configure the path based on site requirements.).

source /opt/Bigdata/client/bigdata_env

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

kinit *MRS cluster user*

Example: **kinit admin**

Step 4 If the Kerberos authentication is disabled for the current cluster, run the following commands to log in. Note that you should create a user that belongs to the **supergroup** group by referring to **Creating a User** and replace *XXXX* with the username:

mkdir /home/XXXX

chown XXXX /home/XXXX

su - XXXX

Step 5 Access OBS. You do not need to configure the AK, SK, and endpoint. The OBS path format is **obs://buck_name/*XXX***.

Example: **hadoop fs -ls "obs://obs-example/job/hadoop-mapreduce-examples-3.1.2.jar"**

NOTE

- If you want to use **hadoop fs** to delete files on OBS, use **hadoop fs -rm -skipTrash** to delete the files.
- If data import is not involved when a table is created using spark-sql and spark-beeline, OBS will not be accessed. That is, if you create a table in an OBS directory on which you do not have permission, the **CREATE TABLE** operation will still be successful, but the error message "**403 AccessDeniedException**" is displayed when you insert data.

----End

Creating a Policy and an Agency on IAM

Step 1 Create a policy on IAM.

1. Log in to the IAM console.
2. Choose **Permissions**. On the displayed page, click **Create Custom Policy**.
3. Set parameters according to **Table 5-6**. Obtain the customized OBS policy samples that are frequently used by referring to **OBS Custom Policies**.

Table 5-6 Policy parameters

Parameter	Description
Policy Name	Only letters, digits, spaces, and special characters (-_.,) are allowed.
Scope	Select **Global services**, because OBS is a global service.
Policy View	Select **Visual editor**.
Policy Content	1. **Allow**: Select **Allow**. 2. **Select service**: Select **Object Storage Service (OBS)**. 3. **Select action**: Select **WriteOnly**, **ReadOnly**, and **ListOnly**. 4. **Specific resources**: a. Set **object** to **Specify resource path**, click **Add Resource Path**, and enter *obs_bucket_name*/**tmp/** and *obs_bucket_name*/**tmp/***. The **/tmp** directory is used as an example. If you need to add permissions for other directories, perform the following steps to add the directories and resource paths of all objects in the directories. b. Set **bucket** to **Specify resource path**, click **Add Resource Path**, and enter *obs_bucket_name*. 5. (Optional) Add request condition, which does not need to be added currently.
Description	(Optional) Brief description about the policy.

NOTE

If the data write operation of each component is implemented in **rename** mode, the permission to delete objects must be configured when data is written.

4. Click **OK** to save the policy.

Step 2 Create an agency on IAM.

1. Log in to the IAM console.
2. Choose **Agencies**. On the displayed page, click **Create Agency**.
3. Set parameters according to **Table 5-7**.

Table 5-7 Agency parameters

Parameter	Description
Agency Name	Only letters, digits, spaces, and special characters (-_.,) are allowed.
Agency Type	Select **Common account**.
Delegated Account	Enter your cloud account, that is, the account you register using your mobile phone number. It cannot be a federated user or an IAM user created using your cloud account.
Validity Period	Set this parameter as required.
Description	(Optional) Brief description about the agency.
Permissions	1. In the **Project [Region]** column, locate the row where **OBS** is, click **Attach Policy**. 2. Select the policy created in **Step 1** to display it in **Selected Policies**. 3. Click **OK**.

4. Click **OK** to save the agency.

NOTE

If you modify an agency and policies bound to it after using the agency to access OBS, the modification will take effect within 15 minutes.

----End

5.2.6 Accessing OBS from a Client on a Node Outside the Cluster

Scenario

In the OBS decoupled storage and compute scenario, obtain a temporary AK/SK through an agency to access the OBS server. To access OBS from a client on a node outside the cluster, obtain an AK/SK through Guardian. Guardian is an MRS self-developed component that helps clients outside the cluster access OBS through temporary AKs/SKs.

NOTE

This function is available in MRS 3.1.5 or later.

How Does Guardian Enable a Client on a Node Outside the Cluster to Access OBS?

Step 1 Install Guardian in the cluster.

- For a new cluster in creation, select **Guardian**.

  Figure 5-34 Selecting Guardian

  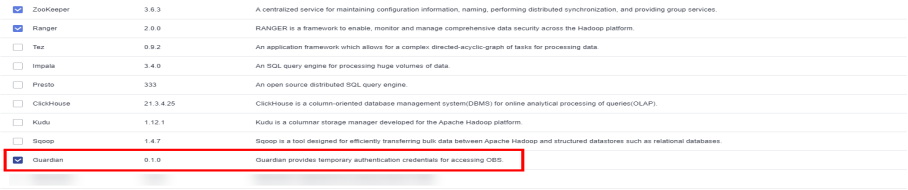

- For an existing cluster, add the Guardian component in the **Components** tab.
  a. On the cluster details page, choose **Components** and click **Add Service**.
  b. In the service list, select the services to be added and click **Next**.

  Figure 5-35 Adding Guardian

  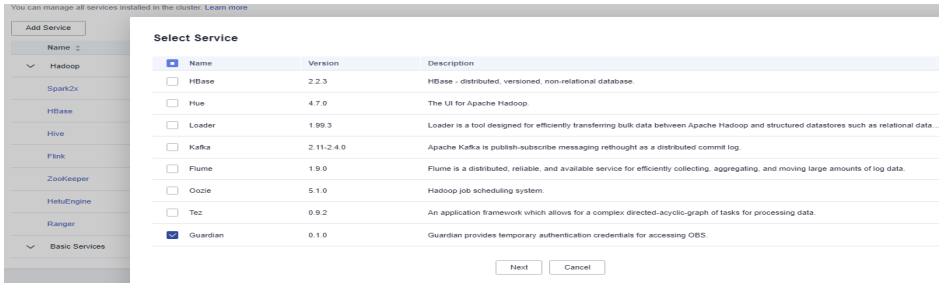

  c. On the **Topology Adjustment** page, select the node where the service is to be deployed. (You are advised to deploy the Guardian service on the master node.)
  d. Click **OK**. After the service is added, you can view the added service on the **Components** page.

  NOTE

  The services added on the console are automatically synchronized to Manager.

Step 2 Configure storage and compute decoupling for the cluster by referring to **Configuring a Storage-Compute Decoupled Cluster (Agency)** or **Configuring Fine-Grained Permissions for MRS Multi-User Access to OBS**.

Step 3 Install or update the client.

- For details about how to install the client on a node outside the cluster, see **Installing a Client (MRS 3.x or Later)**.
- For details about how to update an existing client, see **Updating a Client (Version 3.x or Later)**.

Step 4 After the installation is successful, you can access OBS.

Example:

If the client directory is **/opt/client**, run the **source /opt/client/bigdata_env** command. Use the actual client installation path.

Run the **hdfs dfs -ls** *obs* command.

If the information in the following figure is displayed, the installation is successful:

Figure 5-36 Accessing OBS

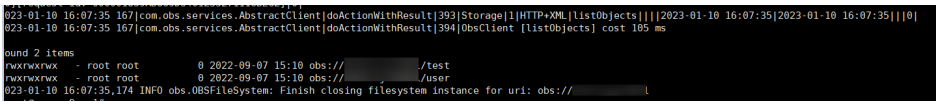

----End

5.3 Interconnecting with OBS Using the Guardian Service

5.3.1 Scenarios

Configuring Storage and Compute Decoupling

1. Create an MRS cluster.

   The MRS cluster must contain basic components such as Guardian, Ranger, and Hadoop.

   NOTE

   Currently, only MRS 3.3.0-LTS and later versions support interconnection with OBS using the Guardian component.

2. Create an OBS agency.

   Create an agency with OBS access permissions, which is used for interconnecting Guardian with OBS.

3. Enable the interconnection between Guardian and OBS and configure parameters.

   Modify the configuration parameters for the Guardian service and configure the IAM agency authentication information.

4. Configure the policy for clearing component data in the recycle bin directory.

   In the storage-compute decoupling scenario, the prevention against accidental deletion is enabled by default for components connected to OBS. When a user deletes data, the deleted object is moved to the corresponding recycle bin directory. You need to configure a lifecycle rule for the corresponding directory in the OBS file system to prevent the storage space from being used up.

5. Interconnect components with OBS.

   Components in the MRS cluster can directly access the corresponding path after the required permissions for accessing OBS buckets are obtained. You can use the component client to directly access resources in the OBS file system in absolute path mode.

Configuring OBS Permissions

If Guardian is deployed with decoupled storage and compute and Ranger authentication is enabled for MRS clusters, Ranger administrators can configure read and write permissions on OBS directories or files for cluster users.

With the Guardian permission model, storage and compute decoupling, and Hive cascading authorization, authorization is not required after the first permission service table authorization on the Ranger page and the system automatically associates the permissions of OBS data storage source in a fine-grained manner. The storage path of the table does not need to be sensed.

NOTE

- On the Ranger page, OBS permission authorization only support Manager custom user groups (built-in user groups are not supported). The user group contains a maximum of 52 characters, including digits 0 to 9, letters A to Z, underscores (_), and number signs (#). Otherwise, the policy fails to be added.
- For clusters with Kerberos authentication enabled, permissions need to be granted based on Ranger. For clusters with Kerberos authentication disabled, OBS permissions are granted by default, and no additional configuration is required.
- If Kerberos authentication is not enabled for the current cluster, the user who accesses OBS must belong to the **supergroup** group.

5.3.2 Interconnecting the Guardian Service with OBS

Scenario

This section describes how to enable storage and compute decoupling for the Guardian component. After this feature is enabled, Guardian can provide temporary authentication credentials for services such as HDFS, Hive, Spark, Loader, and HetuEngine to access OBS when decoupled storage and compute are used.

Perform the following steps to interconnect Guardian with OBS:

1. **Creating an OBS Parallel File System**
2. **Creating a Cloud Service Agency and Binding It to a Cluster**
3. **Creating an Agency for a Regular Account**
4. **Configuring a Cloud Service Agency**

5. **Granting OBS Access Permission to Guardian**
6. **Enabling Cascading Authorization for Hive Tables**
7. **Configuring the Recycle Bin Cleanup Policy**

Prerequisites

- Components such as Guardian, Ranger, and Hadoop have been installed in the cluster.
- If Guardian is installed after components such as Hadoop, HetuEngine, Hive, and Spark are installed, the Guardian client must be downloaded again and the default client for job submission on the management plane must be refreshed.

Impact on the System

- After the configuration is complete, you need to refresh the configuration of the original client or reinstall the client.
- To submit a job on console, log in to the active OMS node as user **omm** and run the **sh /opt/executor/bin/refresh-client-config.sh** command to refresh the built-in client of the cluster.

Creating an OBS Parallel File System

1. Log in to the OBS console.
2. Choose **Parallel File Systems** > **Create Parallel File System**.
3. Enter a file system name, for example, **guardian-obs**.

   The name of an enterprise project must be the same as that of the MRS cluster. Set other parameters.
4. Click **Create Now**.

Creating a Cloud Service Agency and Binding It to a Cluster

1. Log in to the Huawei Cloud management console.
2. In the service list, choose **Management & Governance** > **Identity and Access Management**.
3. Choose **Agencies**. On the displayed page, click **Create Agency**.
4. Set the agency name, for example, **mrs_ecs_obs**.
5. Set **Agency Type** to **Cloud service** and select **Elastic Cloud Server (ECS) and Bare Metal Server (BMS)** to authorize ECS or BMS to invoke OBS.

Figure 5-37 Creating an agency

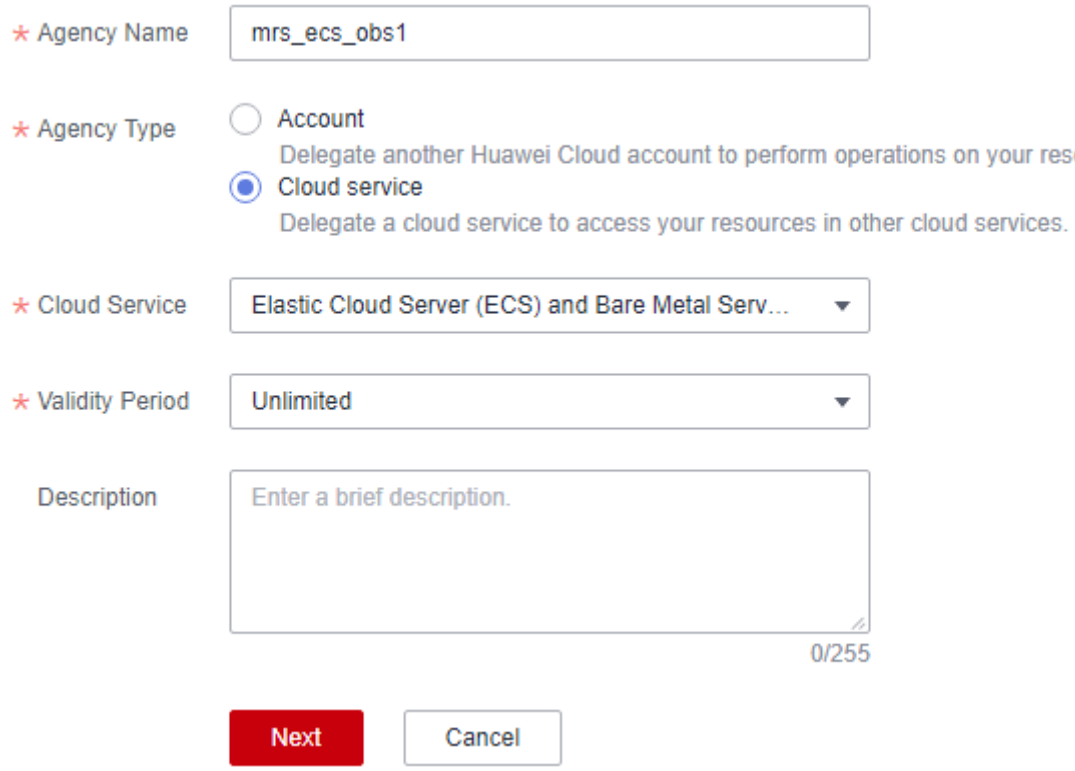

6. Set **Validity Period** to **Unlimited** and click **Next**.
7. On the displayed page, search for the **OBS OperateAccess** policy and enable it.

Figure 5-38 Configuring permissions

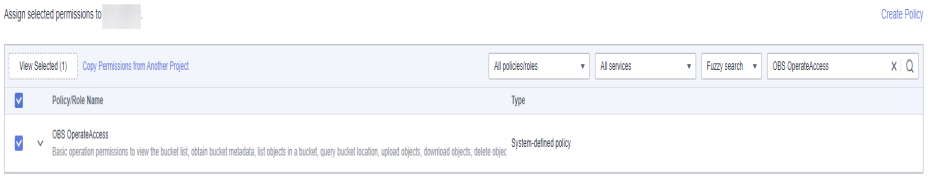

8. Click **Next**, select **All resources**, click **Show More**, select **Global resources**, and click **OK**.
9. In the dialog box that is displayed, click **OK** to start authorization. After "**Authorization successful.**" is displayed, click **Finish**. The agency is successfully created.
10. Log in to the MRS console. In the navigation pane on the left, choose **Clusters** > **Active Clusters**.
11. Click the name of the target cluster to go to details page.
12. In the **Dashboard** tab, click **Synchronize** on the right of **IAM User Sync** to synchronize the IAM user.
13. In the **Dashboard** tab, click **Manage Agency** on the right of **Agency**, select the created agency, for example, **mrs_ecs_obs**, and click **OK** to bind the agency to the cluster.

Figure 5-39 Binding an agency

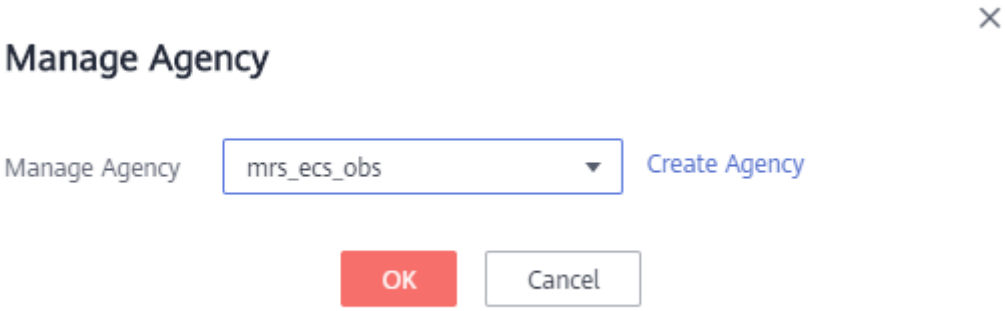

Creating an Agency for a Regular Account

1. Log in to the Huawei Cloud management console.
2. In the service list, choose **Management & Governance** > **Identity and Access Management**.
3. Choose **Agencies**. On the displayed page, click **Create Agency**.
4. Enter an agency name, for example, **agency-MRS-to-OBS**.
5. Set **Agency Type** to **Account**.
6. Enter your cloud account for **Delegated Account**, that is, the account you register using your mobile phone number. It cannot be a federated user or an IAM user created using your cloud account.
7. Set **Validity Period** to **Unlimited** and click **Next**.
8. On the displayed page, search for the **OBS Administrator** policy and enable it.

Figure 5-40 Configuring permission

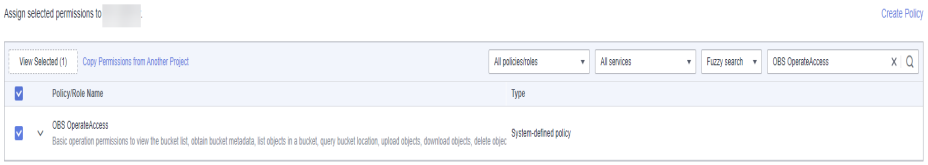

9. Click **Next**, select **All resources**, click **Show More**, select **Global resources**, and click **OK**.
10. After the agency is created, check and record the agency ID.

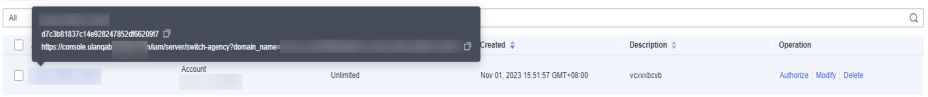

Configuring a Cloud Service Agency

1. Log in to the Huawei Cloud management console.
2. In the service list, choose **Management & Governance** > **Identity and Access Management**.

3. Select **Agencies** and click the agency **mrs_ecs_obs** created in **Creating a Cloud Service Agency and Binding It to a Cluster**
4. Choose **Permissions** > **Authorize**, click **Create Policy** in the upper right corner, and set the parameters as follows:
   - **Policy Name**: Enter a policy name, for example, **guardian-assume-policy**.
   - **Policy View**: Select **JSON**.
   - **Policy Content**: Configure the policy as follows. *{Agency ID}* indicates the ID recorded in **10**.

```
{
   "Version": "1.1",
   "Statement": [
      {
         "Action": [
            "iam:agencies:assume"
         ],
         "Resource": {
            "uri": [
               "/iam/agencies/{Agency ID}"
            ]
         },
         "Effect": "Allow"
      }
   ]
}
```

5. Click **Next**. On the **Select Policy/Role** page, select the policy created in **4**.
6. Click **Next**, select **All resources**, click **Show More**, select **Global resources**, and click **OK**.

Granting OBS Access Permission to Guardian

1. Log in to FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**, click **Configurations**, and then **All Configurations**. On the displayed page, search for and modify the following parameters.

Parameter	Description	Value
fs.obs.guardian.accesslabel.enabled	Whether to enable **access label** for using Guardian to connect to OBS.	true
fs.obs.guardian.enabled	Whether to enable Guardian.	true
fs.obs.delegation.token.providers	Delegation token generator. If **fs.obs.guardian.enabled** is set to **true**, configure both **com.huawei.mrs.dt.MRSDelegationTokenProvider** and **com.huawei.mrs.dt.GuardianDTProvider**.	com.huawei.mrs.dt.MRSDelegationTokenProvider and com.huawei.mrs.dt.GuardianDTProvider
token.server.access.label.agency.name	Name of the specified IAM agency, which is the agency created in **Creating an Agency for a Regular Account**.	agency-MRS-to-OBS

2. Save the service configuration, choose **More** > **Restart Configuration-Expired Instances** on the FusionInsight Manager home page, and restart all service instances whose configurations have expired as prompted.
3. To submit jobs on the MRS console, log in to the active OMS node as user **omm** and run the following command to refresh the built-in client configuration:

   sh /opt/executor/bin/refresh-client-config.sh

Enabling Cascading Authorization for Hive Tables

1. Log in to FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** and click **Configurations**.
2. Search for the **ranger.ext.authorization.cascade.enable** parameter and set it to **true**.

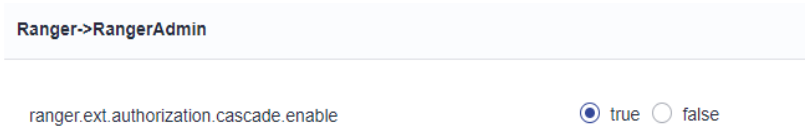

3. Click **Save**.
4. Click **Instance** and select all RangerAdmin instances. Click **More** and select **Restart Instance**. Enter the password, and click **OK** to restart all RangerAdmin instances.

Configuring the Recycle Bin Cleanup Policy

1. Log in to the OBS Console.
2. Select **Parallel File Systems** and click the file system created in **Creating an OBS Parallel File System**.
3. Choose **Basic Configurations** > **Lifecycle Rules** and click **Create** to create a lifecycle rule for the **/user/.Trash** directory.

CAUTION

For clusters that use decoupled storage and compute, configure a lifecycle policy for the related directories by referring to this chapter. Otherwise, the storage space may be used up and storage fees may increase.

Table 5-8 Parameters for creating a lifecycle rule

Parameter	Description	Example
Status	Whether to enable the lifecycle rule.	Enable
Rule Name	User-defined rule name, which is used to identify different lifecycle configurations.	rule-test

Parameter	Description	Example
Prefix	Prefix of the objects to which the lifecycle rule applies. Generally, the recycle bin directory of MRS components is **/user/.Trash**.	user/.Trash
Transition to Infrequent Access After (Days)	Number of days before transitioning to infrequent access after the object is last updated. The value of this parameter must be at least **30**.	30 days
Transition to Archive After (Days)	Number of days before transitioning to archive after the object is last updated. If **Transition to Infrequent Access After (Days)** is also configured, after the lifecycle is transitioned to infrequent access, wait at least 30 days before transitioning it to archive. If only **Transition to Archive After (Days)** is configured, there is no time limit.	31 days
Delete Files After (Days)	Number of days before being deleted by OBS after the object is last updated. This parameter must be larger than the above two parameters.	32 days
Delete Fragments After (Days)	Number of days before fragments are expired and deleted by OBS automatically.	30 days

4. Click **OK** to complete the lifecycle rule configuration.

   To modify the lifecycle configuration, locate the lifecycle rule, click **Edit** or **Disable** on the right. Click **Enable** to enable the lifecycle rule.

5.3.3 Interconnecting Components with OBS Using Guardian

5.3.3.1 Interconnecting Hive with OBS

Interconnecting with OBS

MRS clusters allow Hive to connect to OBS through Metastore.

Interconnecting Hive with OBS through Metastore

Step 1 You have configured storage and compute decoupling by referring to **Interconnecting the Guardian Service with OBS**.

Step 2 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Hive**, and click **Configurations**.

Step 3 Search for **hive.metastore.warehouse.dir** in the search box and change the parameter value to an OBS path, for example, **obs://hivetest/user/hive/warehouse/**. **hivetest** indicates the OBS file system name.

Figure 5-41 hive.metastore.warehouse.dir configuration

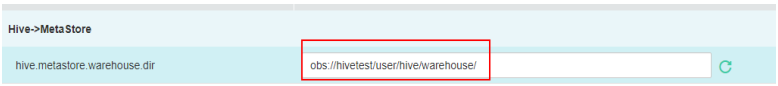

Step 4 Save the configuration, choose **Cluster** > **Services**, and restart the Hive service in the service list.

Step 5 Update the client configuration file.

1. Log in to the node where the Hive client is located and run the following command to modify **hivemetastore-site.xml** in the Hive client configuration file directory:

   vi *Client installation directory*/**Hive/config/hivemetastore-site.xml**

2. Change the value of **hive.metastore.warehouse.dir** to the corresponding OBS path, for example, **obs://hivetest/user/hive/warehouse/**.

   ```
   </property>
   <property>
   <name>hive.metastore.warehouse.dir</name>
   <value>obs://hivetest/user/hive/warehouse</value>
   </property>
   <property>
   ```

3. Change the value of **hive.metastore.warehouse.dir** of **hivemetastore-site.xml** in the HCatalog client configuration file directory to the corresponding OBS path, for example, **obs://hivetest/user/hive/warehouse/**.

   vi *Client installation directory*/**Hive/HCatalog/conf/hivemetastore-site.xml**

Step 6 Go to the Hive Beeline CLI, create a database, and ensure that the location is an OBS path.

cd *Client installation directory*

kinit *Component operation user*

beeline

create database testdb1;

show create database testdb1;

```
+----------------------------------------------------+
|                   createdb_stmt                    |
+----------------------------------------------------+
| CREATE DATABASE `testdb1`                          |
| LOCATION                                           |
|   'obs://0417x86/user/hive/warehouse/testdb1.db'   |
+----------------------------------------------------+
```

----End

Configuring Ranger Permissions

- Granting the read and write permissions on OBS paths to the **hive** user group

  a. Log in to the Ranger web UI as the Ranger administrator. On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area, and assign the **Read** and **Write** permissions on

the OBS storage path to the **hive** user group. If this operation is successful, all users in the **hive** group can access the Hive data warehouse path.

For example, assign the **Read** and **Write** permissions on the **obs://hivetest/user/hive/warehouse/** directory to the **hive** user group:

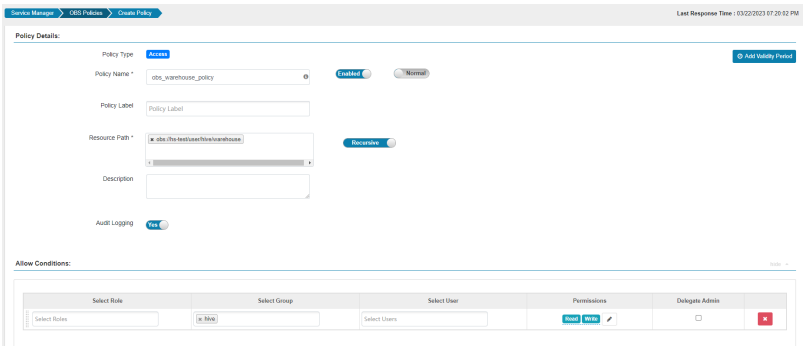

b. Choose **Settings** > **Roles**, click **Add New Role**, and create a role whose **Role Name** is **hive**.

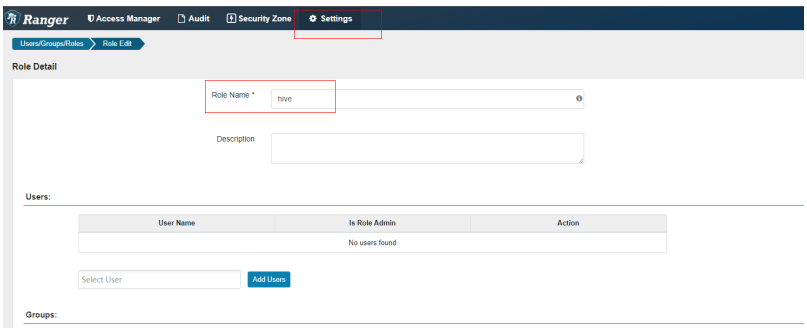

- Granting the read and write permissions on OBS paths to a custom user group

  a. Log in to FusionInsight Manager and choose **System** > **Permission** > **User Group**. On the displayed page, click **Create User Group**.

  b. Create a user group without a role, for example, **hiveobs1**, and bind the user group to the corresponding user.

  c. Log in to the Ranger management page as the **rangeradmin** user.

  d. On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area.

  e. Grant the **Read** and **Write** permissions on the OBS storage path to the **hiveobs1** user group. In this case, all users bound to the **hiveobs1** user group can access the Hive data warehouse path.

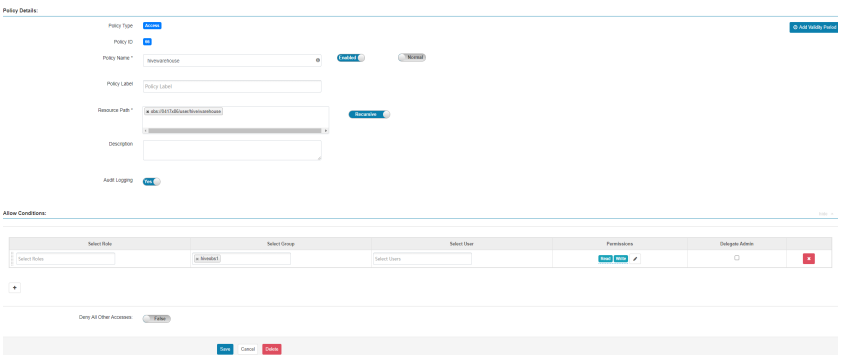

- Creating a database,table, or partition in a custom location and granting read and write permissions on OBS paths

  a. Log in to the Ranger web UI as the Ranger administrator.

  b. On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area, and assign the **Read** and **Write** permissions on the OBS storage path to the user group of the corresponding user.

  For example, assign the **Read** and **Write** permissions on the **obs://obs-test/test/** directory to the **hgroup1** user group, as shown in the following figure.

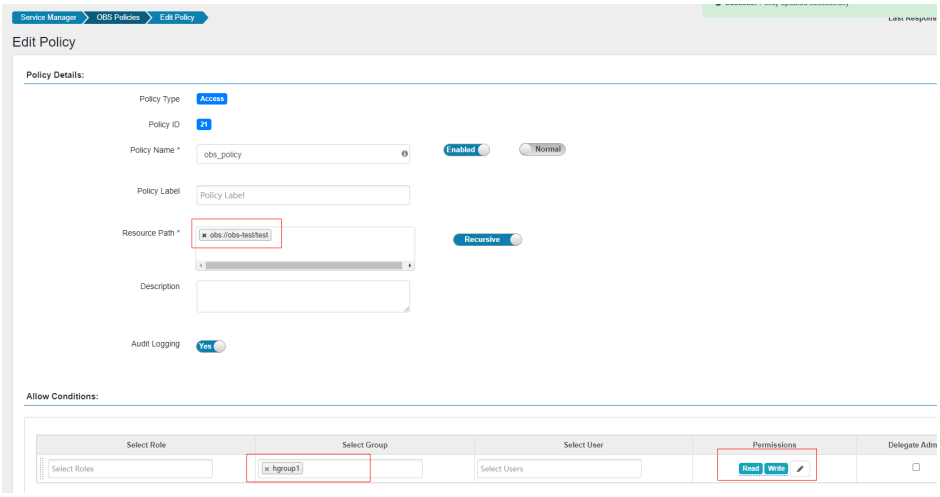

  c. On the home page, click the component plug-in name **Hive** in the **HADOOP SQL** area, and add a URL policy that grants the **Read** and **Write** permissions on the OBS path to the user group of the corresponding user.

  For example, create the **hive_url_policy URL** policy for the **hgroup1** user group and assign the **Read** and **Write** permissions on the **obs://obs-test/test/** directory to the user group, as shown in the following figure.

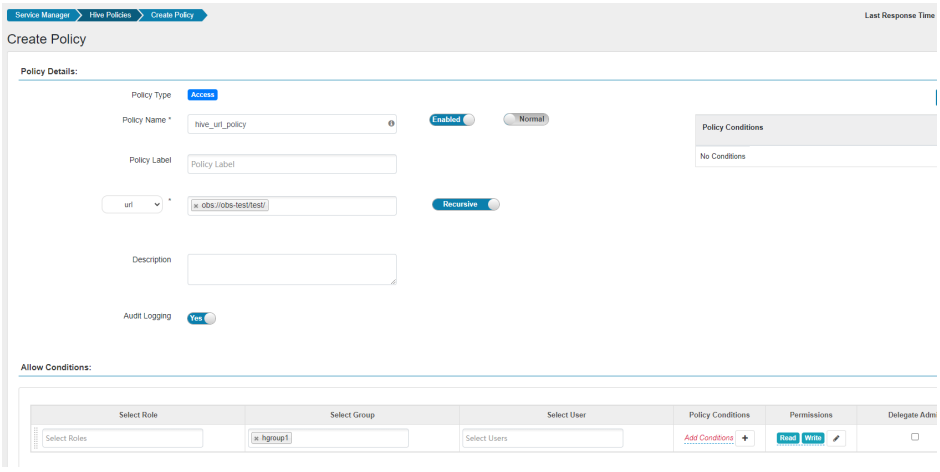

  d. Log in to the beeline client and set **Location** to the OBS file system path when creating a table.

  cd *Client installation directory*

  kinit *Component operation user*

  beeline

For example, to create a table named **test** whose **Location** is **obs://obs-test/test/***Database name*/*Table name*, run the following command:

create external table test(name string) location "obs://obs-test/test/*Database name*/*Table name***";**

NOTE

- To authorize a view chart, you need to grant the view chart permission and the physical table path permission corresponding to the view chart.
- Cascading authorization can be performed only on databases and tables, and cannot be on partitions. If a partition path is not in the table path, you need to manually authorize the partition path.
- Cascading authorization for **Deny Conditions** in the Hive Ranger policy is not supported. That is, the Deny Conditions permission only restricts the table permission and cannot generate the permission of the HDFS/OBS storage source.
- The permission of the HDFS Ranger policy is prior to that of the HDFS/OBS storage source generated by cascading authorization. If the HDFS Ranger permission has been set for the HDFS storage source of the table, the cascading permission does not take effect.
- **alter** operations cannot be performed on tables whose storage source is OBS after cascading authorization. To perform the **alter** operation, you need to grant the **Read** and **Write** permissions of the parent directory of the OBS table path to the corresponding user group.

5.3.3.2 Interconnecting Flink with OBS

Interconnecting with OBS

Step 1 Log in to the Flink client installation node as the client installation user.

Step 2 Run the following command to initialize environment variables:

source *Client installation directory*/**bigdata_env**

Step 3 Configure the Flink client.

Step 4 Start a session.

- Normal cluster (Kerberos authentication disabled)

  yarn-session.sh -nm "*session-name***" -d**

- Security cluster (Kerberos authentication enabled)
  - If the **flink.keystore** and **flink.truststore** file paths are relative paths:

    Run the following command in the directory at the same level as **ssl** to start the session. **ssl/** is a relative path.

    cd /opt/hadoopclient/Flink/flink/conf/

    yarn-session.sh -t ssl/ -nm "*session-name***" -d**

    ```
    ...
    Cluster started: Yarn cluster with application id application_1624937999496_0017
    JobManager Web Interface: http://192.168.1.150:32261
    ```

  - If the **flink.keystore** and **flink.truststore** file paths are absolute paths:

    Run the following command to start a session:

    cd /opt/hadoopclient/Flink/flink/conf/

    yarn-session.sh -nm "*session-name***" -d**

Step 5 For a security cluster, run the following command to perform user authentication. If Kerberos authentication is not enabled for the current cluster, you do not need to run this command.

kinit *Username*

Step 6 Explicitly add the OBS file system to be accessed in the Flink command line.

echo -e 'test' >/tmp/test

hdfs dfs -mkdir -p obs://*Parallel file system name***/tmp/flinkjob**

hdfs dfs -put /tmp/test/ obs://*Parallel file system name***/tmp/flinkjob/**

flink run Client installation directory**/Flink/flink/examples/batch/WordCount.jar -input obs://***Parallel file system name***/tmp/flinkjob/test -output obs://***Parallel file system name***/tmp/flinkjob/output**

----End

NOTE

- Flink jobs are running on Yarn. Before configuring Flink to interconnect with the OBS file system, ensure that the interconnection between Yarn and the OBS file system is normal.
- *Name of the OBS parallel file system/File name*: The OBS file path must be written to the directory level.
- If Kerberos authentication has been enabled (security mode) for the cluster, grant the **Read** and **Write** permissions on OBS paths to component users in Ranger by referring to **Ranger Permission Configuration**.

Ranger Permission Configuration

Step 1 Log in to FusionInsight Manager and choose **System** > **Permission** > **User Group**. On the displayed page, click **Create User Group**.

Step 2 Create a user group without a role, for example, **obs_flink**, and bind the user group to the corresponding user.

Step 3 Log in to the Ranger management page as the **rangeradmin** user.

Step 4 On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area.

Step 5 Click **Add New Policy** to add the **Read** and **Write** permissions on OBS paths to the user group created in **Step 2**. If there are no OBS paths, create one in advance (wildcard character * is not allowed).

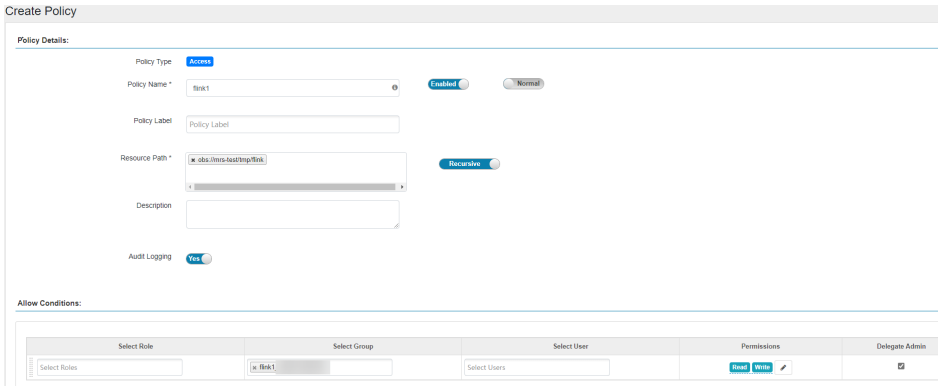

----End

5.3.3.3 Interconnecting Spark with OBS

Interconnecting with OBS

In an MRS cluster, **Location** can be set to an OBS file system path during Spark table creation and Spark can connect to OBS through Hive Metastore.

- Setting the location to an OBS path during table creation:
  a. Log in to the node where the client is installed as the client installation user and access the **spark-sql** client.

     cd *Client installation directory*

     kinit *Component operation user*

     spark-sql --master yarn

  b. Set **Location** to the OBS file system path when creating a table.

     For example, to create a table named **test** whose **Location** is **obs://obs-test/test/***Database name***/***Table name*, run the following command:

     create external table testspark(name string) location "obs://obs-test/test/*Database name***/***Table name***";**
- Interconnecting Spark with OBS through Hive Metastore:
  a. Complete the configurations by referring to **Interconnecting Hive with OBS using MetaStore**.
  b. Log in to FusionInsight Manager, choose **Cluster** > **Services** > **Spark** and choose **Configurations** > **All Configurations**.
  c. In the navigation pane on the left, choose **SparkResource** > **Customization**. In the custom configuration items, add **spark.sql.warehouse.location.first** to the **custom** parameter and set its value to **true**.

Figure 5-42 spark.sql.warehouse.location.first configuration

d. In the navigation pane on the left, choose **JDBCServer** > **Customization**. In the custom configuration items, add **spark.sql.warehouse.location.first** to the **custom** parameter and set its value to **true**.

Figure 5-43 spark.sql.warehouse.location.first configuration

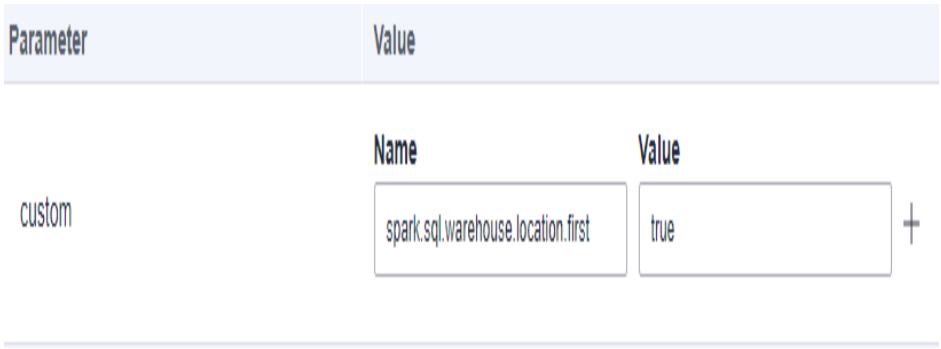

e. Click **Save** to save the configuration. Click the **Dashboard** tab choose **More** > **Restart Service**, enter the password, click **OK**, and click **OK** again to restart Spark.

f. After Spark is restarted, choose **More** > **Download Client** to download and install the Spark client again. Then, go to **g**.

If you do not download and install the client again, you can perform the following steps to update the Spark client configuration file (assume that the client directory is **/opt/client**):

i. Log in to the node where the Spark client is deployed as user **root** and switch to the client installation directory.

cd /opt/client

ii. Run the following command to modify **hive-site.xml** in the configuration file directory of the Spark client:

vi Spark/spark/conf/hive-site.xml

Change the value of **hive.metastore.warehouse.dir** to the corresponding OBS path, for example, **obs://hivetest/user/hive/warehouse/**.

```
<property>
<name>hive.metastore.warehouse.dir</name>
<value>obs://hivetest/user/hive/warehouse/<value>
</property>
```


iii. Run the following command to modify the **spark-defaults.conf** file in the configuration file directory of the Spark client and set **spark.sql.warehouse.location.first** to **true**:

vi Spark/spark/conf/spark-defaults.conf

g. Configure the OBS directory permission for the component operation user in clusters with Kerberos authentication enabled by referring to **Configuring Ranger Permissions**.

h. Go to the SparkSQL CLI and spark-beeline, create a table, and check whether the location of the table is an OBS path.

source bigdata_env

kinit *Service user* (skip this step for normal clusters)

- Go to the SparkSQL CLI.

  spark-sql

  create table *d*(**a int**);

  desc formatted *d*;

  As shown in the following figure, the location of table **d** is in the specified OBS path.

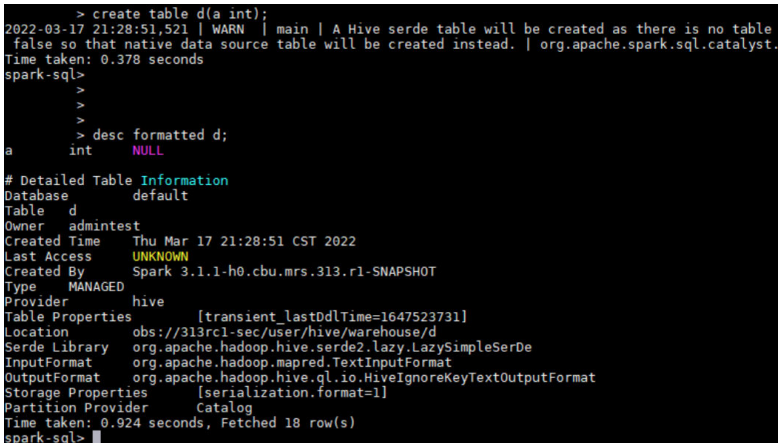

- Go to spark-beeline.

  spark-beeline

  create table *e*(**a int**);

  desc formatted *e*;

  As shown in the following figure, the location of table **e** is in the specified OBS path.


```
0: jdbc:hive2://BMS-ARM-node-master3Xxky:22550/> create table e(a int) ;
+---------+
| Result  |
+---------+
+---------+
No rows selected (0.763 seconds)
0: jdbc:hive2://BMS-ARM-node-master3Xxky:22550/> desc formatted e;
+-------------------------------+----------------------------------------------------+----------+
|           col_name            |                     data_type                      | comment  |
+-------------------------------+----------------------------------------------------+----------+
| a                             | int                                                | NULL     |
|                               |                                                    |          |
| # Detailed Table Information  |                                                    |          |
| Database                      | default                                            |          |
| Table                         | e                                                  |          |
| Owner                         | admintest                                          |          |
| Created Time                  | Fri Mar 18 09:37:17 CST 2022                       |          |
| Last Access                   | UNKNOWN                                            |          |
| Created By                    | Spark 3.1.1-h0.cbu.mrs.313.r1-SNAPSHOT             |          |
| Type                          | MANAGED                                            |          |
| Provider                      | hive                                               |          |
| Table Properties              | [transient_lastDdlTime=1647567437]                 |          |
| Location                      | obs://313rc1-sec/user/hive/warehouse/e             |          |
| Serde Library                 | org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe |          |
| InputFormat                   | org.apache.hadoop.mapred.TextInputFormat           |          |
| OutputFormat                  | org.apache.hadoop.hive.ql.io.HiveIgnoreKeyTextOutputFormat |          |
| Storage Properties            | [serialization.format=1]                           |          |
| Partition Provider            | Catalog                                            |          |
+-------------------------------+----------------------------------------------------+----------+
18 rows selected (1.418 seconds)
```

Configuring Ranger Permissions

Step 1 Log in to FusionInsight Manager and choose **System** > **Permission** > **User Group**. On the displayed page, click **Create User Group**.

Step 2 Create a user group without a role, for example, **obs_spark**, and bind the user group to the corresponding user.

Step 3 Log in to the Ranger management page as the **rangeradmin** user.

Step 4 On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area.

Step 5 Click **Add New Policy** to add the **Read** and **Write** permissions on OBS paths to the user group created in **Step 2**. If there are no OBS paths, create one in advance (wildcard character * is not allowed).

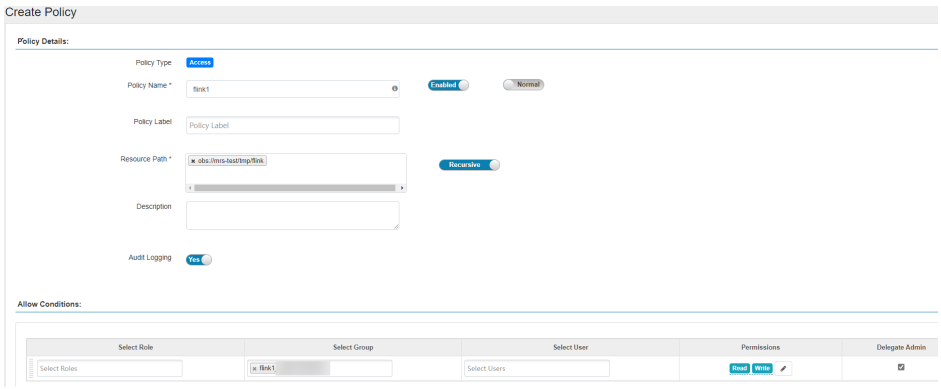

----End

NOTE

- Cascading authorization is not supported for view tables.
- Cascading authorization can be performed only on databases and tables, and cannot be on partitions. If a partition path is not in the table path, you need to manually authorize the partition path.
- Cascading authorization for Deny Conditions in the Hive Ranger policy is not supported. That is, the Deny Conditions permission only restricts the table permission and cannot generate the permission of the HDFS storage source.
- The permission of the HDFS Ranger policy is prior to that of the HDFS storage source generated by cascading authorization. If the HDFS Ranger permission has been set for the HDFS storage source of the table, the cascading permission does not take effect.

Configuring Permissions for CDL Service Users

If Kerberos authentication is enabled for the cluster (the cluster is in security mode) and you need to store real-time data to OBS after the interconnection, perform the following operations to grant the **Read** and **Write** permissions on the corresponding OBS path to the specific user:

Step 1 Log in to FusionInsight Manager and choose **System** > **Permission** > **User Group**. On the displayed page, click **Create User Group**.

Step 2 Create a user group without a role, for example, **obs_cdl**, and bind the user group to the corresponding CDL service user, for example, **cdluser**.

Step 3 Log in to the Ranger management page as the **rangeradmin** user.

Step 4 On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area.

Step 5 Click **Add New Policy** to add the **Read** and **Write** permissions on OBS paths to the created user group. If there are no OBS paths, create one in advance (wildcard character * is not allowed).

The following figure shows the configurations needed for adding the **Read** and **Write** permissions on **obs://***OBS parallel file system name***/cdldata** to user group **obs_cdl**.

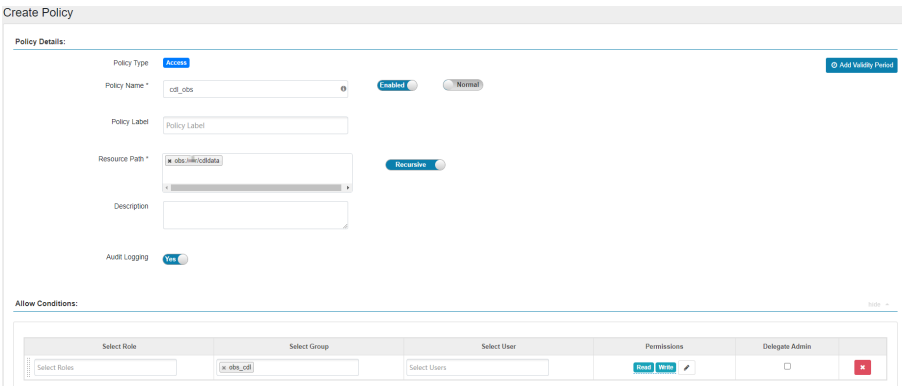

----End

5.3.3.4 Interconnecting Hudi with OBS

Interconnecting with OBS

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

Step 2 Run the following commands to configure environment variables:

source *Client installation directory*/**bigdata_env**

source *Client installation directory*/**Hudi/component_env**

Step 3 Modify the configuration file.

vim *Client installation directory*/**Hudi/hudi/conf/hdfs-site.xml**

```
<property>
<name>dfs.namenode.acls.enabled</name>
<value>false</value>
</property>
```

Step 4 If Kerberos authentication has been enabled(security mode) for the cluster, run the following command to perform authentication as a user who has the Read and Write permissions on the corresponding OBS path. If Kerberos authentication has not been enabled (normal mode) for the cluster, you do not need to run this command.

kinit *Username*

Step 5 Start **spark-shell** and run the following commands to create a COW table and store it to OBS:

spark-shell --master yarn

import org.apache.hudi.QuickstartUtils._

import scala.collection.JavaConversions._

import org.apache.spark.sql.SaveMode._

import org.apache.hudi.DataSourceReadOptions._

import org.apache.hudi.DataSourceWriteOptions._

import org.apache.hudi.config.HoodieWriteConfig._

val tableName = "hudi_cow_table"

val basePath = "obs://testhudi/cow_table/"

val dataGen = new DataGenerator

val inserts = convertToStringList(dataGen.generateInserts(10))

val df = spark.read.json(spark.sparkContext.parallelize(inserts, 2))

df.write.format("org.apache.hudi").

options(getQuickstartWriteConfigs).

option(PRECOMBINE_FIELD_OPT_KEY, "ts").

option(RECORDKEY_FIELD_OPT_KEY, "uuid").

option(PARTITIONPATH_FIELD_OPT_KEY, "partitionpath").

option(TABLE_NAME, tableName).

mode(Overwrite).

save(basePath);

NOTE

"obs://testhudi/cow_table/" is the OBS path, and **testhudi** is the name of the OBS parallel system file. Change them based on site requirements.

Step 6 Use DataSource to check whether the table is created and whether the data is normal.

val roViewDF = spark.

read.

format("org.apache.hudi").

load(basePath + "/*/*/*/*")

roViewDF.createOrReplaceTempView("hudi_ro_table")

spark.sql("select * from hudi_ro_table").show()

Step 7 Run the **:q** command to exit the **spark-shell** CLI.

----End

Configuring Ranger Permissions

Step 1 Log in to FusionInsight Manager and choose **System** > **Permission** > **User Group**. On the displayed page, click **Create User Group**.

Step 2 Create a user group without a role, for example, **obs_hudi**, and bind the user group to the corresponding user.

Step 3 Log in to the Ranger management page as the **rangeradmin** user.

Step 4 On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area.

Step 5 Click **Add New Policy** to add the **Read** and **Write** permissions on OBS paths to the user group created in **Step 2**. If there are no OBS paths, create one in advance (wildcard character * is not allowed).

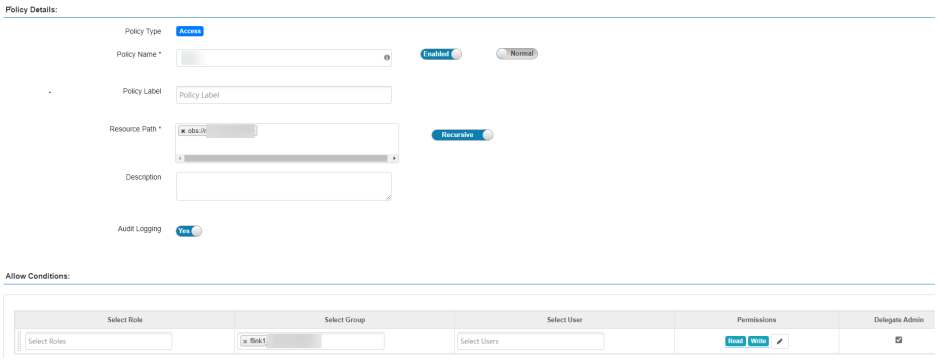

----End

5.3.3.5 Interconnecting HetuEngine with OBS

Interconnecting with OBS

In an MRS cluster, **Location** can be set to an OBS file system path during HetuEngine table creation and HetuEngine can connect to OBS through Hive Metastore.

- Setting **Location** to the OBS file system path when creating a table
  a. If a HetuEngine compute instance is running, restart it.

     Log in to FusionInsight Manager as a user who has permission to access the HetuEngine web UI. Choose **Cluster** > **Services** > **HetuEngine**. In the **Basic Information** area in the Dashboard tab, click the link next to **HSConsole WebUI**. On the displayed HSConsole page, click **Compute Instance**. In the instance list, click **Restart** in the **Operation** column and operate as prompted.

  b. Log in to the node where the HetuEngine service client is located as the client installation user and run the following command:

     source *Client installation directory*/**bigdata_env**

  c. Log in to the HetuEngine client based on the cluster authentication mode.

     - Kerberos authentication has been enabled for the cluster (security mode): Run the following command to complete user authentication and log in to the HetuEngine client:

       kinit *User performing HetuEngine operations*

       hetu-cli --catalog hive --tenant default --schema default

       For details about how to assign permissions to users in the Ranger, see **Configuring Ranger Permissions**.

     - Kerberos authentication is not enabled for the cluster (normal mode): Run the following command to log in to the HetuEngine client:

       hetu-cli --catalog hive --tenant default --schema default --user *User performing HetuEngine operations*

  d. Set **Location** to the OBS file system path when creating a table.

create table test(name string) with (location = 'obs://*Name of the OBS parallel file system***/user/hive/warehouse/test');**

- Interconnecting with OBS through Hive Metastore
  a. Complete the configurations by referring to **Interconnecting Hive with OBS using MetaStore**.
  b. Log in to FusionInsight Manager, choose **Cluster** > **Services** > **HetuEngine**. On the displayed page, choose **More** > **Synchronize Configuration**. After the synchronization is complete, choose **More** > **Synchronize Configuration** again and then restart the HetuEngine service as prompted.

  NOTICE

  If a HetuEngine compute instance is running, stop it before restarting the service. After the service is restarted, start this compute instance.

  c. No location needs to be specified when you log in to the HetuEngine client to create a schema or table. The schema or table is stored on OBS by default.

Configuring Ranger Permissions

For HetuEngine clusters with Kerberos authentication enabled (security mode), the methods to grant Ranger permission are the same for both storage-compute decoupled architecture and and storage-compute coupled architecture.

5.3.3.6 Interconnecting HDFS with OBS

Interconnecting with OBS

Step 1 Log in to the node on which the HDFS client is installed as a client installation user.

Step 2 Run the following command to switch to the client installation directory:

cd *Client installation directory*

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If the cluster is in security mode, run the following command to authenticate the user. The user must have the read and write permissions on the OBS directory. Skip user authentication for normal clusters.

kinit *User performing HDFS operations*

Step 5 Explicitly add the OBS file system to be accessed in the HDFS command line.

The following commands are examples.

- Access the OBS file system.
  hdfs dfs -ls obs://*OBS parallel file system name*/*Path*

- Create a directory in the OBS file system.

  hdfs dfs -mkdir obs://*OBS parallel file system name/hadoop*
- Upload the **/opt/test.txt** file on the client node to the **obs://***OBS parallel file system name***/hadoop** directory.

  hdfs dfs -put /opt/test.txt obs://*OBS parallel file system name***/hadoop**

----End

NOTE

If a large number of logs are printed in the OBS file system, the read and write performance may be affected. You can adjust the log level of the OBS client as follows:

cd *Client installation directory***/HDFS/hadoop/etc/hadoop**

vi log4j.properties

Add the OBS log level configuration to the file as follows:

log4j.logger.org.apache.hadoop.fs.obs=WARN

log4j.logger.com.obs=WARN

```
[root@node-master1AuKK hadoop]# tail -4 log4j.properties
log4j.logger.org.apache.commons.beanutils=WARN

log4j.logger.org.apache.hadoop.fs.obs=WARN
log4j.logger.com.obs=WARN
[root@node-master1AuKK hadoop]#
```

Configuring Ranger Permissions

Step 1 Log in to FusionInsight Manager and choose **System** > **Permission** > **User Group**. On the displayed page, click **Create User Group** to create a user group without any roles, for example, **obs_hadoop**.

Step 2 Back to FusionInsight Manager and choose **System** > **Permission** > **User**. On the displayed page, click **Create User** to create a user that is associated only with the **obs_hadoop** user group, for example, **hadoopuser**.

Step 3 Log in to the Ranger management page as the **rangeradmin** user.

Step 4 On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area.

Step 5 Click **Add New Policy** and add the **Read** and **Write** permissions on the desired OBS paths to the created user group.

The following figure shows the configurations needed for adding the **Read** and **Write** permissions on **obs://***OBS parallel file system name***/hadoop** to user group **obs_hadoop**.

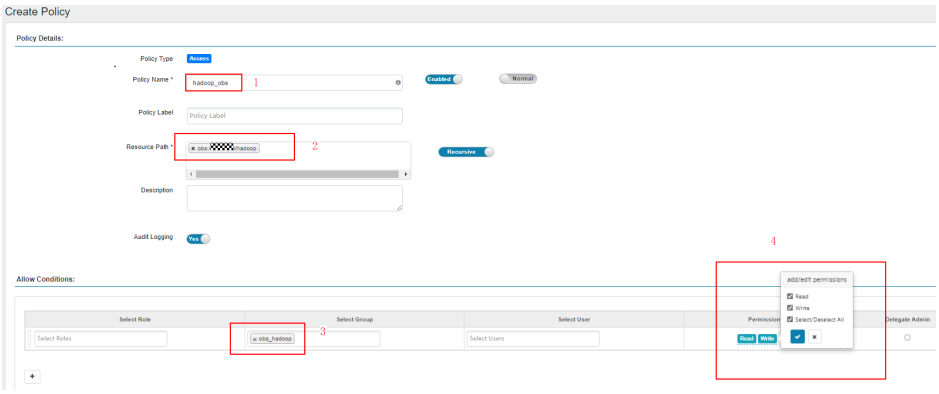

----End

5.3.3.7 Interconnecting Yarn with OBS

Interconnecting with OBS

Step 1 Log in to the node on which the Yarn client is installed as a client installation user.

Step 2 Run the following command to switch to the client installation directory.

cd *Client installation directory*

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 If the cluster is in security mode, run the following command to authenticate the user. The user must have the read and write permissions on the OBS directory. Skip user authentication for normal clusters.

kinit *User performing HDFS operations*

Step 5 Explicitly add the OBS file system to be accessed in the Yarn command line.

- Access the OBS file system.

  hdfs dfs -ls obs://*OBS parallel file system name***/***Path*

- Create a directory in the OBS file system.

  hdfs dfs -mkdir obs://*OBS parallel file system name/hadoop1*

- Execute the Yarn task to access OBS.

  yarn jar *Client installation directory***/HDFS/hadoop/share/hadoop/mapreduce/hadoop-mapreduce-examples-*.jar pi -Dmapreduce.job.hdfs-servers=NAMESERVICE -fs obs://***OBS parallel file system name* **1 1**

NAMESERVICE indicates the NameService in HDFS. The default value is **hdfs://hacluster**. If there are multiple NameServices, separate them with **,**.

The following command is an example:

yarn jar /opt/hadoopclient/HDFS/hadoop/share/hadoop/mapreduce/hadoop-mapreduce-examples-*.jar pi -Dmapreduce.job.hdfs-servers=hdfs://hacluster -fs obs://bucketname 1 1

- Run the following command to write data to OBS:

yarn jar *Client installation directory*/**HDFS/hadoop/share/hadoop/mapreduce/hadoop-mapreduce-examples-*.jar teragen 100 obs://***OBS parallel file system name***/hadoop1/teragen1**

- Run the following command to copy data from OBS to HDFS:

  hadoop distcp obs://*OBS parallel file system name***/hadoop1/teragen1 /tmp**

----End

NOTE

If a large number of logs are printed in the OBS file system, the read and write performance may be affected. You can adjust the log level of the OBS client as follows:

cd *Client installation directory*/**Yarn/config**

vi log4j.properties

Add the OBS log level configuration to the file. (If an application uses the built-in **log4j.properties** file, add the same configuration.)

log4j.logger.org.apache.hadoop.fs.obs=WARN

log4j.logger.com.obs=WARN

```
[root@node-master1AuKK config]# tail -4 log4j.properties
log4j.logger.org.apache.commons.beanutils=WARN

log4j.logger.org.apache.hadoop.fs.obs=WARN
log4j.logger.com.obs=WARN
[root@node-master1AuKK config]#
```

Configuring Ranger Permissions

Step 1 Log in to FusionInsight Manager and choose **System** > **Permission** > **User Group**. On the displayed page, click **Create User Group** to create a user group without any roles, for example, **obs_hadoop1**.

Step 2 Back to FusionInsight Manager and choose **System** > **Permission** > **User**. On the displayed page, click **Create User** to create a user that is associated with the **obs_hadoop1** user group and the **default** role, for example, **hadoopuser1**.

Step 3 Log in to the Ranger management page as the **rangeradmin** user.

Step 4 On the home page, click component plug-in name **OBS** in the **EXTERNAL AUTHORIZATION** area.

Step 5 Click **Add New Policy** and add the **Read** and **Write** permissions on the desired OBS paths to the user group created in **Step 1**.

The following figure shows the configurations needed for adding the **Read** and **Write** permissions on **obs://***OBS parallel file system name***/hadoop1** to user group **obs_hadoop1**.

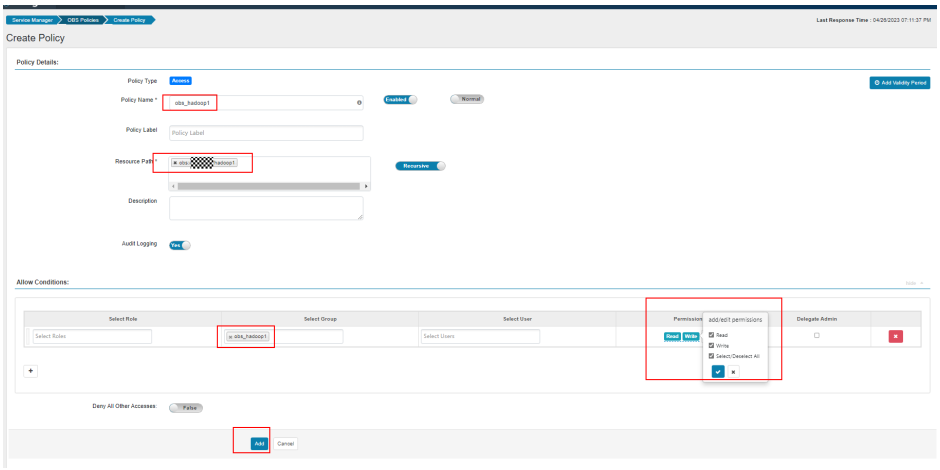

----End

5.3.3.8 Interconnecting MapReduce with OBS

Interconnecting with OBS

Step 1 Log in to FusionInsight Manager, choose **Cluster** > **Services** > **MapReduce** and choose **Configurations** > **All Configurations**. In the navigation tree, choose **MapReduce** > **Customization**. In the customized configuration items, add the configuration item **mapreduce.jobhistory.always-scan-user-dir** to **core-site.xml** and set the parameter to **true**.

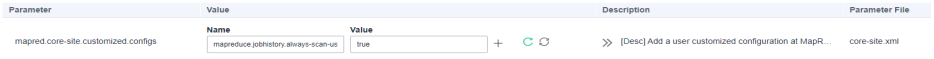

Step 2 Save the configurations and restart the MapReduce service.

----End

6 Accessing Web Pages of Open Source Components Managed in MRS Clusters

6.1 Web UIs of Open Source Components

Scenario

Web UIs of different components are created and hosted on the Master or Core nodes in the MRS cluster by default. You can view information about the components on these web UIs.

Procedure for accessing the web UIs of open-source component:

1. Select an access method.

   MRS provides the following methods for accessing the web UIs of open-source components:

   - **EIP-based Access**: This method is recommended because it is easy to bind an EIP to a cluster.
   - **Access Using a Windows ECS**: Independent ECSs need to be created and configured.
   - **Creating an SSH Channel for Connecting to an MRS Cluster and Configuring the Browser**: Use this method when the user and the MRS cluster are on different networks.

2. Access the web UIs. For details, see **Table 6-1**.

Web UIs

NOTE

For clusters with Kerberos authentication enabled, user **admin** does not have the management permission on each component. To access the web UI of each component, create a user who has the management permission on the corresponding component.

Table 6-1 Web UI addresses of open-source components

Cluster Type	Web UI Type	Address (Earlier than MRS 3.*x*)	Address (MRS 3.*x* and Later)
All Types	MRS Manager	• Method 1: Clusters of all versions **https://***Floating IP address of Manager***:28443/web** **NOTE** 1. Ensure that the local host can communicate with the MRS cluster. 2. Log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of MRS Manager. Record the value of **inet**. If the floating IP address of MRS Manager cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the cluster manager IP address of the Master node. • Method 2: Earlier than MRS 3.*x* https://<*EIP*>:9022/mrsmanager?locale=en-us For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.	• Method 1: Clusters of all versions **https://***Floating IP address of Manager***:28443/web** **NOTE** 1. Ensure that the local host can communicate with the MRS cluster. 2. Log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of MRS Manager. Record the value of **inet**. If the floating IP address of MRS Manager cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one master node, query and record the cluster manager IP address of the master node. • Method 2: MRS 3.*x* or later For details, see **Accessing FusionInsight Manager (MRS 3.x or Later)**.
Analysis cluster	HDFS NameNode	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **HDFS** > **NameNode Web UI** > **NameNode (Active)**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **HDFS** > **NameNode Web UI** > **NameNode (***Host name***, Active)**.
	HBase HMaster	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **HBase** > **HMaster Web UI** > **HMaster (Active)**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **HBase** > **HMaster Web UI** > **HMaster (***Host name***, Active)**.

Cluster Type	Web UI Type	Address (Earlier than MRS 3.*x*)	Address (MRS 3.*x* and Later)
	MapReduce JobHistory Server	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **MapReduce** > **JobHistoryServer Web UI** > **JobHistoryServer**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **MapReduce** > **JobHistoryServer Web UI** > **JobHistoryServer (***Host name***, Active)**.
	YARN ResourceManager	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **Yarn** > **ResourceManager Web UI** > **ResourceManager (Active)**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **Yarn** > **ResourceManager Web UI** > **ResourceManager (***Host name***, Active)**.
	Spark JobHistory	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **Spark** > **Spark Web UI** > **JobHistory**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **Spark2x** > **Spark2x Web UI** > **JobHistory2x (***Host name***, Active)**.
	Hue	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **Hue** > **Hue Web UI** > **Hue (Active)**. Loader is a graphical data migration management tool based on the open-source Sqoop web UI, and its interface is hosted on the Hue web UI.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **Hue** > **Hue Web UI** > **Hue (***Host name***, Active)**. Loader is a graphical data migration management tool based on the open-source Sqoop web UI, and its interface is hosted on the Hue web UI.
	Tez	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **Tez** > **Tez Web UI** > **TezUI**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **Tez** > **Tez Web UI** > **TezUI (***Host name***)**.
	Presto	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **Presto** > **Presto Web UI** > **Coordinator (Active)**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **Presto** > **Coordinator Web UI** > **Coordinator(Coordinator)**.
	Ranger	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **Ranger** > **Ranger Web UI** > **RangerAdmin (Active)**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **Ranger** > **Ranger Web UI** > **RangerAdmin**.

Cluster Type	Web UI Type	Address (Earlier than MRS 3.*x*)	Address (MRS 3.*x* and Later)
Stream processing cluster	Storm	Earlier than MRS 3.*x*: On the cluster details page, choose **Components** > **Storm** > **Storm Web UI** > **UI**.	MRS 3.*x* or later: On the Manager homepage, choose **Cluster** > **Services** > **Storm** > **Storm Web UI** > **UI (***Host name***)**.

6.2 Common Ports of Components

Scenario

When you **buy a custom cluster** of an LTS version, you can customize the component port. If you do not want to customize a port, an open source port is used.

- **Open source**: Find the default port of the component in the Default Open Source Port column of the following table.
- **Custom**: Find the default port of the component in the Default Custom Port column of the following table.
- If there is only the Default Port column, the open source port of the component is the same as the default custom port.

If the cluster is not of an LTS version, the **Component Port** parameter is unavailable and only an open source port can be used. For details, see the Default Open Source Port or Default Port column.

Common HBase Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
hbase.master.port	16000	21300	HMaster RPC port. This port is used to connect the HBase client to HMaster. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.master.info.port	16010	21301	HMaster HTTPS port. This port is used by the remote web client to connect to the HMaster UI. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.regionserver.port	16020	21302	RegoinServer (RS) RPC port. This port is used to connect the HBase client to RegionServer. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
hbase.regionserver.info.port	16030	21303	HTTPS port of the Region server. This port is used by the remote web client to connect to the RegionServer UI. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.thrift.info.port	9095	21304	Thrift Server listening port of Thrift Server This port is used for: Listening when the client is connected **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hbase.regionserver.thrift.port	9090	21305	Thrift Server listening port of RegionServer This port is used for: Listening when the client is connected to the RegionServer **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
hbase.rest.info.port	8085	21308	Port of the RegionServer RESTServer native web page
-	21309	21309	REST port of RegionServer RESTServer

Common HDFS Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.namenode.rpc.port	• 9820 (versions earlier than MRS 3.*x*) • 8020 (MRS 3.*x* and later)	25000	NameNode RPC port This port is used for: 1. Communication between the HDFS client and NameNode 2. Connection between the DataNode and NameNode **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.namenode.http.port	9870	25002	HDFS HTTP port (NameNode) This port is used for: 1. Point-to-point NameNode checkpoint operations 2. Connecting the remote web client to the NameNode UI **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.namenode.https.port	9871	25003	HDFS HTTPS port (NameNode) This port is used for: 1. Point-to-point NameNode checkpoint operations 2. Connecting the remote web client to the NameNode UI **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes
dfs.datanode.ipc.port	9867	25008	IPC server port of DataNode This port is used for: Connection between the client and DataNode to perform RPC operations. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes
dfs.datanode.port	9866	25009	DataNode data transmission port This port is used for: 1. Transmitting data from HDFS client from or to the DataNode 2. Point-to-point DataNode data transmission **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.datanode.http.port	9864	25010	DataNode HTTP port This port is used for: Connecting to the DataNode from the remote web client in security mode **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.datanode.https.port	9865	25011	HTTPS port of DataNode This port is used for: Connecting to the DataNode from the remote web client in security mode **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
dfs.JournalNode.rpc.port	8485	25012	RPC port of JournalNode This port is used for: Client communication to access multiple types of information **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
dfs.journalnode.http.port	8480	25013	JournalNode HTTP port This port is used for: Connecting to the JournalNode from the remote web client in security mode **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes
dfs.journalnode.https.port	8481	25014	HTTPS port of JournalNode This port is used for: Connecting to the JournalNode from the remote web client in security mode **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes
httpfs.http.port	14000	25018	Listening port of the HttpFS HTTP server This port is used for: Connecting to the HttpFS from the remote REST API **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes

Common HetuEngine Ports

The protocol type of the port in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
server.port (HSBroker)	29860	29860	Specifies the port number that HSBroker listens to.
server.port (HSConsole)	29880	29880	Specifies the port number that HSConsole listens to.
server.port (HSFabric)	29900	29900	Specifies the port number that HSFabric listens to, which is used for cross-domain connections
gateway.port	29902	29902	Specifies the port number that HSFabric listens to, which is used for JDBC connections

Common Hive Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
templeton.port	9111	21055	Port used for WebHCat to provide the REST service This port is used for: Communication between the WebHCat client and WebHCat server • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
hive.server2.thrift.port	10000	21066	Port for HiveServer to provide Thrift services This port is used for: Communication between the HiveServer and HiveServer client • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hive.metastore.port	9083	21088	Port for MetaStore to provide Thrift services This port is used for: Communication between the MetaStore client and MetaStore, that is, communication between HiveServer and MetaStore. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
hive.server2.webui.port	10002	-	Web UI port of Hive This port is used for: HTTPS/HTTP communication between Web requests and the Hive UI server

Common Hue Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
HTTP_PORT	8888	21200	Port for Hue to provide HTTPS services This port is used to provide web services in HTTPS mode, which can be changed. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Kafka Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
port	9092	21005	Port for a broker to receive data and obtain services
ssl.port	9093	21008	SSL port used by a broker to receive data and obtain services
sasl.port	21007	21007	SASL security authentication port provided by a broker, which provides the secure Kafka service
sasl-ssl.port	21009	21009	Port used by a broker to provide encrypted service based on the SASL and SSL protocols

Common Loader Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Port	Port Description
LOADER_HTTPS_PORT	21351	This port is used to provide REST APIs for configuration and running of Loader jobs. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Manager Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Port	Port Description
-	8080	Port provided by WebService for user access This port is used to access the web UI over HTTP. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes
-	28443	Port provided by WebService for user access This port is used to access the web UI over HTTPS. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes

Common MapReduce Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
mapreduce.jobhistory.webapp.port	19888	26012	Web HTTP port of the JobHistory server This port is used for: viewing the web page of the JobHistory server **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes

Parameter	Default Open Source Port	Default Custom Port	Port Description
mapreduce.jobhistory.port	10020	26013	Port of the JobHistory server This port is used for: 1. Task data restoration in the MapReduce client 2. Obtaining task report in the Job client **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
mapreduce.jobhistory.webapp.https.port	19890	26014	Web HTTPS port of the JobHistory server This port is used to view the web page of the JobHistory server. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common Spark Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
hive.server2.thrift.port	22550	22550	JDBC thrift port This port is used for: Socket communication between Spark2.1.0 CLI/JDBC client and server **NOTE** If **hive.server2.thrift.port** is occupied, an exception indicating that the port is occupied is reported. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes
spark.ui.port	4040	22950	Web UI port of JDBC This port is used for: HTTPS/HTTP communication between Web requests and the JDBC Server Web UI server **NOTE** The system verifies the port configuration. If the port is invalid, the value of the port plus 1 is used till the calculated value is valid. (A maximum number of 16 attempts are allowed. The number of attempts is specified by **spark.port.maxRetries**.) ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes
spark.history.ui.port	18080	22500	JobHistory Web UI port This port is used for: HTTPS/HTTP communication between Web requests and Spark2.1.0 History Server **NOTE** The system verifies the port configuration. If the port is invalid, the value of the port plus 1 is used till the calculated value is valid. (A maximum number of 16 attempts are allowed. The number of attempts is specified by **spark.port.maxRetries**.) ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes

Common Storm Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
nimbus.thrift.port	6627	29200	Port for Nimbus to provide thrift services
supervisor.slots.ports	6700,6701,6702,6703	29200-29499	Port for receiving service requests that are forwarded from other servers
logviewer.https.port	29248	29248	Port for LogViewer to provide HTTPS services
ui.https.port	29243	29243	Port for Storm UI to provide HTTPS services (ui.https.port)

Common Yarn Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
yarn.resourcemanager.webapp.port	8088	26000	Web HTTP port of the ResourceManager service
yarn.resourcemanager.webapp.https.port	8090	26001	Web HTTPS port of the ResourceManager service This port is used to access the Resource Manager web applications in security mode. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes
yarn.nodemanager.webapp.port	8042	26006	NodeManager Web HTTP port

Parameter	Default Open Source Port	Default Custom Port	Port Description
yarn.nodemanager.webapp.https.port	8044	26010	NodeManager Web HTTPS port This port is used for: Accessing the NodeManager web application in security mode **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes

Common ZooKeeper Ports

The protocol type of all ports in the table is TCP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
clientPort	2181	24002	ZooKeeper client port This port is used for: Connection between the ZooKeeper client and server. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. ● Is the port enabled by default during the installation: Yes ● Is the port enabled after security hardening: Yes

Common Kerberos Ports

The protocol type of all ports in the table is TCP and UDP.

Parameter	Default Port	Port Description
kdc_ports	21732	Kerberos server port This port is used for performing Kerberos authentication for components. This parameter may be used during the configuration of mutual trust between clusters. **NOTE** The port ID is a recommended value and is specified based on the product. The port range is not restricted in the code. • Is the port enabled by default during the installation: Yes • Is the port enabled after security hardening: Yes

Common OpenTSDB Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
tsd.network.port	4242	Web UI port of OpenTSDB This port is used for: HTTPS/HTTP communication between web requests and the OpenTSDB UI server

Common Tez Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
tez.ui.port	28888	Web UI port of Tez

Common KafkaManager Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
kafka_manager_port	9099	Web UI port of KafkaManager

Common Presto Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
http-server.http.port	7520	HTTP port for Presto coordinator to provide services to external systems
http-server.https.port	7521	HTTPS port for Presto coordinator to provide services to external systems
http-server.http.port	7530	HTTP port for Presto worker to provide services to external systems
http-server.https.port	7531	HTTPS port for Presto worker to provide services to external systems

Common Flink Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
jobmanager.web.port	32261-32325	Web UI port of Flink This port is used for: HTTP/HTTPS communication between the client web requests and Flink server

Common ClickHouse Ports

The protocol type of the port in the table is TCP and HTTP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
interserver_http_port	9009	9009	HTTP port for the communication between ClickHouse servers.
interserver_https_port	9010	9010	HTTPS port for the communication between ClickHouse servers.
http_port	8123	8123	Port for connecting to the ClickHouse server through HTTP.
https_port	8443	8443	Port for connecting to the ClickHouse server through HTTPS.
tcp_port	9000	9000	Port for connecting the client to the ClickHouse server through TCP.
tcp_port_secure	9440	9440	Port for connecting the client to the ClickHouse server through TCP SSL.
lb_tcp_port	21424	21424	TCP port listened by ClickHouseBalancer
lb_http_port	21425	21425	HTTP port listened by ClickHouseBalancer
lb_https_port	21426	21426	HTTPS port listened by ClickHouseBalancer
lb_tcp_secure_port	21428	21428	TCP SSL port listened by ClickHouseBalancer

Common Impala Ports

The protocol type of the port in the table is TCP.

Parameter	Default Port	Port Description
--beeswax_port	21000	Port for impala-shell communication
--hs2_port	21050	Port for Impala application communication

Parameter	Default Port	Port Description
--hs2_http_port	28000	Port used by Impala to provide the HiveServer2 protocol for external systems

Common Doris Ports

The protocol type of the port in the table is TCP and HTTP.

Parameter	Default Open Source Port	Default Custom Port	Port Description
http_port	8030	29980	HTTP port of the FE service
https_port	8050	29991	HTTPS port of the FE service
query_port	9030	29982	Port used by the Doris FE to query connections through the MySQL protocol
rpc_port	9020	29981	Thrift Server port of the FE service
be_port	9060	29984	Thrift Server Port on BE for receiving requests from FE
brpc_port	8060	29987	BRPC port on BE, which is used for communications between BE instances.
heartbeat_service_port	9050	29985	Thrift heartbeat service port on BE, which is used to receive heartbeat messages from FE
webserver_port	8040	29986	HTTP server port on BE
broker_ipc_port	8000	29990	Thrift Server communication port on Broker, which is used to receive requests.
single_replica_load_brpc_port	9070	29988	RPC port used for the communication between the master and slave replicas to import single-replica data
single_replica_load_download_port	8050	29989	Port used by the slave replica to download data files from the master replica through HTTP for single-copy data import

6.3 Access Through Direct Connect

MRS allows you to access MRS clusters using Direct Connect. Direct Connect is a high-speed, low-latency, stable, and secure dedicated network connection that connects your local data center to an online cloud VPC. It extends online cloud services and existing IT facilities to build a flexible, scalable hybrid cloud computing environment.

Prerequisites

Direct Connect is available, and the connection between the local data center and the online VPC has been established.

Accessing an MRS Cluster Using Direct Connect

Step 1 Log in to the MRS console.

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

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

Step 4 Set **Access Mode** to **Direct Connect** and select **I confirm that the network between the local PC and the floating IP address is connected and that MRS Manager is accessible using the Direct Connect connection**.

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

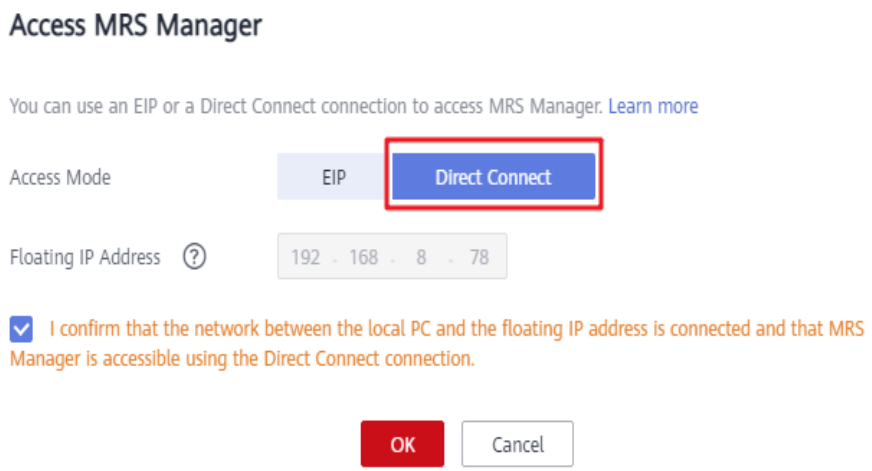

Step 5 Click **OK**. The MRS Manager login page is displayed. Enter the username **admin** and the password set during cluster creation.

----End

Switching the MRS Manager Access Mode

To facilitate user operations, the browser cache records the selected Manager access mode. To change the access mode, perform the following steps:

Step 1 Log in to the MRS console.

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

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

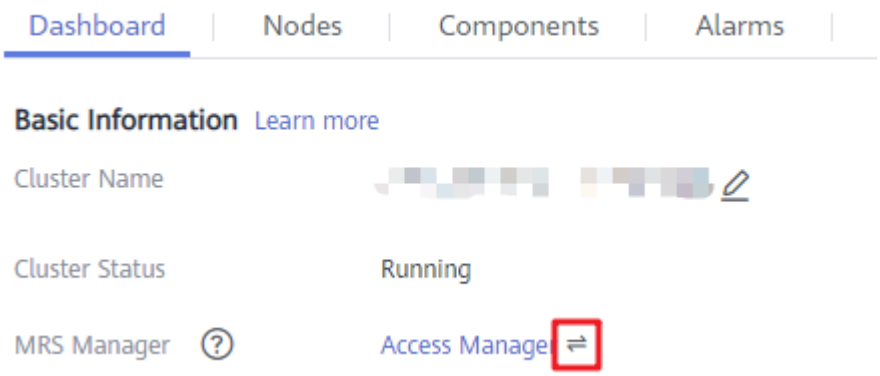

Step 4 On the displayed page, set **Access Mode**.

- To change **EIP** to **Direct Connect**, ensure that the network for direct connections is available, set **Access Mode** to **Direct Connect**, and select **I confirm that the network between the local PC and the floating IP address is connected and that MRS Manager is accessible using the Direct Connect connection**. Click **OK**.

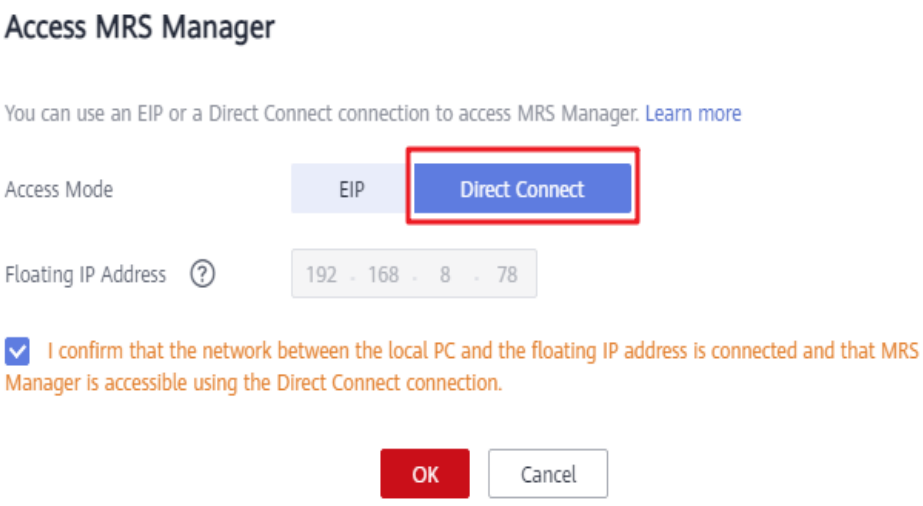

- To change **Direct Connect** to **EIP**, set **Access Mode** to **EIP** and configure the EIP by referring to **Accessing Manager Using an EIP**. If a public IP address

has been configured for the cluster, click **OK** to access MRS Manager using an EIP.

----End

6.4 EIP-based Access

You can bind an EIP to a cluster to access the web UIs of the open-source components managed in the MRS cluster. This method is simple and easy to use and is recommended for accessing the web UIs of the open-source components.

Binding an EIP to a Cluster and Adding a Security Group Rule

1. On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users. After the IAM users are synchronized, the **Components** tab is available.

Figure 6-1 Synchronizing users to access the **Component** tab page

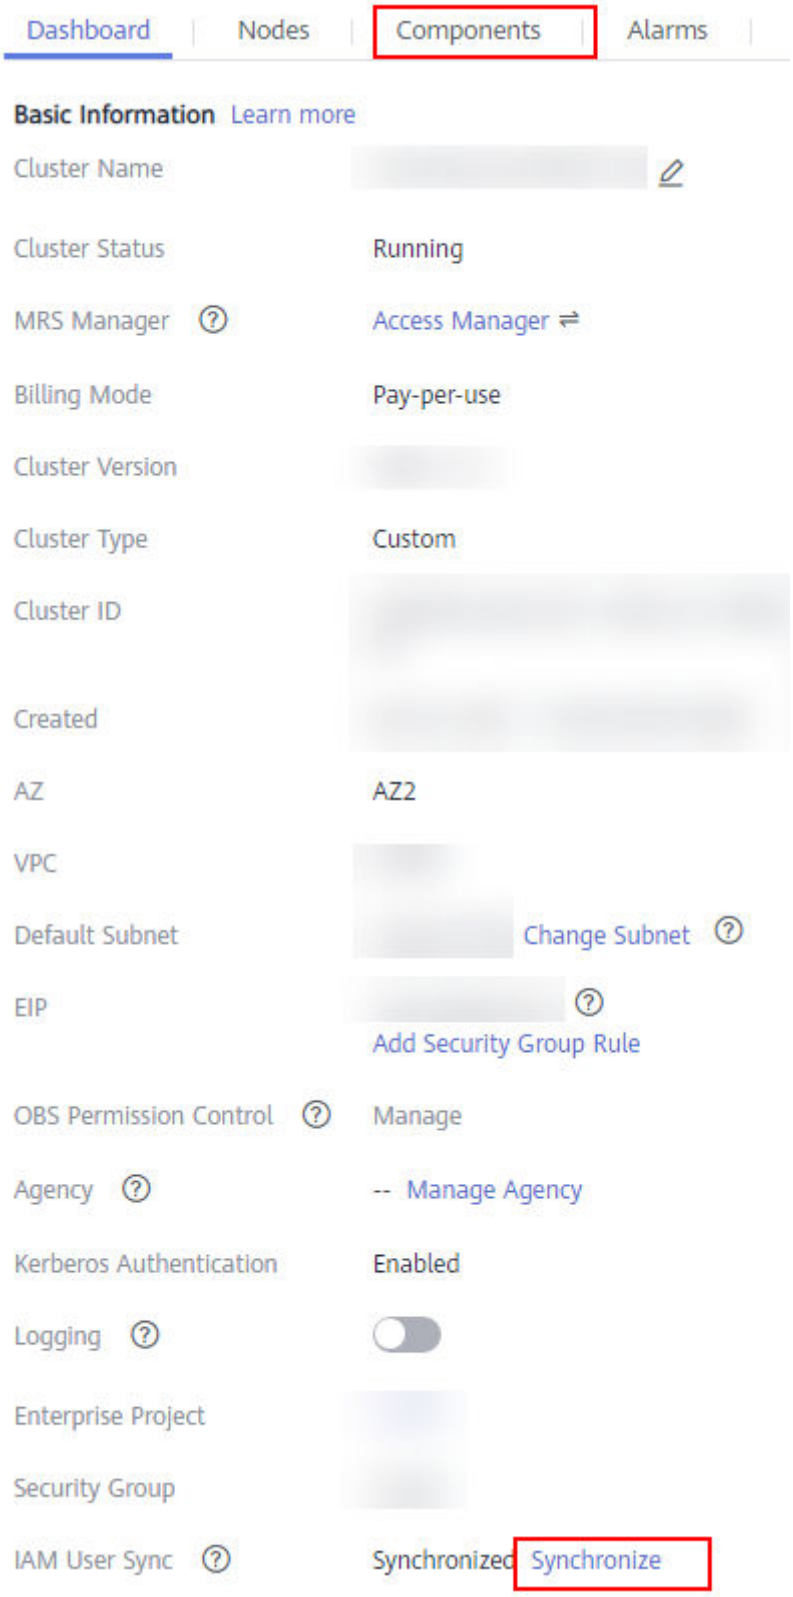

2. Click **Access Manager** on the right of **MRS Manager**.

3. The page for accessing MRS Manager is displayed. Bind an EIP and add a security group rule. Perform the following operations only when you access the web UIs of the open-source components of the cluster for the first time.

   a. Select an available EIP from the EIP drop-down list to bind it. If there is no available EIP, click **Manage EIP** to buy an EIP. If an EIP has been bound during cluster creation, skip this step.

   b. Select the security group to which the security group rule to be added belongs. The security group is configured when the group is created.

   c. Add a security group rule. By default, your public IP address used for accessing port 9022 is filled in the rule. If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

      NOTE

      - It is normal that the automatically generated public IP address is different from the local IP address and no action is required.
      - If port 9022 is a Knox port, you need to enable the permission of port 9022 to access Knox for accessing MRS components.

   d. Select the checkbox stating that **I confirm that xx.xx.xx.xx is a trusted public IP address and MRS Manager can be accessed using this IP address**.

      Figure 6-2 Accessing MRS Manager

      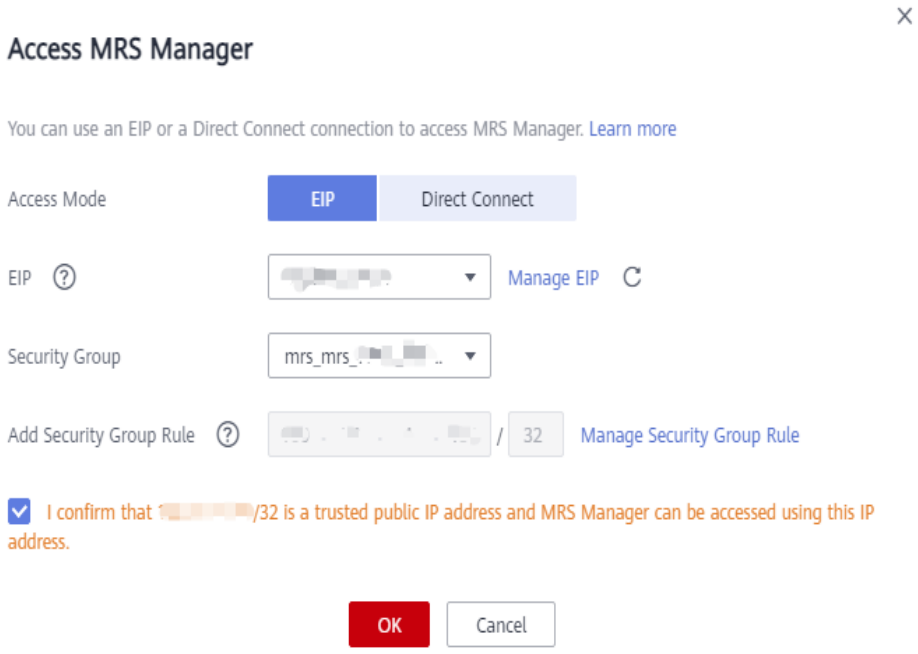

   e. Click **OK**. The login page is displayed. Enter the username **admin** and the password set during cluster creation.

4. Log in to FusionInsight Manager and choose **Cluster** > **Services** > **HDFS**. On the displayed page, click **NameNode(***Host name***, active)** to access the HDFS web UI. The HDFS NameNode is used as an example. For details about the web UIs of other components, see **Web UIs of Open Source Components**.

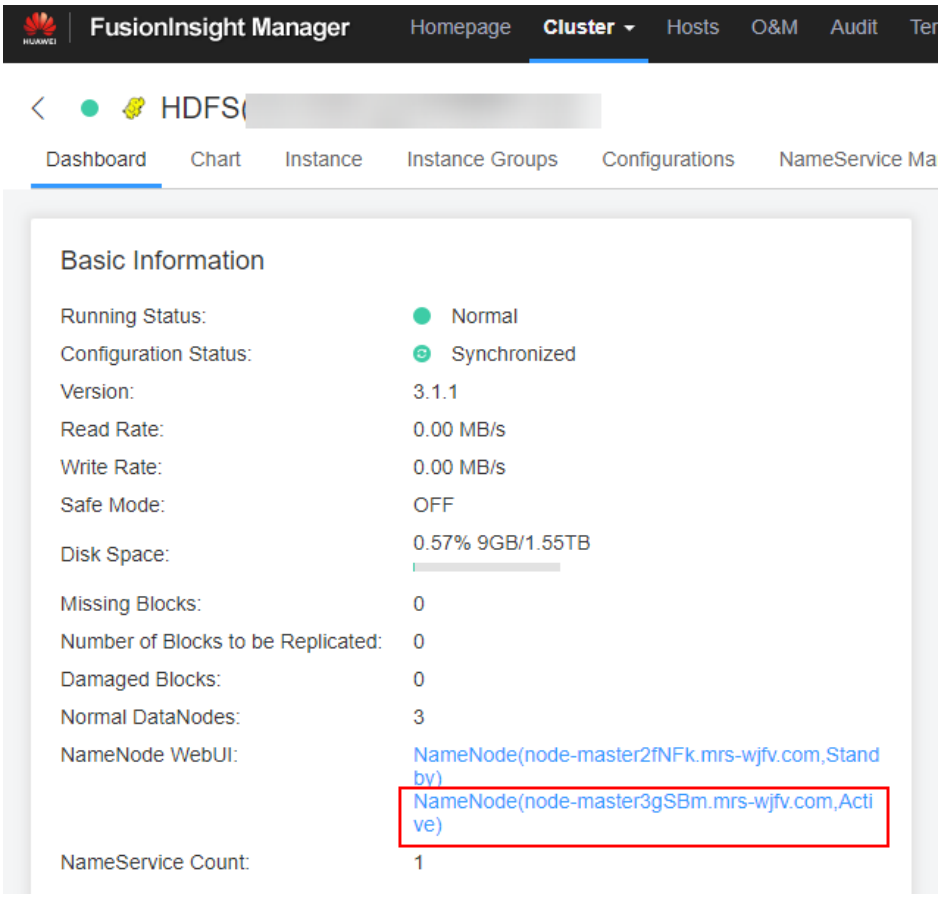

6.5 Access Using a Windows ECS

MRS allows you to access the web UIs of open-source components through a Windows ECS. This method is complex and is recommended for MRS clusters that do not support the EIP function.

Step 1 On the MRS management console, click **Clusters**.

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

On the cluster details page, record the **AZ**, **VPC**, **Floating IP Address of OMS**, and **Security Group** of the cluster.

NOTE

To obtain the floating IP address of OMS, log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of OMS. Record the value of **inet**. If the floating IP address of OMS cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the cluster manager IP address of the Master node.

Step 3 On the ECS management console, create an ECS.

- The **AZ**, **VPC**, and **Security Group** of the ECS must be the same as those of the cluster to be accessed.
- Select a Windows public image. For example, select the standard image **Windows Server 2012 R2 Standard 64bit(40GB)**.

- For details about other configuration parameters, see **Purchasing an ECS**.

NOTE

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

- Change the security group of the ECS to the security group of the MRS cluster. For details, see .
- Add two security group rules to the security groups of the Master and Core nodes to enable the ECS to access the cluster. Set **Protocol** to **TCP** and **ports** of the two security group rules to **28443** and **20009**, respectively. For details, see **Creating a Security Group**.

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

For details, see **Assigning an EIP and Binding It to an ECS**.

Step 5 Log in to the ECS.

The Windows system account, password, EIP, and security group rules are required for logging in to the ECS. For details, see .

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

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

NOTE

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

Step 7 Visit the web UIs of the open-source components by referring to the addresses listed in **Web UIs of Open Source Components**.

----End

Related Tasks

Configuring the Mapping Between Cluster Node Names and IP Addresses

Step 1 Log in to MRS Manager, and choose **Host Management**.

Record the host names and management IP addresses of all nodes in the cluster.

Step 2 In the work environment, use Notepad to open the **hosts** file and add the mapping between node names and IP addresses to the file.

Fill in one row for each mapping relationship, as shown in the following figure.


```
192.168.4.127 node-core-Jh3ER
192.168.4.225 node-master2-PaWVE
192.168.4.19 node-core-mtZ81
192.168.4.33 node-master1-zbYN8
192.168.4.233 node-core-7KoGY
```

Save the modifications.

----End

6.6 Creating an SSH Channel for Connecting to an MRS Cluster and Configuring the Browser

Scenario

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

The MAC system does not support this function. For details about how to access MRS, see **EIP-based Access**.

Prerequisites

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

Procedure

Step 1 Log in to the MRS management console and choose **Clusters** > **Active Clusters**.

Step 2 Click the specified MRS cluster name.

Record the security group of the cluster.

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

Step 4 Query the primary management node of the cluster. For details, see **Determining Active and Standby Management Nodes**.

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

For details, see .

Step 6 Start Git Bash locally and run the following command to log in to the active management node of the cluster: **ssh root@*Elastic IP address*** or **ssh -i *Path of the key file* root@*Elastic IP address***.

Step 7 Run the following command to view data forwarding configurations:

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

- If **net.ipv4.ip_forward=1** is displayed, the forwarding function has been configured. Go to **Step 9**.
- If **net.ipv4.ip_forward=0** is displayed, the forwarding function has not been configured. Go to **Step 8**.
- If **net.ipv4.ip_forward** fails to be queried, this parameter has not been configured. Run the following command and then go to **Step 9**:

  echo "net.ipv4.ip_forward = 1" >> /etc/sysctl.conf

Step 8 Modify forwarding configurations on the node.

1. Run the following command to switch to user **root**:

   sudo su - root

2. Run the following commands to modify forwarding configurations:

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

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

   sysctl -w net.ipv4.ip_forward=1

3. Run the following command to modify the **sshd** configuration file:

   vi /etc/ssh/sshd_config

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

   ```
   AllowTcpForwarding yes
   GatewayPorts yes
   ```

4. Run the following command to restart the sshd service:

   service sshd restart

Step 9 Run the following command to view the floating IP address:

ifconfig

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

Run the **exit** command to exit.

Step 10 Run the following command on the local PC to create an SSH channel supporting dynamic port forwarding:

ssh -i *Path of the key file* -v -ND *Local port* root@*Elastic IP address* or ***ssh -v -ND Local port root@Elastic IP address***. After running the command, enter the password you set when you create the cluster.

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

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

Step 11 Run the following command to configure the browser proxy.

1. Go to the Google Chrome client installation directory on the local PC.

2. Press **Shift** and right-click the blank area, choose **Open Command Window Here** and enter the following command:

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

   NOTE

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

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

Address format: **https://***Floating IP address of FusionInsight Manager***:28443/web**

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

When accessing Manager for the first time, you must add the address to the trusted site list.

Step 13 Prepare the website access address.

1. Obtain the website address format and the role instance according to **Web UIs**.
2. Click **Services**.
3. Click the specified service name, for example, HDFS.
4. Click **Instance** and view **Service IP Address** of **NameNode(Active)**.

Step 14 In the address bar of the browser, enter the website address to access it.

Step 15 When logging out of the website, terminate and close the SSH tunnel.

----End

7 Accessing Manager

7.1 Accessing FusionInsight Manager (MRS 3.*x* or Later)

Scenario

In MRS 3.*x* or later, FusionInsight Manager is used to monitor, configure, and manage clusters. After a cluster is installed, you can use an account to log in to FusionInsight Manager.

Currently, you can access FusionInsight Manager using the following methods:

- **Accessing FusionInsight Manager Using EIP**
- **Accessing FusionInsight Manager Using Direct Connect**
- **Accessing FusionInsight Manager from an ECS**

You can switch the access methods between **EIP** and **Direct Connect** on the MRS console by performing the following steps:

Log in to the MRS management console and click the desired cluster. On the displayed page, click ⇌ next to **MRS Manager** on the **Dashboard** tab, and switch the access method.

NOTE

If you cannot log in to the WebUI of the component, access FusionInsight Manager by referring to **Accessing FusionInsight Manager from an ECS**.

FusionInsight Manager cannot be accessed when the cluster is in any of the following states:

Starting, Stopping, Stopped, Deleting, Deleted, and Frozen.

Accessing FusionInsight Manager Using EIP

If the EIP address function is enabled for the cluster, perform the following steps:

Step 1 Log in to the MRS management console.

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

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

1. If no EIP is bound during MRS cluster creation, select an available EIP from the drop-down list on the right of **IEP**. If you have bound an EIP when you create a cluster, go to **Step 3.2**.

   NOTE

   - If no EIPs are available, click **Manage EIP** to buy one. Then, select the EIP from the drop-down list.
   - To unbind or release an EIP after using it, log in to the **EIPs** page, locate the row containing the target EIP, and click **Unbind** or choose **More** > **Release** in the **Operation** column.
   - If an EIP has been created but cannot be found during binding, the EIP may have been bound to another cluster. In this case, unbind the EIP on the **EIPs** page and then bind it to the current cluster.

2. In **Security Group**, select the security group to which the current cluster belongs. The security group is configured during cluster creation or is automatically created by the cluster.

   NOTE

   - When creating a custom cluster, you can configure a security group created in advance or retain the default value **Auto create**. When you quickly create a cluster, the security group is automatically created by the cluster.
   - You can view the security group name in **Security Group** on the **Dashboard** tab page of the cluster.

3. Add a security group rule. By default, the filled-in rule is used to access the EIP. To enable multiple IP address segments to access Manager, see steps **Step 6** to **Step 9**. If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

   NOTE

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

   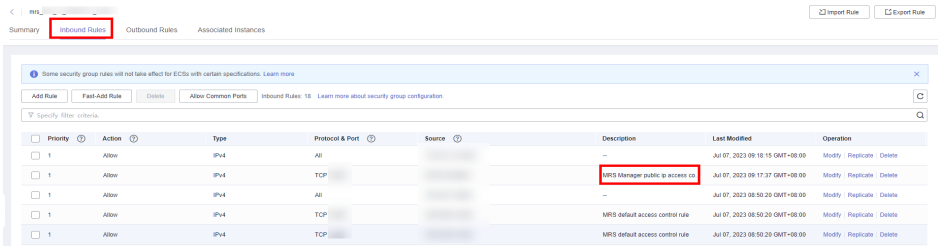

4. Select the information to be confirmed and click **OK**.

   NOTE

   Click ⇌ on the right of **Access Manager** to change the FusionInsight Manager access mode. For details about how to access FusionInsight Manager by using **Direct Connect**, see **Accessing FusionInsight Manager by Using Direct Connect**.

Step 4 Click **OK**. The Manager login page is displayed.

Step 5 Enter the default username **admin** and the password set during cluster creation, and click **Log In**. The Manager page is displayed.

Step 6 On the MRS management console, choose **Clusters** > **Active Clusters**. Click the target cluster name to access the cluster details page.

NOTE

To grant other users the permission to access Manager, perform **Step 6** to **Step 9** to add the users' public IP addresses to the trusted IP address range.

Step 7 Click **Add Security Group Rule** next to **EIP**.

Figure 7-1 Cluster details

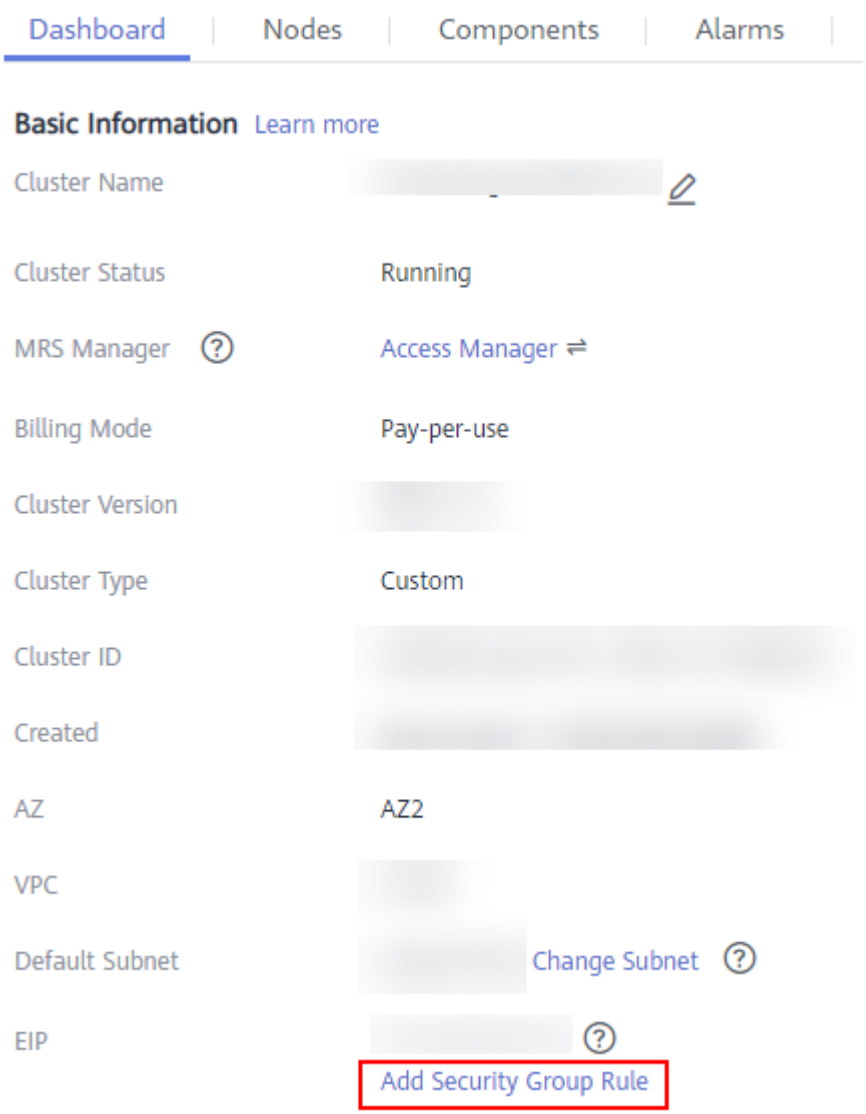

Step 8 On the **Add Security Group Rule** page, add the IP address segment for users to access the public network and select **I confirm that *public network IP*|*port* is a trusted public IP address. I understand that using 0.0.0.0/0. poses security risks**. See **Figure 7-2**.

Figure 7-2 Adding a security group rule

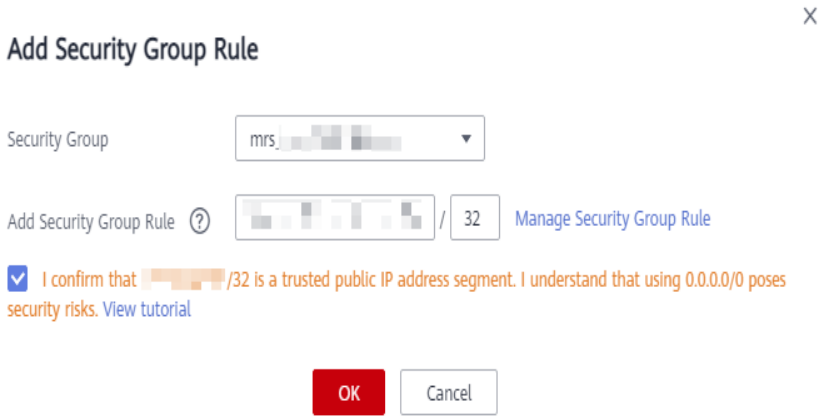

By default, the IP address used for accessing the public network is filled. You can change the IP address segment as required. To enable multiple IP address segments, repeat steps **Step 6** to **Step 9**. If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

Step 9 Click **OK**.

----End

Accessing FusionInsight Manager by Using Direct Connect

You need to ensure that Direct Connect is available, and the connection between the local data center and the online VPC has been established.

Step 1 Log in to the MRS console.

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

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

Step 4 Set **Access Mode** to **Direct Connect** and select **I confirm that the network between the local PC and the floating IP address is connected and that MRS Manager is accessible using the Direct Connect connection**.

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

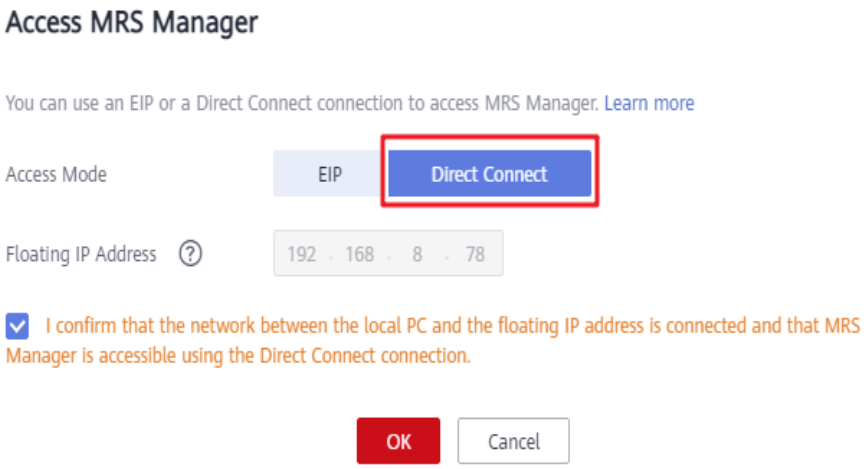

Step 5 Click **OK**. The MRS Manager login page is displayed. Enter the username **admin** and the password set during cluster creation.

----End

Accessing FusionInsight Manager from an ECS

Step 1 On the MRS management console, click **Clusters**.

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

Record the **AZ**, **VPC**, **MRS ManagerSecurity Group** of the cluster.

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

- The **AZ**, **VPC**, and **Security Group** of the ECS must be the same as those of the cluster to be accessed.
- Select a Windows public image. For example, a standard image **Windows Server 2012 R2 Standard 64bit(40GB)**.
- For details about other configuration parameters, see **Purchasing an ECS with Customized Configurations**.

NOTE

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

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

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

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

For details, see **Assigning an EIP and Binding It to an ECS**.

Step 5 Log in to the ECS.

The Windows system account, password, EIP, and security group rules are required for logging in to the ECS.

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

The address for accessing Manager is the address of the **MRS Manager** page. Enter the name and password of the cluster user, for example, user **admin**.

NOTE

- If you access Manager with other cluster usernames, change the password upon your first access. The new password must meet the requirements of the current password complexity policies. For details, contact the administrator.
- By default, a user is locked after inputting an incorrect password five consecutive times. The user is automatically unlocked after 5 minutes.

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

----End

7.2 Accessing MRS Manager (MRS 2.*x* or Earlier)

Scenario

MRS uses Manager to monitor, configure, and manage clusters. You can go to the Manager management page on the MRS console and use the admin account and password configured during cluster creation to log in to Manager.

NOTE

The Manager cannot be accessed when the cluster is in any of the following states: Starting, Stopping, Stopped, Deleting, Deleted, and Frozen.

Accessing Manager Using an EIP

Step 1 Log in to the MRS management console.

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

Step 3 Click **Access Manager** next to **MRS Manager**. In the displayed dialog box, set **Access Mode** to **EIP**. For details about **Direct Connect**, see **Access Through Direct Connect**.

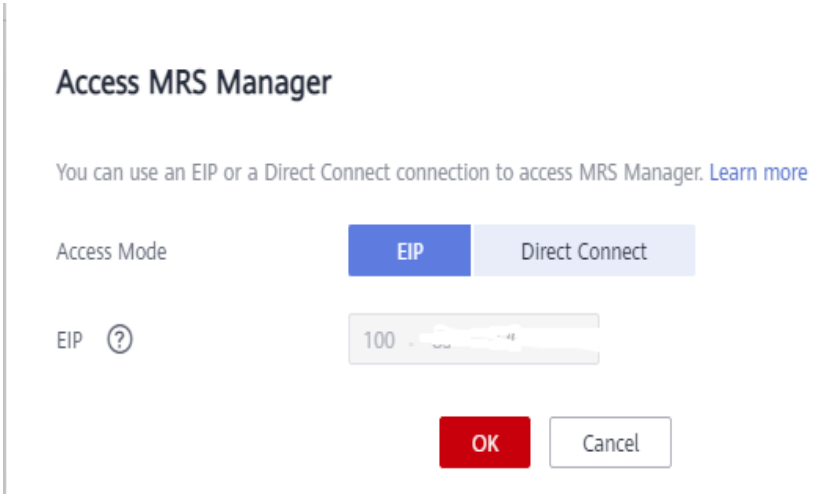

1. If no EIP is bound during MRS cluster creation, select an available EIP from the drop-down list on the right of **IEP**. If you have bound an EIP when creating a cluster, go to **Step 3.2**.

   NOTE

   - If no EIPs are available, click **Manage EIP** to buy one. Then, select the EIP from the drop-down list.
   - To unbind or release an EIP after using it, log in to the **EIPs** page, locate the row containing the target EIP, and click **Unbind** or choose **More** > **Release** in the **Operation** column.
   - If an EIP has been created but cannot be found during binding, the EIP may have been bound to another cluster. In this case, unbind the EIP on the **EIPs** page and then bind it to the current cluster.

2. In **Security Group**, select the security group to which the current cluster belongs. The security group is configured during cluster creation or is automatically created by the cluster.

   NOTE

   - When creating a custom cluster, you can configure a security group created in advance or retain the default value **Auto create**. When you quickly create a cluster, the security group is automatically created by the cluster.
   - You can view the security group name in **Security Group** on the **Dashboard** tab page of the cluster.

3. Add a security group rule. By default, your public IP address used for accessing port 9022 is filled in the rule. To enable multiple IP address segments to access MRS Manager, see **Step 6** to **Step 9**. If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

   NOTE

   - It is normal that the automatically generated public IP address is different from the local IP address and no action is required.
   - If port 9022 is a Knox port, you need to enable the permission of port 9022 to access Knox for accessing MRS Manager.

4. Select the checkbox stating that **I confirm that xx.xx.xx.xx is a trusted public IP address and MRS Manager can be accessed using this IP address**.

Figure 7-3 Binding an EIP

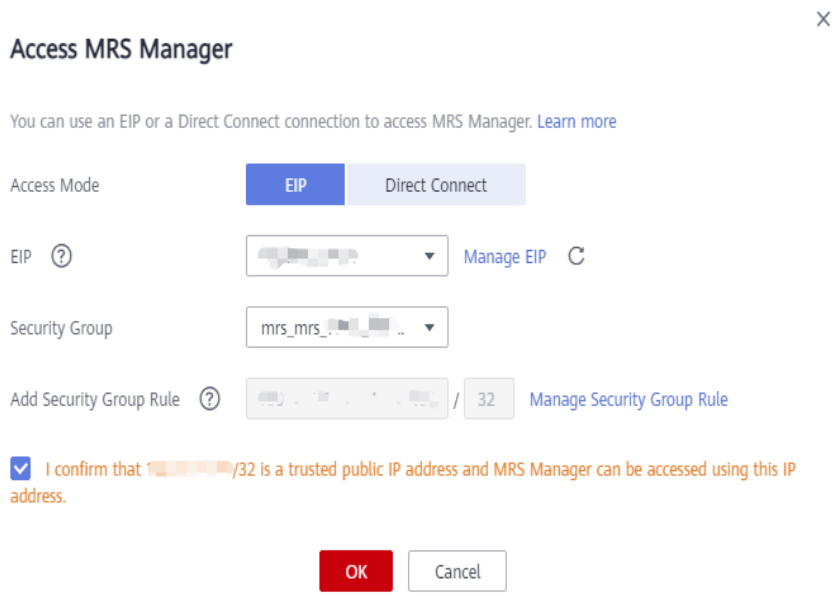

Step 4 Click **OK**. The MRS Manager login page is displayed.

Step 5 Enter the default username **admin** and the password set during cluster creation, and click **Log In**. The MRS Manager page is displayed.

Step 6 On the MRS management console, choose **Clusters** > **Active Clusters**, and click the target cluster name to access the cluster details page.

NOTE

To assign MRS Manager access permissions to other users, follow instructions from **Step 6** to **Step 9** to add the users' public IP addresses to the trusted range.

Step 7 Click **Add Security Group Rule** next to **EIP**.

Figure 7-4 Cluster details

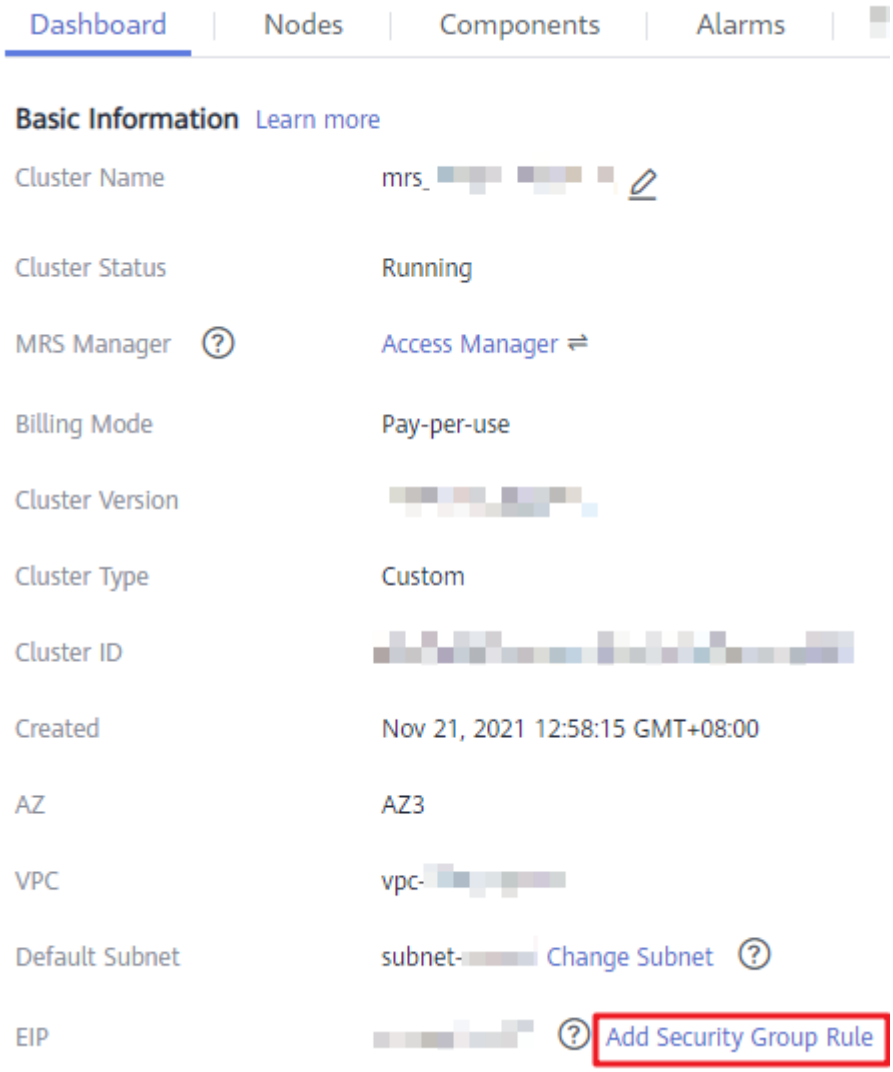

Step 8 On the **Add Security Group Rule** page, add the IP address segment for users to access the public network and select **I confirm that the authorized object is a trusted public IP address range. Do not use 0.0.0.0/0. Otherwise, security risks may arise**. See **Figure 7-5**.

Figure 7-5 Adding a security group rule

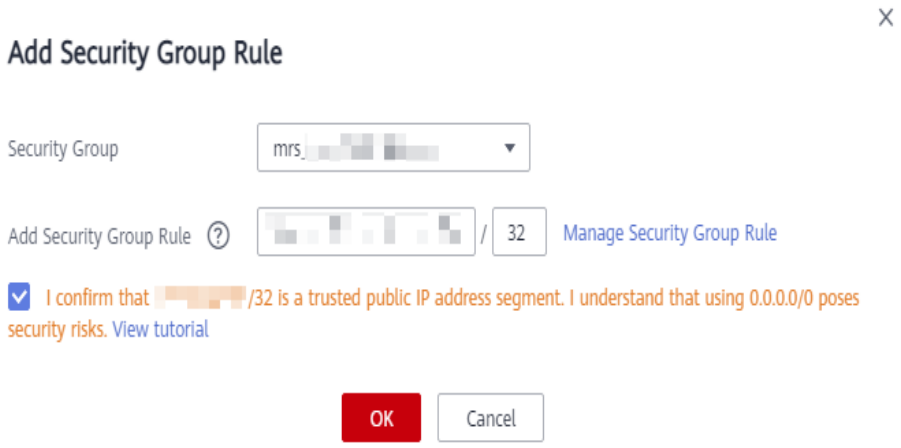

By default, the IP address used for accessing the public network is filled. You can change the IP address segment as required. To enable multiple IP address segments, repeat steps **Step 6** to **Step 9**. If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

Step 9 Click **OK**.

----End

Accessing MRS Manager Using an ECS

Step 1 On the MRS management console, click **Clusters**.

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

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

Step 3 On the ECS management console, create an ECS.

- The **AZ**, **VPC**, and **Security Group** of the ECS must be the same as those of the cluster to be accessed.
- Select a Windows public image. For example, select the standard image **Windows Server 2012 R2 Standard 64bit(40GB)**.
- For details about other configuration parameters, see **Creating an ECS**.

NOTE

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

- Change the default security group of the ECS to the security group of the MRS cluster. For details, see **Changing a Security Group**.
- Add two security group rules to the security groups of the Master and Core nodes to enable the ECS to access the cluster. Set **Protocol** to **TCP** and **ports** of the two security group rules to **28443** and **20009**, respectively. For details, see .

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

For details, see **Assigning an EIP and Binding It to an ECS**.

Step 5 Log in to the ECS.

The Windows system account, password, EIP, and security group rules are required for logging in to the ECS.

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

For example, you can use Internet Explorer 11 in the Windows 2012 OS.

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

NOTE

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

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

----End

Changing an EIP for a Cluster

Step 1 On the MRS management console, choose **Clusters** > **Active Clusters**, and click the target cluster name to access the cluster details page.

Step 2 View EIPs

Step 3 Log in to the VPC management console.

Step 4 Choose **Elastic IP and Bandwidth** > **EIPs**.

Step 5 Search for the EIP bound to the MRS cluster and click **Unbind** in the **Operation** column to unbind the EIP from the MRS cluster.

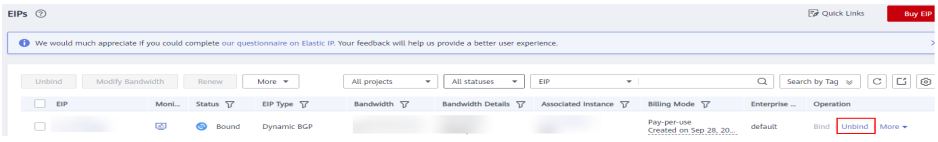

Step 6 Log in to the MRS management console, choose **Clusters** > **Active Clusters**, and click the target cluster name to access the cluster details page.

EIP on the cluster details page is displayed as **Unbound**.

Step 7 Click **Access Manager** next to **MRS Manager**. In the displayed dialog box, set **Access Mode** to **EIP**.

Step 8 Select a new EIP from the EIP drop-down list and configure other parameters. For details, see **Accessing Manager Using an EIP**.

----End

Granting the Permission to Access MRS Manager to Other Users

Step 1 On the MRS management console, choose **Clusters** > **Active Clusters**, and click the target cluster name to access the cluster details page.

Step 2 Click **Add Security Group Rule** on the right of **EIP**, as shown in **Figure 7-6**.

Figure 7-6 Cluster details

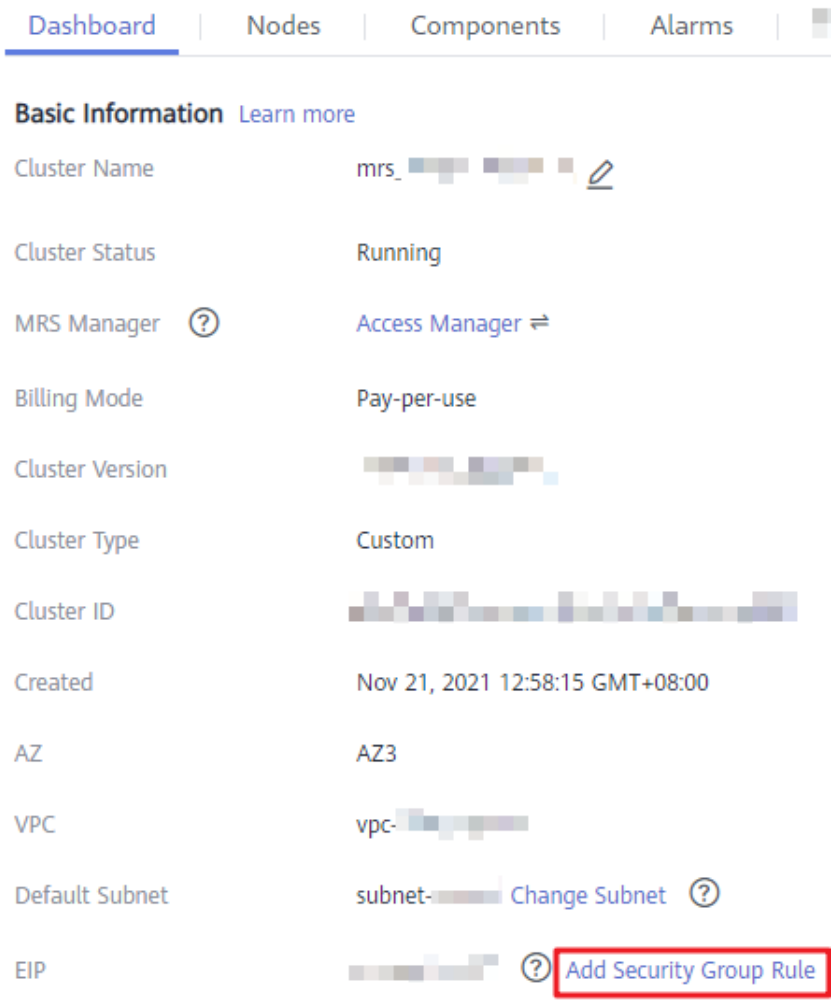

Step 3 On the **Add Security Group Rule** page, add the IP address segment for users to access the public network and select **I confirm that the authorized object is a trusted public IP address range. Do not use 0.0.0.0/0. Otherwise, security risks may arise**. See **Figure 7-7**.

Figure 7-7 Adding a security group rule

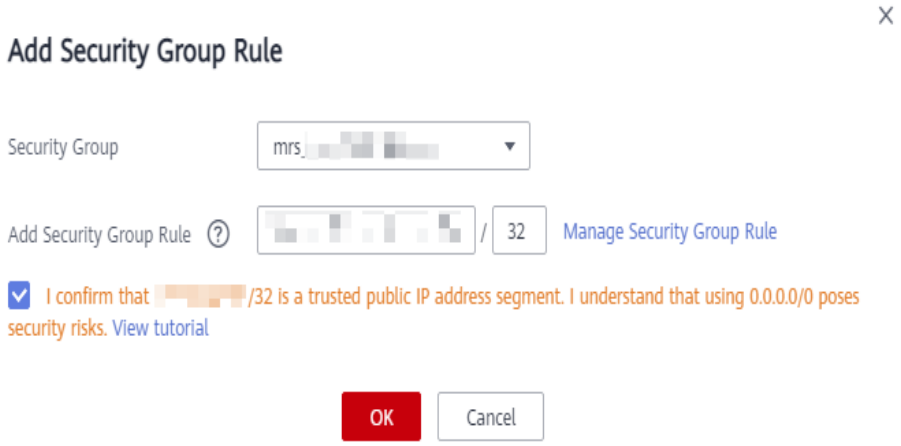

By default, the IP address used for accessing the public network is filled. You can change the IP address segment as required. To enable multiple IP address segments, repeat steps **Step 1** to **Step 4**. If you want to view, modify, or delete a security group rule, click **Manage Security Group Rule**.

Step 4 Click **OK**.

----End

8 FusionInsight Manager Operation Guide (Applicable to 3.x)

8.1 Homepage

8.1.1 Overview

After you log in to FusionInsight Manager, the home page is displayed by default. You can preview the status of each cluster, view the monitoring information, and view alarm statistics and analysis in the cluster.

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

  Figure 8-1 Cluster status information

  

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

    NOTE

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

  - Click About under  to view the Manager version.
- The taskbar at the bottom of the home page displays the language options of FusionInsight Manager and the current cluster time and time zone information. You can switch the system language as needed.

  Figure 8-2 Taskbar at the bottom of the home page

  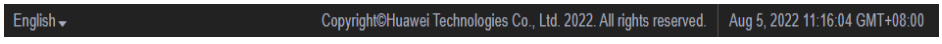

Service Status Preview Area

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

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

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

The icon displayed on the right of the service name indicates that the service configuration has expired.

Monitoring Status Report Area

The chart area is on the right of the homepage, which displays key monitoring metric reports, such as the status of all hosts in the cluster, host CPU usage, and host memory usage. You can customize monitoring reports to display in this area. For details about how to manage monitoring metrics, see **Managing Monitoring Metric Reports**.

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

Figure 8-3 Monitoring status report

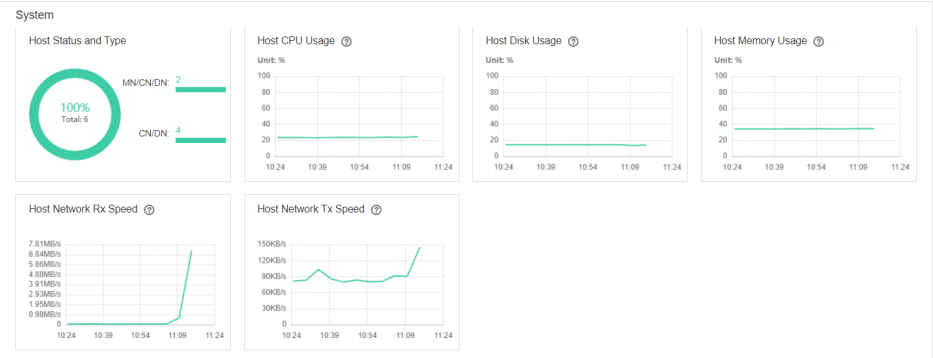

Alarm Analysis

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

Figure 8-4 Alarm Analysis tab page

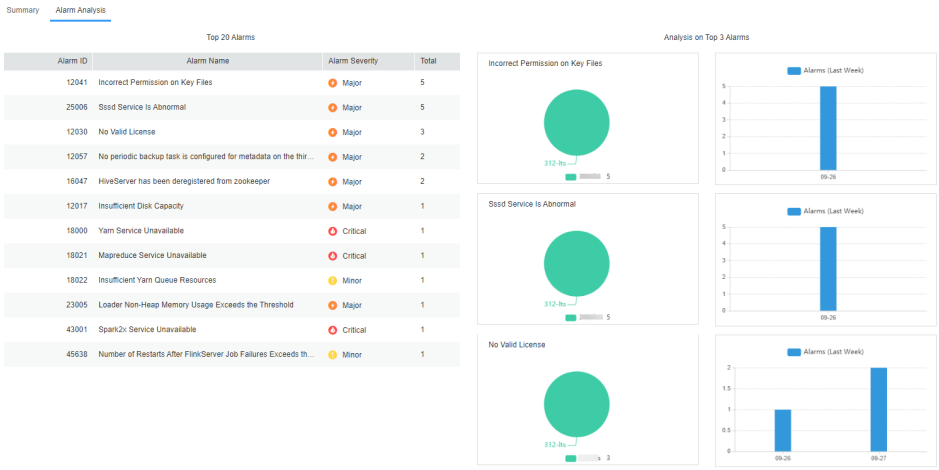

NOTE

For MRS 3.3.0 and later versions, the alarm information page is different from that of other versions.

8.1.2 Managing Monitoring Metric Reports

Scenario

On FusionInsight Manager, you can customize monitoring items to display on the homepage and export monitoring data.

NOTE

The interval on the horizontal axis of the chart varies depending on the time period you specify. Data monitoring rules are as follows:

- If the disk usage of the partition where GaussDB is deployed exceeds 80%, real-time monitoring data and monitoring data whose interval is 5 minutes will be deleted.
- **Storage resources (HDFS) in Tenant Resources (0 to 300 hours)**: The interval is 1 hour. The cluster must have been installed for at least 1 hour, and monitoring data of a maximum of 3 months is saved.

For clusters of versions earlier than MRS 3.3.0:

- **0 to 25 hours**: The interval is 5 minutes. The cluster must have been installed for at least 10 minutes, and monitoring data of a maximum of 15 days is saved.
- **25 to 150 hours**: The interval is 30 minutes. The cluster must have been installed for at least 30 minutes, and monitoring data of a maximum of 3 months is saved.
- **150 to 300 hours**: The interval is 1 hour. The cluster must have been installed for at least 1 hour, and monitoring data of a maximum of 3 months is saved.
- **300 hours to 300 days**: The interval is 1 day. The cluster must have been installed for at least 1 day, and monitoring data of a maximum of 6 months is saved.
- **Over 300 days**: The interval is 7 days. The cluster must have been installed for more than 7 days, and monitoring data of a maximum of 1 year is saved.

For clusters of MRS 3.3.0 or later versions:

- **0 to 21 hours and 20 minutes**: The interval is 5 minutes. The cluster must have been installed for at least 10 minutes, and monitoring data of a maximum of 90 days is saved.
- **21 hours and 20 minutes to 128 hours**: The interval is 30 minutes. The cluster must have been installed for at least 30 minutes, and monitoring data of a maximum of 90 days is saved.
- **128 to 256 hours**: The interval is 1 hour. The cluster must have been installed for at least 1 hour, and monitoring data of a maximum of 90 days is saved.
- **256 hours to 256 days**: The interval is 1 day. The cluster must have been installed for at least 1 day, and monitoring data of a maximum of 90 days is saved.

Customizing a Monitoring Metric Report

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Homepage**.

Step 3 In the upper right corner of the chart area, click ▼ and choose **Customize** from the displayed menu.

NOTE

Monitoring data of the past 1 hour is displayed at an interval of 5 minutes. After you enter the **Real-time Monitoring** page, you can view that real-time monitoring data is displayed on the right of the monitoring chart at an interval of 5 minutes.

Step 4 In the left pane of the **Customize Statistics** dialog box, select a resource to monitor.

Step 5 Select one or multiple monitoring metrics in the right pane.

Figure 8-5 Customizing a monitoring metric report

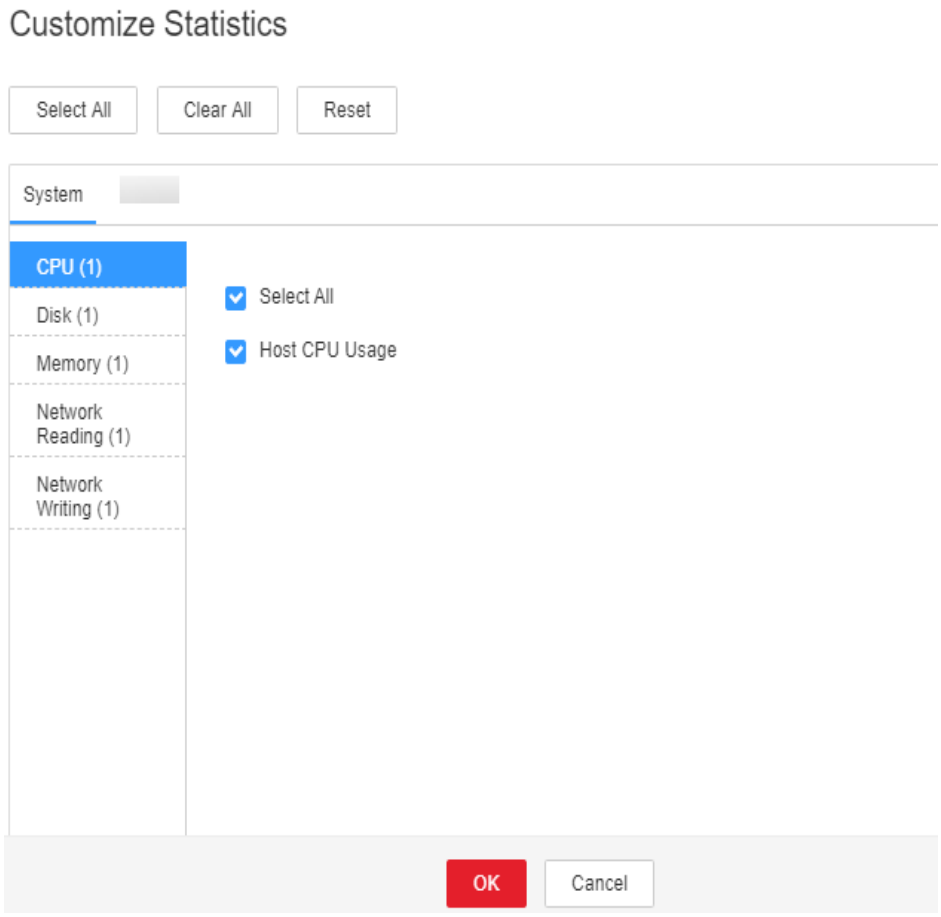

Step 6 Click **OK**.

----End

Exporting All Monitoring Data

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Homepage**.

Step 3 In the upper right corner of the chart area, select a time range to obtain monitoring data, for example, **1w**.

Real-time data is displayed by default, which cannot be exported. You can click to customize a time range.

Figure 8-6 Customizing a time range

Step 4 In the upper right corner of the chart area, click and choose **Export** from the displayed menu.

----End

Exporting Monitoring Data of a Specified Monitoring Item

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Homepage**.

Step 3 Click in the upper right corner of any monitoring report pane in the chart area of the target cluster.

Step 4 Select a time range to obtain monitoring data, for example, **1w**.

Real-time data is displayed by default, which cannot be exported. You can click to customize a time range.

Figure 8-7 Customizing a time range for a specified monitoring item

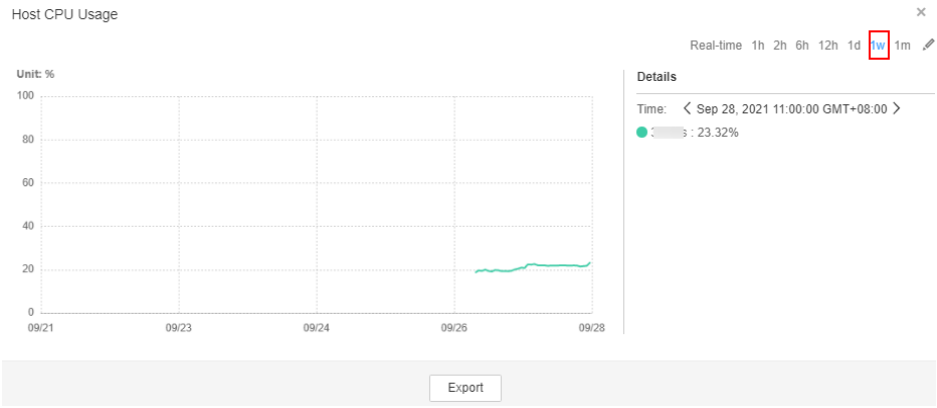

Step 5 Click **Export**.

----End

8.1.3 Querying the FusionInsight Manager Version

By viewing the FusionInsight Manager version, you can prepare for system upgrade and routine maintenance.

- Using the GUI:

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

  Figure 8-8 Viewing the version

- Using the CLI
  a. Log in to the FusionInsight Manager active OMS node as user **root**. To obtain the IP address of the active OMS node, view the host information marked with ★ in the **Hosts** tab of FusionInsight Manager.
  b. Run the following commands to check the version and platform information of FusionInsight Manager:

     su - omm

     cd ${BIGDATA_HOME}/om-server/om/sbin/pack

     ./queryManager.sh

     The following information is displayed:

Version	Package	Cputype
***	FusionInsight_Manager_***	x86_64

NOTE

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

8.2 Cluster

8.2.1 Cluster Management

8.2.1.1 Overview

Dashboard

Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Dashboard** to view the status of the cluster.

In the **Dashboard** tab, you can start, stop, perform a rolling restart of, synchronize configurations to the cluster, and perform other basic operations.

Figure 8-9 Dashboard

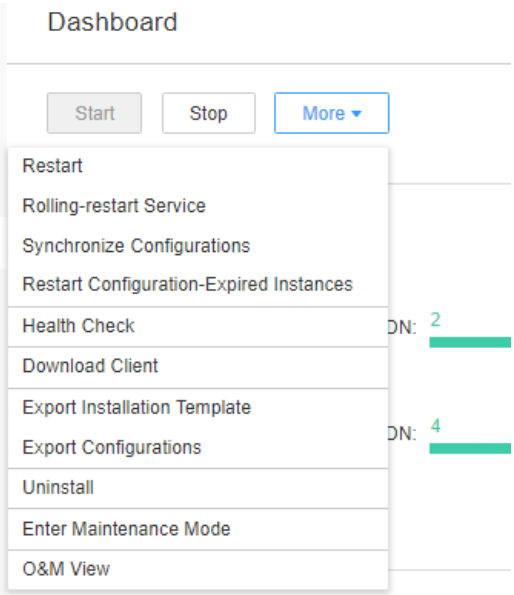

NOTE

For MRS 3.3.0 or later, the **Cluster** > **Dashboard** page has been removed from Manager. You can choose **More** in the upper right corner of the **Homepage** to access cluster maintenance and management functions.

Table 8-1 Maintenance and management operations

UI Portal	Description
Start	Starts all services in the cluster.
Stop	Stops all services in the cluster.
More > **Restart**	Restarts all services in the cluster.
More > **Rolling-restart Service**	Restarts all services in the cluster one at a time without interrupting workloads. For details about how to perform a rolling restart, see **Performing a Rolling Restart of a Cluster**.
More > **Synchronize Configurations**	Enables new configuration parameters for all services in the cluster.
More > **Restart Configuration-Expired Instances**	Restarts expired instances for all services in the cluster. For details, see **Managing Expired Configurations**.
More > **Health Check**	Performs a health check on the OMS nodes, all services, and the rest nodes in the cluster. There are three types of check items: running status, related alarms, and custom monitoring metrics. The health check results are not always the same as the values of **Running Status** displayed on the GUI. You can export check results by clicking **Export** in the upper left corner of the checklist. If any issues are detected, you can click **View Help** to find a troubleshooting method.
More > **Download Client**	Downloads the default client. For details, see **Downloading the Client**.
More > **Export Installation Template**	Batch exports all installation configurations of the cluster, such as the cluster authentication mode, node information, and service configuration. You can use this function when you need to reinstall the cluster in the same environment.
More > **Export Configurations**	Batch exports configurations of all services in the cluster.
More > **Enter Maintenance Mode** and **More** > **Exit Maintenance Mode**	Enters or exits the cluster maintenance mode.

UI Portal	Description
More > **O&M View**	Allows you to view services or hosts that are in the maintenance mode.

8.2.1.2 Performing a Rolling Restart of a Cluster

Scenario

A rolling restart is batch restarting all services in a cluster after they are modified or upgraded without interrupting workloads.

You can perform a rolling restart of a cluster as needed.

NOTE

- Certain services in a cluster do not support rolling restart. These services are restarted in normal mode during the rolling restart of the cluster. As a result, workloads may be interrupted. So, you need to determine whether to perform this operation as prompted.
- Configurations that must take effect immediately, for example, server port configurations, should be restarted in normal mode.

Impact on the System

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

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > **Dashboard**. In the upper right corner, click **More** > **Service Rolling Restart**.

NOTE

For MRS 3.3.0 or later, the **Cluster** > **Dashboard** page has been removed from Manager. You can choose **More** in the upper right corner of the **Homepage** to access cluster maintenance and management functions.

Step 3 In the dialog box that is displayed, enter the password of the current login user and click **OK**.

Step 4 Configure the parameters based on site requirements.

Figure 8-10 Rolling-restart Cluster

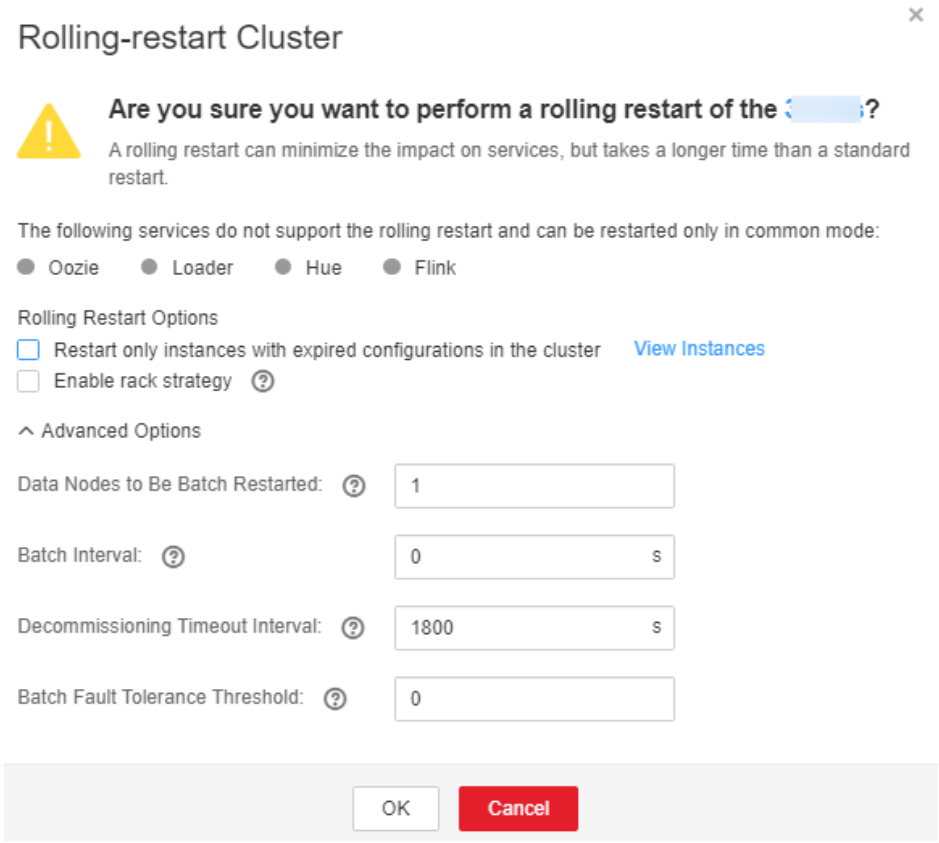

Table 8-2 Rolling restart parameters

Parameter	Description
Restart only instances with expired configurations in the cluster	Whether to restart only the modified instances in a cluster
Enable rack strategy	Whether to enable the concurrent rack rolling restart strategy. This parameter takes effect only for roles that meet the rack rolling restart strategy. (The roles support rack awareness, and instances of the roles belong to two or more racks.) **NOTE** This parameter is configurable only when a rolling restart is performed on HDFS or YARN.

Parameter	Description
Data Nodes to Be Batch Restarted	Number of instances that are restarted in each batch when the batch rolling restart strategy is used. The default value is **1**. **NOTE** • This parameter is valid only when the batch rolling restart strategy is used and the instance type is DataNode. • This parameter is invalid when the rack strategy is enabled. In this case, the cluster uses the maximum number of instances (20 by default) configured in the rack strategy as the maximum number of instances that are concurrently restarted in a rack. • This parameter is configurable only when a rolling restart is performed on HDFS, HBase, YARN, Kafka, Storm, or Flume. • This parameter for the RegionServer of HBase cannot be manually configured. Instead, it is automatically adjusted based on the number of RegionServer nodes. Specifically, if the number of RegionServer nodes is less than 30, the parameter value is **1**. If the number is greater than or equal to 30 and less than 300, the parameter value is **2**. If the number is greater than or equal to 300, the parameter value is 1% of the number (rounded-down).
Batch Interval	Interval between two batches of instances to be roll-restarted. The default value is **0**.
Decommissioning Timeout Interval	Decommissioning interval for role instances during a rolling restart. The default value is **1800s**. Some roles (such as HiveServer and JDBCServer) stop providing services before the rolling restart. Stopped instances cannot be connected to new clients. Existing connections will be completed after a period of time. An appropriate timeout interval can ensure service continuity. **NOTE** This parameter is configurable only when a rolling restart is performed on Hive or Spark2x.
Batch Fault Tolerance Threshold	Tolerance times when the rolling restart of instances fails to be batch executed. The default value is **0**, which indicates that the rolling restart task ends after any batch of instances fails to restart.

NOTE

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

Example:

- If **Data Nodes to Be Batch Restarted** is set to an unnecessarily large value, a large number of instances are restarted concurrently. As a result, services are interrupted or cluster performance is severely affected due to too few working instances.
- If **Batch Fault Tolerance Threshold** is too large, services will be interrupted because a next batch of instances will be restarted after a batch of instances fails to restart.

Step 5 Click **OK**.

----End

8.2.1.3 Managing Expired Configurations

Scenario

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

If the configuration of the services in the cluster has been synchronized but do not take effect, you need to restart the instances whose configuration has expired.

Impact on the System

- After synchronizing the cluster configuration, you need to restart the services whose configuration has expired. These services are unavailable during restart.
- The instances whose configuration has expired are unavailable during restart.

Procedure

Synchronize the configuration.

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Dashboard**.

NOTE

For MRS 3.3.0 or later, the **Cluster** > **Dashboard** page has been removed from Manager. You can choose **More** in the upper right corner of the **Homepage** to access cluster maintenance and management functions.

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

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

----End

Restart configuration-expired instances.

Step 1 Choose **More** > **Restart Configuration-Expired Instances**.

Step 2 In the dialog box that is displayed, enter the password of the current login user and click **OK**.

Step 3 In the displayed dialog box, click **OK**.

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

----End

8.2.1.4 Downloading the Client

Scenario

Use the default client provided by MRS clusters to manage the cluster, run services, and perform secondary development. Before you use this client, you need to download its software package.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > **Dashboard**. In the upper right corner, click **More** > **Download Client**.

NOTE

For MRS 3.3.0 or later, the **Cluster** > **Dashboard** page has been removed from FusionInsight Manager.

The **Download Cluster Client** dialog box is displayed.

Figure 8-11 Downloading the cluster client

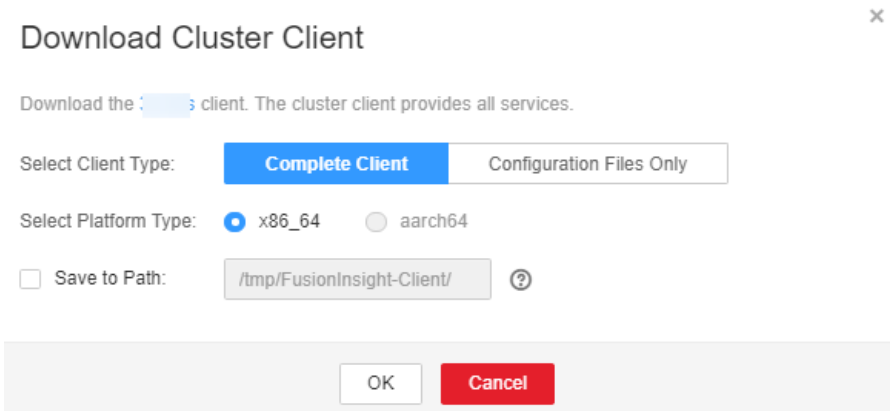

Step 3 Select a client type for **Select Client Type**.

- **Complete Client**: the package contains scripts, compilation files, and configuration files.
- **Configuration Files Only**: the package contains only the client configuration files.

This type is applicable to application development tasks. For example, after a complete client is downloaded and installed, the cluster administrator modifies the service configuration on FusionInsight Manager, and developers need to update the client configuration files.

NOTE

The two platform types (**x86_64** and **aarch64**) correspond to x86 and ARM nodes, respectively. Select the platform type based on the architecture types of the nodes where you will install the client.

Step 4 Determine whether to generate a client software package file on the cluster node.

- If yes, select **Save to Path** and click **OK** to generate the client file.

  The generated file is stored in the **/tmp/FusionInsight-Client** directory on the active management node by default. You can also store the client file in other directories, and user **omm** has the read, write, and execute permissions on the directory. If the client file already exists in the path, the existing client file will be replaced.

  After the file is generated, copy the obtained package to another directory, for example, **/opt/Bigdata/hadoopclient**, as user **omm** or client installation user.
- If no, click **OK** to download the client file to the local host.

  The system starts to download the client software package.

----End

8.2.1.5 Modifying Cluster Attributes

Scenario

View basic cluster attributes on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Cluster Properties**.

By default, you can view the cluster name, cluster description, product type, cluster ID, authentication mode, creation time, and installed components.

Figure 8-12 Cluster Properties page

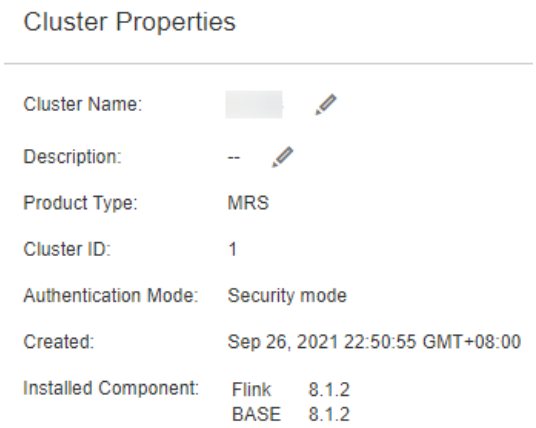

Step 3 Change the cluster name.

1. Click and enter a new name.

   Enter 2 to 199 characters. Only letters, digits, underscores (_), hyphens (-), and spaces are allowed, and the name cannot start with a space.

2. Click **OK** for the new cluster name to take effect.

Step 4 Modify the cluster description.

1. Click and enter a new description.

   Enter a maximum of 199 characters. Only letters, digits, commas (,), periods (.), underscores (_), spaces, and newline characters (\n) are allowed.

2. Click **OK** for the new description to take effect.

----End

8.2.1.6 Managing Cluster Configurations

Scenario

FusionInsight Manager allows you to view the changes of service configuration parameters in a cluster with one click, helping you quickly locate faults and improve configuration management efficiency.

You can quickly view all non-default values of each service in the cluster, non-uniform values between instances of the same role, historical records of cluster configuration modifications, and expired parameters in the cluster on the configuration page.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Configurations**.

Step 3 Select an operation page based on the scenario.

- To view all non-default values:
  a. Click **All Non-default Values**. The system displays the parameters whose values are different from the default values configured for each service, role, or instance in the current cluster.

     You can click next to a parameter value to quickly restore the value to the default one. You can click to view the historical modification records of the parameter.

     If there are a large number of parameters to configure, you can filter the parameters in the filter box in the upper right corner of the page or enter keywords in the search box.

  b. To change the values of the parameters, change the values according to the parameter description and click **Save**. In the dialog box that is displayed, click **OK**.
- To view all non-uniform values:

a. Click **All Non-uniform Values**. The system displays parameters with different role, service, instance group, or instance configurations in the current cluster.

   You can click  next to a parameter value and view the differences in the dialog box that is displayed.

b. To change the value of a parameter, click  to cancel the configuration difference or manually adjust the parameter value, click **OK**, and then click **Save**. In the dialog box that is displayed, click **OK**.

- To check expired configurations:
  a. Click **Expired Configurations**. Expired configuration items in the current cluster are displayed.
  b. You can filter services using the service filter box in the upper part of the page to view expired configurations of different services. Alternatively, you can enter keywords in the search box.
  c. Expired configuration items do not take effect completely. Restart the services or instances whose configurations have expired in a timely manner.
- To view historical configuration records:
  a. Click **Historical Configurations**. The historical configuration change records of the current cluster are displayed. You can view details about parameter value changes, including the service to which the parameter belongs, parameter values before and after the modification, and parameter files.
  b. To restore a configuration change, click **Restore Configuration** in the **Operation** column of the target record. In the dialog box that is displayed, click **OK**.

NOTE

Some configuration items take effect only after the corresponding services are restarted. After the configurations are saved, restart the services or instances whose configurations have expired in a timely manner.

----End

8.2.1.7 Managing Static Service Pools

8.2.1.7.1 Static Service Resources

Overview

A cluster allocates static service resources to services Flume, HBase, HDFS, and YARN. The total volume of computing resources allocated to each service is fixed, and they are static. A tenant can exclusively use or share a service to obtain the resources required for running this service.

Static Service Pool

Static service pools are used to specify service resource configurations.

Static service pools centrally manage resources that can be used by each service.

- Limits the total number of resources that can be used by each service. Specifically, the total number of CPU, I/O, and memory resources can be configured on the nodes where services Flume, HBase, HDFSIoTDB, Kafka (Kafka supports static service pools only in MRS 3.2.0 or later), and YARN are deployed.
- Isolates the resources of services in a cluster from those of other services. In this way, the load of one service has very limited impact on other services.

Scheduling Mechanism

The time-based dynamic resource scheduling mechanism enables different volumes of static resources to be configured for services at different time, optimizing service running environments and improving the cluster efficiency.

In a complex cluster environment, multiple services share resources in the cluster, but the resource service period of each service may be different.

The following use a bank customer as an example:

- The HBase query service is heavy in the daytime.
- The query service is light, but the Hive analysis service is heavy at night.

If fixed resources are allocated to each service, the following problems may occur:

- The query service cannot obtain sufficient resources while the resources for the analysis service are idle in the daytime.
- The analysis service cannot obtain sufficient resources while the resources for the query service are idle at night.

As a result, the cluster resource utilization is low and the service capability is weak. Resolve the problem in the following ways:

- Sufficient resources need to be configured for HBase in the daytime.
- Sufficient resources need to be configured for Hive at night.

The time-based dynamic scheduling mechanism can efficiently utilize resources and run tasks.

8.2.1.7.2 Configuring Cluster Static Resources

Scenario

You can adjust resource base on FusionInsight Manager and customize resource configuration groups if you need to control service resources used on each node in a cluster or the available CPU or I/O quotas on each node at different time segments.

Impact on the System

- After a static service pool is configured, the configuration status of affected services is displayed as **Expired**. You need to restart the services. Services are unavailable during restart.
- After a static service pool is configured, the maximum number of resources used by each service and role instance cannot exceed the upper limit.

Procedure

Modify the Resource Adjustment Base

Step 1 On FusionInsight Manager, choose **Cluster**, click the name of the target cluster, and choose **Static Service Pool Configurations**.

Step 2 Click **Configuration** in the upper right corner. The page for configuring resource pools is displayed.

Step 3 Change the values of **CPU (%)** and **Memory (%)** in the **System Resource Adjustment Base** area.

Modifying the system resource adjustment base changes the maximum physical CPU and memory usage on nodes by services. If multiple services are deployed on the same node, the maximum physical resource usage of all services cannot exceed the adjusted CPU or memory usage.

Step 4 Click **Next**.

To modify parameters again, click **Previous**.

Modify the Default Resource Configuration Group

Step 5 Click **default**. In the **Configure weight** table, set **CPU LIMIT(%)**, **CPU SHARE(%)**, **I/O(%)**, and **Memory(%)** for each service.

Figure 8-13 Weight Configuration

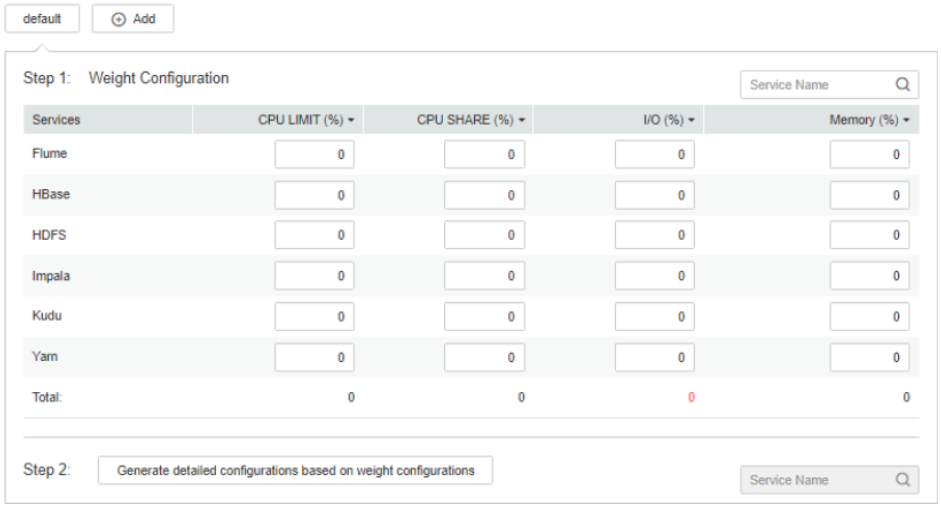

NOTE

- The sum of **CPU LIMIT(%)** and **CPU SHARE(%)** used by all services can exceed 100%.
- The sum of **I/O(%)** used by all services can exceed 100% but cannot be 0.
- The sum of **Memory(%)** used by all services can be greater than, smaller than, or equal to 100%.
- **Memory(%)** cannot take effect dynamically and can only be modified in the default configuration group.
- **CPU LIMIT(%)** is used to configure the ratio of the number of CPU cores that can be used by a service to those can be allocated to related nodes.
- **CPU SHARE(%)** is used to configure the ratio of the time when a service uses a CPU core to the time when other services use the CPU core. That is, the ratio of time when multiple services compete for the same CPU core.

Step 6 Click **Generate detailed configurations based on weight configurations**. FusionInsight Manager generates the actual values of the parameters in the default weight configuration table based on the cluster hardware resources and allocation information.

Step 7 Click **OK**.

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

Add a Customized Resource Configuration Group

Step 8 Determine whether to automatically adjust resource configurations at different time segments.

- If yes, go to **Step 9**.
- If no, use the default configurations, and no further action is required.

Step 9 Click **Configuration**, change the system resource adjustment base values, and click **Next**.

Step 10 Click **Add** to add a resource configuration group.

Figure 8-14 Adding a resource configuration group

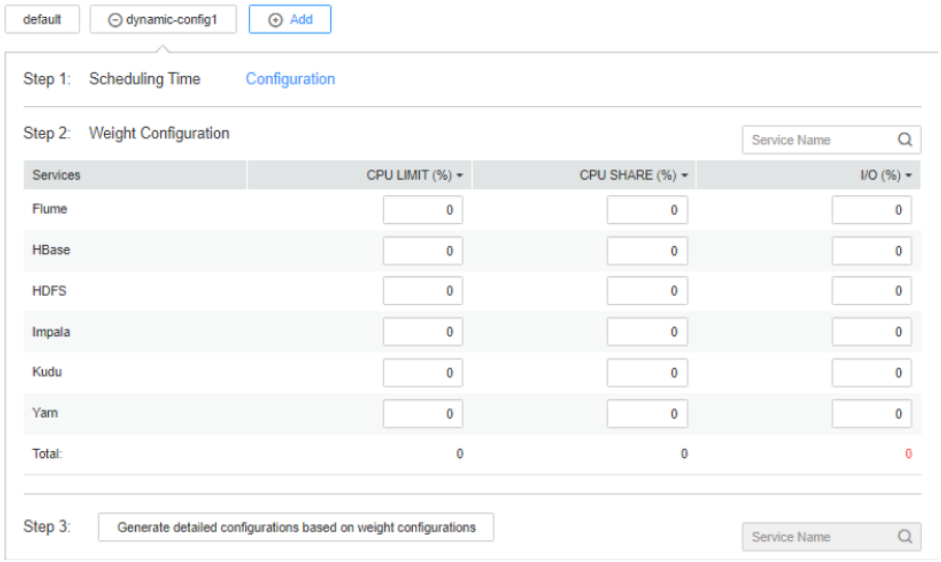

Step 11 In **Step 1: Scheduling Time**, click **Configuration**.

The page for configuring the time policy is displayed.

Modify the following parameters based on service requirements and click **OK**.

- **Repeat**: If this parameter is selected, the customized resource configuration is applied repeatedly based on the scheduling period. If this parameter is not selected, set the date and time when the configuration of the group of resources can be applied.
- **Repeat Policy**: The available values are **Daily**, **Weekly**, and **Monthly**. This parameter is valid only when **Repeat** is selected.
- **On**: indicates the time period between the start time and end time when the resource configuration is applied. Set a unique time range. If the time range overlaps with that of an existing group of resource configuration, the time range cannot be saved.

NOTE

- The default group of resource configuration takes effect in all undefined time segments.
- The newly added resource group is a parameter set that takes effect dynamically in a specified time range.
- The newly added resource group can be deleted. A maximum of four resource configuration groups that take effect dynamically can be added.
- Select a repetition policy. If the end time is earlier than the start time, the resource configuration ends in the next day by default. For example, if a validity period ranges from 22:00 to 06:00, the customized resource configuration takes effect from 22:00 on the current day to 06:00 on the next day.
- If the repetition policy types of multiple configuration groups are different, the time ranges can overlap. The policy types are listed as follows by priority from low to high: daily, weekly, and monthly. The following is an example. There are two resource configuration groups using the monthly and daily policies, respectively. Their application time ranges in a day overlap as follows: 04:00 to 07:00 and 06:00 to 08:00. In this case, the configuration of the group that uses the monthly policy prevails.
- If the repetition policy types of multiple resource configuration groups are the same, the time ranges of different dates can overlap. For example, if there are two weekly scheduling groups, you can set the same time range on different day for them, such as to 04:00 to 07:00, on Monday and Wednesday, respectively.

Step 12 Modify the resource configuration of each service in **Step 2: Weight Configuration**.

Step 13 Click **Generate detailed configuration**. FusionInsight Manager generates the actual values of the parameters in the default weight configuration table based on the cluster hardware resources and allocation information.

Step 14 Click **OK**.

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

----End

8.2.1.7.3 Viewing Cluster Static Resources

Scenario

The big data management platform can manage and isolate service resources that are not running on YARN using static service resource pools. The system supports time-based automatic adjustment of static service resource pools. This enables the

cluster to automatically adjust the parameter values at different periods to ensure more efficient resource utilization.

System administrators can view the monitoring indicators of resources used by each service in the static service pool on FusionInsight Manager. The monitoring indicators are as follows:

- CPU usage of services
- Total disk I/O read rate of services
- Total disk I/O write rate of services
- Total used memory of services

Procedure

Step 1 On FusionInsight Manager, choose **Cluster**, click the name of the target cluster, and choose **Static Service Pool Configurations**.

Step 2 In the configuration group list, click a configuration group, for example, **default**.

Step 3 Check the system resource adjustment base values.

- **System Resource Adjustment Base** indicates the maximum volume of resources that can be used by each node in the cluster. If a node has only one service, the service exclusively occupies the available resources on the node. If a node has multiple services, all services share the available resources on the node.
- **CPU** indicates the maximum number of CPUs that can be used by services on a node.
- **Memory** indicates the maximum memory that can be used by services on a node.

Step 4 In **Chart**, view the metric data of the cluster service resource usage.

NOTE

- You can click **Add Service to Chart** to add static service resource data of specific services (up to 12 services) to the chart.
- For details about how to manage a chart, see **Managing Monitoring Metric Reports**.

----End

8.2.1.8 Managing Clients

8.2.1.8.1 Managing a Client

Scenario

FusionInsight Manager supports unified management of cluster client installation information. After a user downloads and installs a client, FusionInsight Manager automatically records information about the installed (registered) client to facilitate query and management. In addition, you can manually add or modify the information about clients that are not automatically registered, for example, clients installed in earlier versions.

Procedure

View client information.

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster**, click the name of the desired cluster, and choose **Client Management** to view information about clients installed in the cluster.

You can view the IP address, installation path, component list, registration time, and installation user of the node where the client is located.

When the client is downloaded and installed in the cluster of the latest version, the client information is automatically registered.

Figure 8-15 Client information

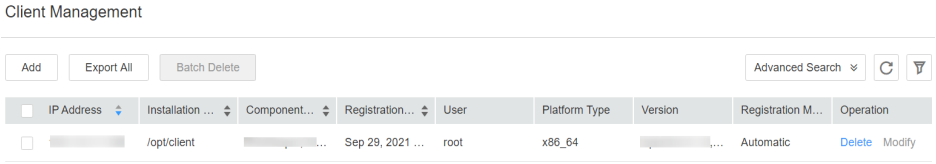

Add client information.

Step 3 To manually add information about an installed client, click **Add** and manually add the IP address, installation path, user, platform information, and registration information of the client as prompted.

Step 4 Configure the client information and click **OK**.

Modify client information.

Step 5 Modify information about the manually registered client.

On the **Client Management** page, select the target client and click **Modify**. After modifying the information, click **OK**.

Delete client information.

Step 6 On the **Client Management** page, select the target client and click **Delete**. In the displayed dialog box, click **OK**.

To delete multiple clients, select the all of them and click **Batch Delete**. In the displayed dialog box, click **OK**.

Export client information.

Step 7 On the **Client Management** page, click **Export All** to export information about all registered clients to the local PC.

NOTE

On the **Client Management** page, only components that have clients are displayed in the component list. Therefore, some components that do not have clients and have special components are not displayed.

The following components are not displayed:

LdapServer, KrbServer, DBService, Hue, MapReduce, and Flume

----End

8.2.1.8.2 Batch Upgrading Clients

Scenario

The client package downloaded from FusionInsight Manager contains the client batch upgrade tool. When multiple clients need to be upgraded after the cluster upgrade or scale-out, you can use this tool to upgrade the clients in batches with a few clicks. In addition, the tool provides the lightweight function of batch updating the **/etc/hosts** file on the nodes where the clients are located.

Procedure

Prepare for the client upgrade.

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > **Overview** > **More** > **Download Client** (for MRS 3.3.0 or later, click **Download Client** on the **Homepage** page) to download the client to a specified directory on the server.

For details, see **Downloading the Client**.

Decompress the downloaded client package and find the **batch_upgrade** directory, for example, **/tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_ClientConfig/batch_upgrade**.

Step 3 Choose **Cluster** > **Client Management**. On the **Client Management** page, click **Export All** to export all client information to the local PC.

Step 4 Decompress the exported client information and upload the **client-info.cfg** file to the **batch_upgrade** directory.

Step 5 Supplement the password in the **client-info.cfg** file by referring to **Reference Information**.

Upgrade clients in batches.

Step 6 Run the **sh client_batch_upgrade.sh -u -f /tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_Client.tar -g /tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_ClientConfig/batch_upgrade/client-info.cfg** command to perform the upgrade.

NOTICE

You are advised to delete the **client-info.cfg** file as soon as possible after the upgrade because the password has been configured.

Step 7 After the upgrade is complete, verify the upgrade result by running the **sh client_batch_upgrade.sh -c** command.

Step 8 If the client is faulty, run the **sh client_batch_upgrade.sh -s** command to roll back the client.

NOTE

- The client batch upgrade tool moves the original client to the backup directory, and then uses the client package specified by the **-f** parameter to install the client. Therefore, if the original client contains customized content, manually save the customized content from the backup directory or move the customized content to the client directory after the upgrade before running the **-c** command. The backup path on the client is *{Original client path}*-**backup**.
- The **-u** command is the prerequisite for the **-c** and **-s** commands. You can run the **-c** command to commit the upgrade or the **-s** command to perform a rollback only after the **-u** command is executed to perform an upgrade.
- You can run the **-u** command multiple times to upgrade only the clients that fail to be upgraded.
- The client batch upgrade tool also supports the clients of earlier versions.
- When upgrading a client installed by a non-root user, ensure that the user has the read and write permissions on the directory where the client is located and the parent directory on the target node. Otherwise, the upgrade will fail.
- The client package specified by the **-f** parameter must be a full client package. The client packages of a single component or some components cannot be used as the input.

----End

Reference Information

Before upgrading clients in batches, you need to manually configure the user password for remotely logging in to the client node.

Run the **vi client-info.cfg** command to add a user password.

Example:

```
clientIp,clientPath,user,password
10.10.10.100,/home/omm/client /home/omm/client2,omm,Password
```

The fields in the configuration file are as follows:

- **clientIp**: indicates the IP address of the node where the client is located.
- **clientPath**: indicates the client installation path. Multiple paths are separated by spaces. Note that the path cannot end with a slash (/).
- **user**: indicates the username of the node.
- **password**: indicates the user password of the node.

NOTE

- If the execution fails, view the **node.log** file in the **work_space/log_*XXX*** directory.
- There can be security risks if a configuration file contains the authentication password. You are advised to delete the configuration file or use other secure methods to keep the password.

8.2.1.8.3 Updating the hosts File in Batches

Scenario

The client package downloaded from FusionInsight Manager contains the client batch upgrade tool. This tool provides the function of upgrading clients in batches

and the lightweight function of batch updating the **/etc/hosts** file on the node where the client is located.

Prerequisites

You have made preparations for the upgrade. For details, see "Prepare for the client upgrade." in **Batch Upgrading Clients**.

Updating the hosts File in Batches

Step 1 Check whether the user configured for the node where the **/etc/hosts** file needs to be updated is **root**.

- If yes, go to **Step 2**.
- If no, change the user to **root** and go to **Step 2**.

Step 2 Run the **sh client_batch_upgrade.sh -r -f /tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_Client.tar -g /tmp/FusionInsight-Client/FusionInsight_Cluster_1_Services_ClientConfig/batch_upgrade/client-info.cfg** command to batch update the **/etc/hosts** file on the nodes where the client resides.

NOTE

- When you batch update the **/etc/hosts** file, the entered client package can be a complete client package or a client package that contains only configuration files (recommended).
- The user configured for the host where the **/etc/hosts** file needs to be updated must be **root**. Otherwise, the update fails.

----End

8.2.2 Managing a Service

8.2.2.1 Overview

Dashboard

Log in to FusionInsight Manager. Choose **Cluster**, click the name of the desired cluster, and choose **Services**. The service management page is displayed, including the functional area and service list.

Figure 8-16 Service management page

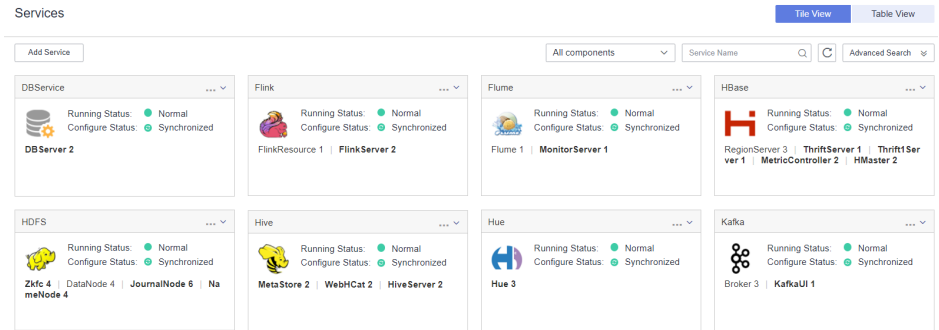

Functional Area

In the functional area of the service management page, you can select a view type and filter and search for services by service type. You can use the advanced search to select required services based on the running status and configuration status.

Service List

The service list on the service management page contains all installed services in the cluster. If the tile view is selected, the services will be displayed in pane style. If you select the list view, the services will be displayed in a table.

NOTE

In this section, the **Tile View** is used by default.

The service list displays the running status, configuration status, role type, and number of instances of each service. On this page, you can perform some service maintenance tasks, such as starting, stopping, and restarting services.

Table 8-3 Service running status

Status	Description
Normal	Indicates that the service is running properly.
Faulty	Indicates that the service cannot run properly.
Partially Healthy	Indicates that some enhanced functions of the service are abnormal.
Not started	Indicates that the service is stopped.
Unknown	Indicates that the initial status of the service cannot be detected.
Starting	Indicates that the service is being started.
Stopping	Indicates that the service is being stopped.
Failed to start	Indicates that the service fails to be started.
Failed to stop	Indicates that the service fails to be stopped.

NOTE

- If the running status of a service is **Faulty**, an alarm is generated. Rectify the fault based on the alarm information.
- HBase, Hive, Spark, and Loader may be in the **Subhealthy** state.
  - If Yarn is installed but is abnormal, HBase is in the **Subhealthy** state.
  - If HBase is installed but is abnormal, Hive, Spark, and Loader are in the **Subhealthy** state.

Table 8-4 Service configuration status

Status	Description
Synchronized	Indicates that all service parameter settings have taken effect in the cluster.
Expired	Indicates that the latest configuration is not synchronized and does not take effect after the service parameters are modified. You need to synchronize the configurations and restart the services. You can click next to **Configuration Status** to view expired configuration items.
Failed	Indicates that a communication or read/write exception occurs during the parameter configuration synchronization. Use **Synchronize Configuration** to rectify the fault.
Synchronizing	Indicates that the service parameter configuration is being synchronized.
Unknown	Indicates that the initial status of the service cannot be detected.

You can click a service in the service list to perform simple maintenance and management operations on the service, as described in **Table 8-5**.

Table 8-5 Basic maintenance and management

Menu Item on the UI	Description
Start Service	Start a specified service in the cluster.
Stop Service	Stop a specified service in the cluster.
Restart Service	Restart a specified service in the cluster. **NOTE** If a service is restarted, other services that depend on this service will be unavailable. Therefore, select **Restart upper-layer services**. Determine whether to perform this operation based on the displayed service list. Services are restarted one by one due to their dependency. **Table 8-6** describes the restart duration of a single service.
Service Rolling Restart	Restart a specified service in the cluster without interrupting services. For details about the parameter settings, see **Table 8-2**.

Menu Item on the UI	Description
Synchronize Configuration	• Enable new configuration parameters for a specified service in the cluster. • Distribute new configuration parameters for services whose **Configuration Status** is **Expired**. **NOTE** After some services are synchronized, restart the services for the settings to take effect.

Table 8-6 Restart duration for reference

Service	Restart Duration	Startup Duration	Remarks
IoTDB	3 min	IoTDBServer: 3 min	-
CDL	2 min	CDLConnector: 1 min CDLService: 1 min	-
ClickHouse	4 min	ClickHouseServer: 2 min ClickHouseBalancer: 2 min	-
HDFS	10min+x	NameNode: 4 min + x DataNode: 2 min JournalNode: 2 min Zkfc: 2 min	*x* indicates the NameNode metadata loading duration. It takes about 2 minutes to load 10,000,000 files. For example, *x* is 10 minutes for 50 million files. The startup duration fluctuates with reporting of DataNode data blocks.
Yarn	5 min + x	ResourceManager: 3 min + x NodeManager: 2 min	*x* indicates the time required for restoring ResourceManager reserved tasks. It takes about 1 minute to restore 10,000 reserved tasks.
MapReduce	2 min + x	JobHistoryServer: 2 min + x	*x* indicates the scanning duration of historical tasks. It takes about 2.5 minutes to scan 100,000 tasks.
ZooKeeper	2 min + x	quorumpeer: 2 min + x	*x* indicates the duration for loading znodes. It takes about 1 minute to load 1 million znodes.

Service	Restart Duration	Startup Duration	Remarks
Hive	3.5 min	HiveServer: 3 min MetaStore: 1 min 30s WebHcat: 1 min Hive service: 3 min	-
Spark2x	5 min	JobHistory2x: 5 min SparkResource 2x: 5 min JDBCServer2x: 5 min	-
Flink	4 min	FlinkResource: 1 min FlinkServer: 3 min	-
Kafka	2 min + x	Broker: 1 min + x	*x* indicates the data restoration duration. It takes about 2 minutes to start 20,000 partitions for a single instance.
Storm	6 min	Nimbus: 3 min UI: 1 min Supervisor: 1 min Logviewer: 1 min	-
Flume	3 min	Flume: 2 min MonitorServer: 1 min	-
Doris	2 min	FE: 1min BE: 1min DBroker: 1min	-

8.2.2.2 Service Management Operations

8.2.2.2.1 Service Details Page

Overview

Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services**. In the service list, click the specified service name to go to the service details page, including the **Dashboard**, **Instance**, **Instance Groups** and **Configurations** tab pages as well as function areas. For some services, the custom management tool page can be displayed.For details about the supported management tools, see **Table 8-7**..

Table 8-7 Customized management tools

Tool	Service	Description
Flume configuration tool	Flume	Configures collection parameters for the Flume server and client.
Flume client management tool	Flume	Views the monitoring information about the Flume client.
Kafka topic monitoring tool	Kafka	Monitors and manages Kafka topics.

The **Dashboard** page is the default page, which contains the basic information, role list, dependency table, and monitoring chart, and more. You can manage services in the upper right corner. For details about basic service management, such as starting, stopping, rolling restart, and synchronization configuration, see **Table 8-5**. For details about other service management operations, see **Table 8-8**.

Table 8-8 Service management operations

Navigation Path	Description
More > **Health Check**	Performs a health check for the current service. The health check items include the health status of each check object, related alarms, and user-defined monitoring indicators. The check result is not the same as the values of **Running Status** displayed on the GUI. To export the result of the health check, click **Export Report** in the upper left corner of the checklist. If you find any problem, click **View Help**.

Navigation Path	Description
More > **Download Client**	Download the default client that contains only specific services and perform management operations, run services, or perform secondary development on the client. For details, see **Downloading the Client**.
More > **Change Service Name**	Changes the name of the current service.
More > **Perform** *XX* **Switchover**	For details, see **Performing Active/Standby Switchover of a Role Instance**.
More > **Enter/Exit Maintenance Mode**	Configures a service to enter/exit the maintenance mode.
Configurations > **Import/Export**	In the scenario where services are migrated to a new cluster or the same services are deployed again, you can import or export all configuration data of a specific service to quickly copy the configuration results.

Basic Information Area

The basic information area on the **Dashboard** tab page contains the basic status data of the service, including the running status, configuration details, version, and key information of the service. If the service supports the open-source web UIs, you can access the open-source web UIs by clicking the links in the basic information area.

NOTE

In the current version, user **admin** does not have the permission to access all the service functions provided on the open source web UI. Create a component service administrator to access the WebUI address.

Role List

The role list on the **Dashboard** tab page contains all roles of the service. The role list displays the running status and the number of instances of each role.

Dependency

The dependency relationship table on the **Dashboard** tab page displays the services on which the current service depends and other services that depend on the service.

Historical Records of Alarms and Events

The alarm and event history area displays the key alarms and events reported by the current service. Up to 20 historical records are displayed.

Chart

The chart area is displayed on the right of the **Dashboard** tab page and contains the key monitoring indicator report of the service. You can customize the monitoring report that is displayed in the chart area, view the description of the monitoring metrics, or export the monitoring data. For a customized resource contribution chart, you can zoom in on the chart and switch between the trend chart and distribution chart.

NOTE

Some services in the cluster provide service-level resource monitoring items. For details, see **Resource Monitoring**.

8.2.2.2.2 Performing Active/Standby Switchover of a Role Instance

Scenario

Some service roles are deployed in active/standby mode. If the active instance needs to be maintained and cannot provide services, or other maintenance is required, you can manually trigger an active/standby switchover.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 On the service details page, expand the **More** drop-down list and select **Perform** *Role Instance* **Switchover**.

Step 5 In the displayed dialog box, enter the password of the current login user and click **OK**.

Step 6 In the displayed dialog box, click **OK** to perform active/standby switchover for the role instance.

NOTE

- The Manager component package only supports the active/standby switchover of DBService role instances.
- The HD component package supports the active/standby switchover of the following service role instances: HDFS, YARN, Storm, HBase, and MapReduce.
- When an active/standby switchover is performed for a NameNode on HDFS, a NameService must be set.
- The Porter component package only supports the active/standby switchover of Loader role instances.
- This function cannot be used for other role instances.

----End

8.2.2.2.3 Resource Monitoring

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**, and click **Resource**. The resource monitoring page is displayed.

Some services in the cluster provide service-level resource monitoring metrics. By default, the monitoring data of the latest 12 hours is displayed. You can click to customize a time range. Time range options are **12h**, **1d**, **1w**, and **1m**. You can click to export the corresponding report information. If a monitoring item has no data, the report cannot be exported. **Table 8-9** lists the services and monitoring items that support resource monitoring.

Table 8-9 Service resource monitoring

Service	Metrics	Description
HDFS	Resource Usage (by Tenant)	• Collects statistics on HDFS resource usage by tenant. • Views the metrics **Capacity** or **Number of File Objects**.
	Resource Usage (by User)	• Collects statistics on HDFS resource usage by user. • Views the metrics **Used Capacity** or **Number of File Objects**.
	Resource Usage (by Directory)	• Collects statistics on HDFS resource usage by directory. • Views the metrics **Used Capacity** or **Number of File Objects**. • You can click to configure space monitoring. Alternatively, you can specify an HDFS file system directory for monitoring.
	Resource Usage (by Replica)	• Collects statistics on HDFS resource usages by replica count. • Views the metrics **Used Capacity** or **File Count**.
	Resource Usage (by File Size)	• Collects statistics on HDFS resource usages by file size. • Views the metrics **Used Capacity** or **File Count**.
	Recycle Bin (by User)	• Collects statistics on the usage of the HDFS recycle bin by user. • Views the metrics **Recycle Bin Capacity** or **Number of File Objects**.
	Operation Count	• Collects the number of operations in HDFS.

Service	Metrics	Description
	Automatic Balancer	• Collects statistics on the execution speed of HDFS automatic balancer and the total capacity of the current balancer migration.
	NameNode RPC Open Connections (by User)	• Displays the number of connections of each user in the Client RPC requests connected to NameNodes.
	Slow DataNodes	Displays DataNode that transmits or processes data slowly in the cluster.
	Slow Disks	Displays the disk that processes data slowly on the DataNode in the cluster.
HBase	Operation Requests in Tables	Displays the number of PUT, DELETE, GET, SCAN, INCREMENT, and APPEND operation requests in all tables on all RegionServers.
	Operation Requests on RegionServers	Displays the number of PUT, DELETE, GET, SCAN, INCREMENT, and APPEND operation requests and number of all operation requests in RegionServer.
	Operation Requests for Service	Displays the number of PUT, DELETE, GET, SCAN, INCREMENT, and APPEND operation requests in all regions on RegionServers.
	HFiles on RegionServers	Displays the number of HFiles in all RegionServers.
HetuEngine	Coordinator Resource Usage	Displays the coordinator resource usage in the selected queue.
	Coordinator Resource Usage Ratio	Displays the coordinator resource usage in the selected queue.
	Worker Resource Usage	Displays the worker resource usage in the selected queue.
	Worker Resource Usage Ratio	Displays the worker resource usage in the selected queue.
	Number of Coordinators and Workers	Displays the number of coordinators and workers in the selected queue.
Hive	HiveServer2-Background-Pool Threads (by IP)	Displays the number of HiveServer2-Background-Pool threads of top users. These threads are measured and displayed in a measurement period.

Service	Metrics	Description
	HiveServer2-Handler-Pool Threads (by IP)	Displays the number of HiveServer2-Handler-Pools of top users collected and displayed in a period.
	Used MetaStore Number (by IP)	Collects statistics on and displays the MetaStore usage of top users in a period.
	Number of Hive jobs	Displays the number of user-related jobs collected by Hive in a period.
	Number of Files Accessed in the Split Phase	Displays the number of files accessed by the underlying file storage system (HDFS by default) in the Split phase in a period.
	Hive Basic Operation Time	Collects time for creating a directory (mkdirTime), creating a file (touchTime), writing a file (writeFileTime), renaming a file (renameTime), moving a file (moveTime), deleting a file (deleteFileTime), and deleting a directory (deleteCatalogTime) in a period of time.
	Table Partitions	Displays the number of partitions in all Hive tables, which is displayed in the following format: *database # table name*, *number of table partitions*.
	HQL Map Count	Collects statistics on HQL statements executed in a period and the number of Map statements invoked during the execution. The displayed information includes users, HQL statements, and the number of Map statements.
	HQL Access Statistics	Displays the number of HQL access times in a period.
Kafka	Kafka Disk Usage Distribution	Displays the disk usage distribution statistics of the Kafka cluster.
Spark/ Spark2x	HQL Access Statistics	Collects HQL access statistics in a period, including the username, HQL statement, and HQL statement execution times.
Yarn	Used resources (by task)	• Displays the number of CPU cores and memory used by a task. • Views the metrics **By memory** or **By CPU**.

Service	Metrics	Description
	Resource usage (by tenant)	• Displays the number of CPU cores and memory used by a tenant. • Views the metrics **By memory** or **By CPU**.
	Resource usage ratio (by tenant)	• Displays the ratio of the number of CPU cores to the memory used by a tenant. • Views the metrics **By memory** or **By CPU**.
	Task Duration Ranking	Displays Yarn tasks sorted by time consumption.
	ResourceManager RPC Open Connections (by User)	Displays the number of client RPC connections to ResourceManager by user.
	Operation Count	Collects statistics on the number and proportion of operations corresponding to each Yarn operation type.
	Ranking of Tasks in a Queue by Resource Usage	• Displays the resources consumed by the tasks running in a queue after the queue (tenant) is selected on the GUI. • Views the metrics **By memory** or **By CPU**.
	Ranking of Users in a Queue by Resource Usage	• Displays the resources consumed by the users who are running tasks in the queue after a queue (tenant) is selected on the GUI. • Views the metrics **By memory** or **By CPU**.
ZooKeeper	Used Resources (By Second-Level Znode)	• Displays the ZooKeeper level-2 znode resource status. • Views the metrics **By Znode quantity** or **By capacity**.
	Number of Connections (by Client IP Address)	Displays the ZooKeeper client connection resource status.

8.2.2.2.4 Collecting Stack Information

Scenario

To meet actual service requirements, the cluster administrator can collect stack information about a specified role or instance on FusionInsight Manager, save the

information to a local directory, and download the information. The following information can be collected:

1. jstack information.
2. jmap -histo information.
3. jmap -dump information.
4. Thr jstack and jmap-histo information can be collected continuously for comparison.

Procedure

Collecting stack information

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Cluster**, click the name of the desired cluster, click **Services**, and click the target service.

Step 3 On the displayed page, Choose **More** > **Collect Stack Information**.

NOTE

- To collect stack information of multiple instances, go to the instance list, select the desired instances in the instance list and choose **More** > **Collect Stack Information**.
- To collect stack information of a single instance, click the desired instance and choose **More** > **Collect Stack Information**.

Step 4 In the displayed dialog box, select the desired role and content, configure advanced options (retain the default settings if there is no special requirement), and click **OK**.

Figure 8-17 Collecting stack information

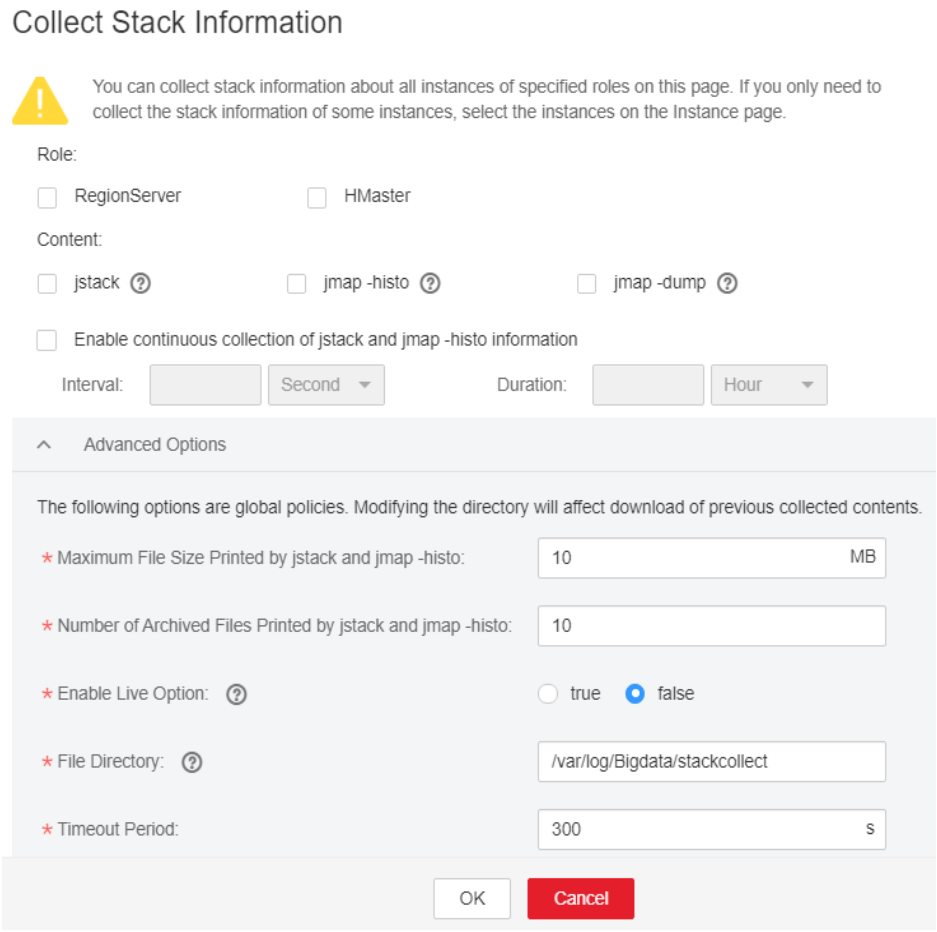

Step 5 After the collection is successful, click **Download**.

Downloading Stack Information

Step 6 Click **Cluster**, click the name of the desired cluster, click **Services**, and click the target service. Choose **More** > **Download Stack Information** in the upper right corner.

Step 7 Select the desired role and content and click **Download** to download the stack information to the local PC.

Figure 8-18 Downloading stack information

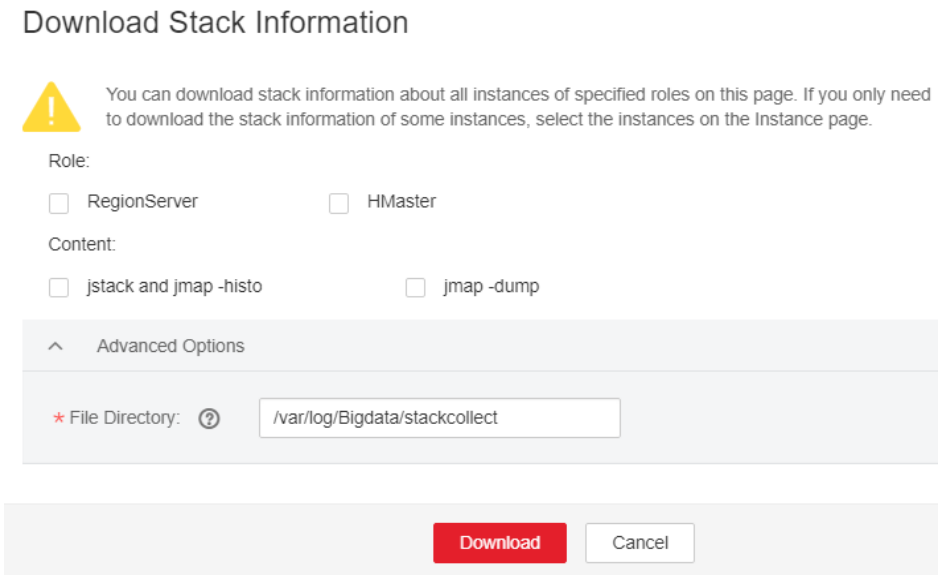

Clearing stack information

Step 8 Click **Cluster**, click the name of the desired cluster, click **Services**, and click the target service.

Step 9 Choose **More** > **Clear Stack Information** in the upper right corner.

Step 10 Select the desired role and content and configure **File Directory**. Click **OK**.

Figure 8-19 Clearing stack information

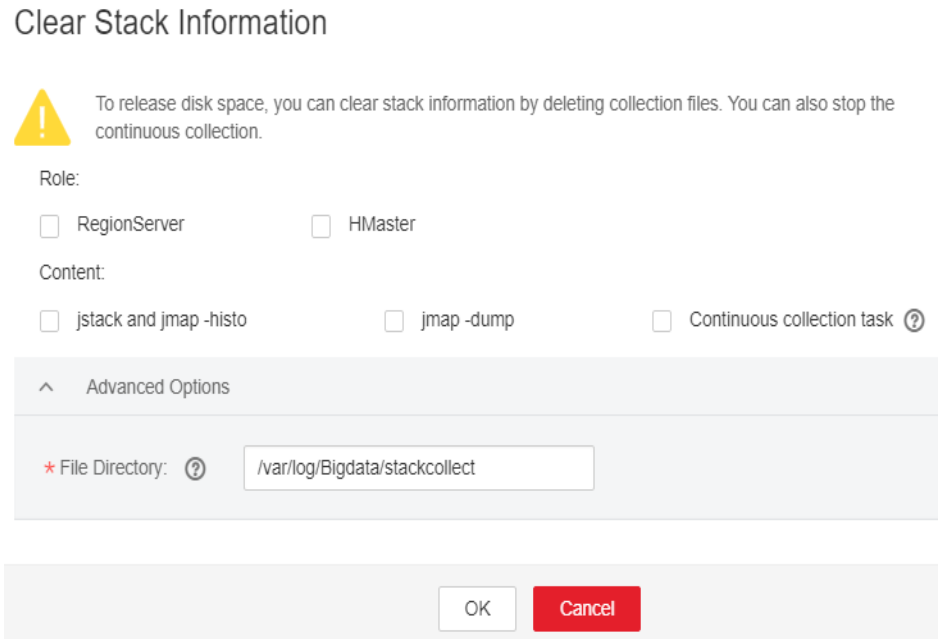

----End

8.2.2.2.5 Switching Ranger Authentication

Scenario

By default, the Ranger service is installed and Ranger authentication is enabled for a newly installed cluster in security mode. You can set fine-grained security access policies for accessing component resources through the permission plug-in of the component. If Ranger authentication is not required, the cluster administrator can manually disable Ranger authentication on the service page. After Ranger authentication is disabled, the system continues to perform permission control based on the role model of FusionInsight Manager when accessing component resources.

In a cluster upgraded from an earlier version, Ranger authentication is not used by default when users access component resources. The cluster administrator can manually enable Ranger authentication after installing the Ranger service.

NOTE

- In a cluster in security mode, the following components support Ranger authentication: HDFS, YARN, Kafka, Hive, HBase, Storm, Impala, HetuEngine, CDL, and Spark/Spark2x.
- In a cluster in non-security mode, Ranger supports permission control on component resources based on OS users. The following components support Ranger authentication: HBase, HDFS, Hive, Spark/Spark2x, and YARN.
- After Ranger authentication is enabled, all authentication of the component will be managed by Ranger. The permissions set by the original authentication plug-in will become invalid (The ACL rules of HDFS and YARN components still take effect). Exercise caution when performing this operation. You are advised to deploy permissions on Ranger in advance.
- After Ranger authentication is disabled, all authentication of the component will be managed by the permission plug-in of the component. The permission set on Ranger will become invalid. Exercise caution when performing this operation. You are advised to deploy permissions on Manager in advance.

Enabling Ranger Authentication

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 On the service details page, expand the **More** drop-down list and select **Enable Ranger**.

Step 5 In the displayed dialog box, enter the password of the current login user and click **OK**.

Step 6 In the service list, restart the service whose configuration has expired.

----End

Disabling Ranger Authentication

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 On the service details page, expand the **More** drop-down list and select **Disable Ranger**.

Step 5 Enter the password of the current login user and click **OK**. In the displayed dialog box, click **OK**.

Step 6 In the service list, restart the service whose configuration has expired.

----End

8.2.2.3 Service Configuration

8.2.2.3.1 Modifying Service Configuration Parameters

Scenario

To meet actual service requirements, cluster administrators can quickly view and modify default service configurations on FusionInsight Manager. Configure parameters based on the information provided in the configuration description.

NOTE

The parameters of DBService cannot be modified when only one DBService role instance exists in the cluster.

Impact on the System

- After configuring properties of a service, you need to restart the service if the service status is **Expired**. The service is unavailable during the restart.
- After the service configuration parameters are modified and then take effect after restart, you need to download and install the client again or download the configuration file to update the client.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 Click **Configuration**.

The **Basic Configuration** page is displayed by default. To modify more parameters, click the **All Configurations** tab. The navigation tree displays all configuration parameters of the service. The level-1 nodes in the navigation tree are service names or role names. The parameter category is displayed after the level-1 node is expanded.

As shown in the following figure, the first node **LdapServer** indicates the service name, and the second node **SlapdServer** indicates the role name. The configuration parameter displayed takes effect for all instances of the role and the service.

Figure 8-20 Configuration parameter navigation tree

Step 5 In the navigation tree, select the specified parameter category and change the parameter values on the right.

NOTE

Select a port parameter value from the value range on the right. Ensure that all parameter values in the same service are within the value range and are unique. Otherwise, the service fails to be started.

If you are not sure about the location of a parameter, you can enter the parameter name in search box in the upper right corner. The system searches for the parameter in real time and displays the result.

Step 6 Click **Save**. In the confirmation dialog box, click **OK**.

Wait until the message "Operation succeeded." is displayed. Click **Finish**.

The configuration is modified.

NOTE

- To update the queue configuration of the YARN service without restarting service, choose **More** > **Refresh Queue** to update the queue for the configuration to take effect.
- During configuration of the **flume.config.file** parameter, you can upload and download files. After a configuration file is uploaded, the old file will be overwritten. If the configuration is not saved and the service is restarted, the configuration does not take effect. Save the configuration in time.
- If you need to restart the service for the configuration to take effect after modifying service configuration parameters, choose **More** > **Restart Service** in the upper right corner of the service page.
- If the  is displayed before a parameter, this parameter takes effect dynamically. After the configuration is saved, the parameter value is automatically updated to the configuration file. This function is supported only in MRS 3.2.0 or later.

----End

8.2.2.3.2 Modifying Custom Configuration Parameters of a Service

Scenario

All open source parameters can be configured for all MRS cluster components. Parameters used in some key application scenarios can be modified on FusionInsight Manager, and some parameters of open source features may not be configured for some component clients. To modify the component parameters that are not directly supported by Manager, cluster administrators can add new parameters for components using the configuration customization function on Manager. Newly added parameters are saved in component configuration files and take effect after restart.

Impact on the System

- After configuring properties of a service, you need to restart the service if the service status is **Expired**. The service is unavailable during the restart.
- After the service configuration parameters are modified and then take effect after restart, you need to download and install the client again or download the configuration file to update the client.

Prerequisites

Cluster administrators have fully understood the meanings of the parameters to be added, configuration files to take effect, and the impact on components.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 Click **Configuration** and click **All Configurations**.

Step 5 In the navigation tree on the left, locate a level-1 node and select **Customization**. The system displays the customized parameters of the current component.

The configuration files that save the newly added custom parameters are displayed in the **Parameter File** column. Different configuration files may have same open source parameters. After the parameters in different files are set to different values, the configuration takes effect depends on the loading sequence of the configuration files by components. You can customize parameters for services and roles as required. Adding customized parameters for a single role instance is not supported.

Step 6 Locate the row where a specified parameter resides, enter the parameter name supported by the component in the **Name** column and enter the parameter value in the **Value** column.

You can click **+** or **-** to add or delete a customized parameter.

Step 7 Click **Save**. In the displayed **Save Configuration** dialog box, confirm the modification and click **OK**. After the system displays "Operation succeeded", click **Finish**. The configuration is saved successfully.

Restart the expired service or instance for the configuration to take effect.

----End

Task Example (Configuring Customized Hive Parameters)

Hive depends on HDFS. By default, Hive accesses the HDFS client. The configuration parameters that have taken effect are controlled by HDFS. For example, the HDFS parameter **ipc.client.rpc.timeout** affects the RPC timeout interval for all clients to connect to the HDFS server. Cluster administrators can modify the timeout interval for Hive to connect to HDFS by configuring custom parameters. After this parameter is added to the **core-site.xml** file of Hive, this

parameter can be identified by the Hive service and its configuration overwrites the parameter configuration in HDFS.

Step 1 On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and click **Services**.

Step 2 On the displayed page, click **Configuration** and click **All Configurations**.

Step 3 In the navigation tree on the left, select **Customization** for the Hive service. The system displays the custom service parameters supported by Hive.

Step 4 In **core-site.xml**, locate the row that contains the **core.site.customized.configs** parameter, enter **ipc.client.rpc.timeout** in the **Name** column, and enter a new value in the **Value** column, for example, 150000. The unit is ms.

Figure 8-21 Adding custom parameters

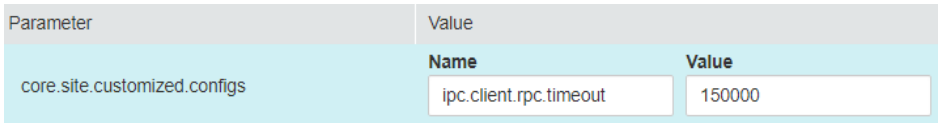

Step 5 Click **Save**. In the displayed **Save Configuration** dialog box, confirm the modification and click **OK**. Wait until the message "Operation succeeded" is displayed, and click **Finish**.

The configuration is saved successfully.

After the configuration is saved, restart the expired service or instance for the configuration to take effect.

----End

8.2.3 Instance Management

8.2.3.1 Overview

Overview

Log in to FusionInsight Manager and choose **Cluster** > **Services** > *Name of the service to be operated*. On the displayed page, click **Instance**. The instance management page is displayed, including the function area and role instance list.

Functional Area

After selecting the instances to be operated in the function area, you can maintain and manage the role instances, such as starting or stopping the instances. **Table 8-10** shows the main operations.

Table 8-10 Instance maintenance and management

UI Portal	Description
Start Instance	Start a specified instance in the cluster. You can start a role instance in the **Not Started**, **Stop Failed**, or **Startup Failed** state to use the role instance.
More > Stop Instance	Stop a specified instance in the cluster. You can stop a role instance that is no longer used or is abnormal.
More > Restart Instance	Restart a specified instance in the cluster. You can restart an abnormal role instance to restore it.
More > Instance Rolling Restart	Restart a specified instance in the cluster without interrupting services. For details about the parameter settings, see **Performing a Rolling Restart of a Cluster**.
More > Decommission/ Recommission	Recommission or decommission a specified instance in the cluster to change the service availability status of the service. For details, see **Decommissioning and Recommissioning an Instance**. **NOTE** Only the role DataNode in HDFS, role NodeManager in Yarn, and role RegionServer in HBase support the recommissioning and decommissioning functions.
Desired instance > **More** > **Synchronize Configuration**	If the **Configuration Status** of a role instance is **Expired**, the role instance has not been restarted after the configuration is modified, and the new configuration is saved only on FusionInsight Manager. In this case, use this function to deliver the new configuration to the specified instance. **NOTE** • After synchronizing the role instance configuration, you need to restart the role instance whose configuration has expired. The role instance is unavailable during the restart. • After the synchronization is complete, restart the instance for the configuration to take effect.
Desired instance > **Instance Configurations**	For details, see **Managing Instance Configurations**.

You can filter instances based on the role they belong to or their running status in the function area.

NOTE

Click **Advanced Search** to search for specified instances by specifying other filter criteria, such as **Host Name**, **Management IP Address**, **Business IP Address**, or **Instance Groups**.

Role Instance List

The role instance list contains the instances of all roles in the cluster. The list displays the running status, configuration status, hosts, and related IP addresses of each instance.

Table 8-11 Instance running status

Status	Description
Normal	Indicates that the instance is running properly.
Faulty	Indicates that the instance cannot run properly.
Decommissioned	Indicates that the instance is out of service.
Not started	Indicates that the instance is stopped.
Unknown	Indicates that the initial status of the instance cannot be detected.
Starting	Indicates that the instance is being started.
Stopping	Indicates that the instance is being stopped.
Restoring	Indicates that an exception may occur in the instance and the instance is being automatically rectified.
Decommissioning	Indicates that the instance is being decommissioned.
Recommissioning	Indicates that the instance is being recommissioned.
Failed to start	Indicates that the service fails to be started.
Failed to stop	Indicates that the service fails to be stopped.

Instance Details

You can click an instance name to go to the instance details page and view the basic information, configuration file, instance logs, and monitoring metric reports of the instance.

8.2.3.2 Decommissioning and Recommissioning an Instance

Scenario

Some role instances provide services for external services in distributed and parallel mode. Services independently store information about whether each instance can be used. Therefore, you need to use FusionInsight Manager to recommission or decommission these instances to change the instance running status.

Some instances do not support the recommissioning and decommissioning functions.

NOTE

The following roles support decommissioning and recommissioning: HDFS DataNode, YARN NodeManager, and HBase RegionServer.

- By default, if the number of the DataNodes is less than or equal to that of HDFS replicas, decommissioning cannot be performed. If the number of HDFS replicas is three and the number of DataNodes is less than four in the system, decommissioning cannot be performed. In this case, an error will be reported and force FusionInsight Manager to exit the decommissioning 30 minutes after FusionInsight Manager attempts to perform the decommissioning.
- You can enable quick decommissioning before decommissioning DataNodes for clusters of MRS 3.3.0 or later. In this case, when the number of DataNodes meets the value of **dfs.namenode.decommission.force.replication.min**, the system decommissions the nodes and adds HDFS copies at the same time. **If data is written during quick decommissioning, data may be lost. Exercise caution when performing this operation.** The parameters related to quick decommissioning are listed as follows. You can search for and view the parameters on the HDFS configuration page on FusionInsight Manager.

  dfs.namenode.decommission.force.enabled: whether to enable quick decommissioning for DataNode. If this parameter is set to **true**, the function is enabled.

  dfs.namenode.decommission.force.replication.min: minimum number of available copies of a block required for DataNode quick decommissioning. The value ranges from 1 to 3.
- During MapReduce task execution, files with 10 replicas are generated. Therefore, if the number of DataNode instances is less than 10, decommissioning cannot be performed.
- If the number of DataNode racks (the number of racks is determined by the number of racks configured for each DataNode) is greater than 1 before the decommissioning, and after some DataNodes are decommissioned, that of the remaining DataNodes changes to 1, the decommissioning will fail. Therefore, before decommissioning DataNode instances, you need to evaluate the impact of decommissioning on the number of racks to adjust the DataNodes to be decommissioned.
- If multiple DataNodes are decommissioned at the same time, and each of them stores a large volume of data, the DataNodes may fail to be decommissioned due to timeout. To avoid this problem, it is recommended that one DataNode be decommissioned each time and multiple decommissioning operations be performed.

Procedure

Step 1 Perform the following steps to perform a health check for the DataNodes before decommissioning:

1. Log in to the client installation node as a client user and switch to the client installation directory.
2. For a security cluster, use user **hdfs** for permission authentication.

   ```
   source bigdata_env          #Configure client environment variables.
   kinit hdfs                  #Configure kinit authentication.
   Password for hdfs@HADOOP.COM:    #Enter the login password of user hdfs.
   ```
3. Run the **hdfs fsck / -list-corruptfileblocks** command, and check the returned result.
   - If "has 0 CORRUPT files" is displayed, go to **Step 2**.
   - If the result does not contain "has 0 CORRUPT files" and the name of the damaged file is returned, go to **Step 1.4**.
4. Run the **hdfs dfs -rm** *Name of the damaged file* command to delete the damaged file.

NOTE

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.

Step 2 Log in to FusionInsight Manager.

Step 3 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 4 Click the specified service name on the service management page. On the displayed page, click the **Instance** tab.

Step 5 Select the specified role instance to be decommissioned.

Step 6 Select **Decommission** or **Recommission** from the **More** drop-down list.

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

Select **I confirm to decommission these instances and accept the consequence of service performance deterioration** and click **OK** to perform the corresponding operation.

NOTE

During the instance decommissioning, if the service corresponding to the instance is restarted in the cluster using another browser, FusionInsight Manager displays a message indicating that the instance decommissioning is stopped, but the operating status of the instance is displayed as **Started**. In this case, the instance has been decommissioned on the background. You need to decommission the instance again to synchronize the operating status.

----End

8.2.3.3 Managing Instance Configurations

Scenario

Configuration parameters of each role instance can be modified. In the scenario where instances are migrated to a new cluster or the service is redeployed, the cluster administrator can import or export all configuration data of a service on FusionInsight Manager to quickly copy configuration results.

FusionInsight Manager can manage configuration parameters of a single role instance. Modifying configuration parameters and importing or exporting instance configurations do not affect other instances.

Impact on the System

After modifying the configuration of a role instance, you need to restart the instance if the instance status is **Expired**. The role instance is unavailable during restart.

Modifying Instance Configuration

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 On the page that is displayed, click the **Instance** tab.

Step 4 Click the specified instance and select **Instance Configuration**.

By default, **Basic Configuration** is displayed. To modify more parameters, click **All Configurations**. All parameter categories supported by the instance are displayed on the **All Configurations** tab page.

Figure 8-22 Instance configurations

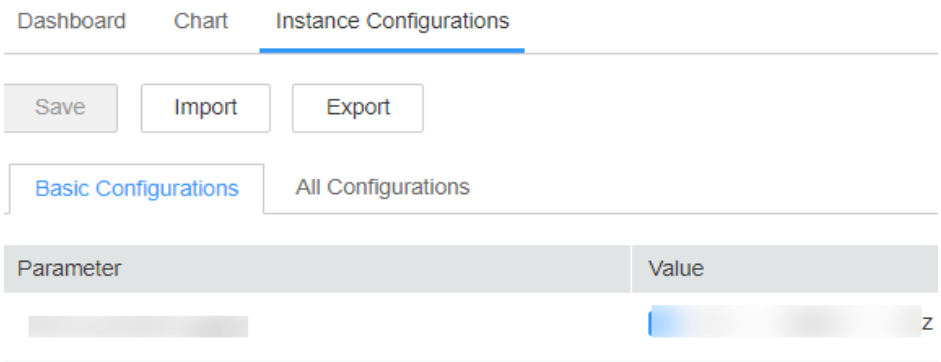

Step 5 In the navigation tree, select the specified parameter category and change the parameter values on the right.

If you are not sure about the location of a parameter, you can enter the parameter name in search box in the upper right corner. The system searches for the parameter in real time and displays the result.

Step 6 Click **Save**. In the confirmation dialog box, click **OK**.

Wait until the message "Operation succeeded." is displayed. Click **Finish**.

The configuration is modified.

NOTE

After the configuration parameters of a role instance are modified, you need to restart the instance if the instance status is **Expired**. You can select the expired instance on the **Instances** page and choose **More** > **Restart Instance**.

If the is displayed before a parameter, this parameter takes effect dynamically. After the configuration is saved, the parameter value is automatically updated to the configuration file. is supported only in MRS 3.2.0 or later.

----End

Exporting/Importing Instance Configuration

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page. On the displayed page, click the **Instance** tab.

Step 4 Click the specified instance and select **Instance Configurations**.

Step 5 Click **Export** to export the configuration parameter file to the local host.

Step 6 On the **Instance Configurations** page, click **Import**, select the configuration parameter file of the instance, and import the file.

----End

8.2.3.4 Viewing the Instance Configuration File

Scenario

FusionInsight Manager allows O&M personnel to view the content configuration files such as environment variables and role configurations of the instance node on the management page. If O&M personnel need to quickly check whether configuration items of the instance are incorrectly configured or when some hidden configuration items need to be viewed, the O&M personnel can directly view the configuration files on FusionInsight Manager. In this case, users quickly analyze configuration problems.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Service**.

Step 3 Click the specified service name on the service management page. On the displayed page, click the **Instance** tab.

Step 4 Click the name of the target instance. In the **Configuration File** area on the **Instance Status** page, the configuration file list of the instance is displayed.

Figure 8-23 Viewing the instance configuration file

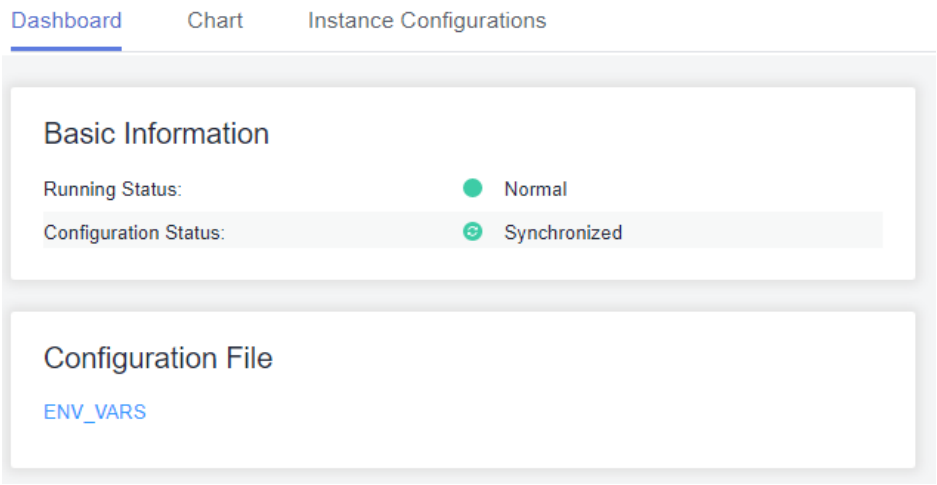

Step 5 Click the name of the configuration file to be viewed to view the parameter values in the configuration file.

To obtain the configuration file, you can download the configuration file to the local PC.

NOTE

If a node in the cluster is faulty, the configuration file cannot be viewed. Rectify the fault before viewing the configuration file again.

----End

8.2.3.5 Instance Group

8.2.3.5.1 Managing Instance Groups

Scenario

Instance groups can be managed on FusionInsight Manager. That is, you can group multiple instances in the same role based on a specified principle, such as the nodes with the same hardware configuration. The modification on the configuration parameters of an instance group applies to all instances in the group.

In a large cluster, instance groups are used to improve the capability of managing instances in batches in the heterogeneous environment. After instances are grouped, the instances can be configured repeatedly to reduce redundant instance configuration items and improve system performance.

Creating an Instance Group

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 On the displayed page, click the **Instance Groups** tab.

Click  and configure parameters as prompted.

Figure 8-24 Creating an Instance Group

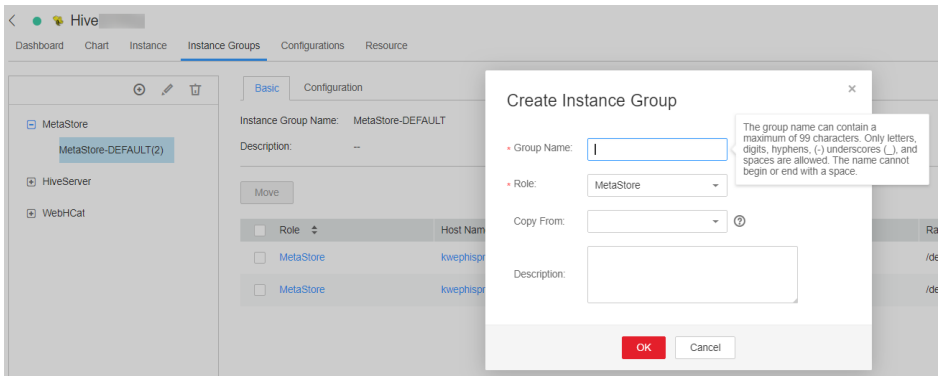

Table 8-12 Instance group configuration parameters

Parameter	Description
The group name	Indicates the instance group name. The value can contain only letters, digits, underscores (_), hyphens (-), and spaces. It must start with a letter, digit, underscore (_), or hyphen (-) and cannot ends with a space. It can contain a maximum of 99 characters.
Role	Indicates the role to which an instance group belongs.
Copy From	Indicates that the parameter values of a specified instance group are copied to the parameters of a new group. If the value is null, the default values are used for the parameters of the new group.
Description	Indicates the instance group description. It can contain only letters, digits, commas (,), periods (.), underscores (_), spaces, and line breaks, and can contain a maximum of 200 characters.

NOTE

- Each instance must belong to only one instance group. When an instance is installed for the first time, it belongs to the instance group *Role name*-**DEFAULT** by default.
- You can delete unnecessary or unused instance groups. Before deleting an instance group, migrate all instances in the group to other instance groups, and then delete the instance group by referring to **Deleting an Instance Group**. The default instance group cannot be deleted.

Step 5 Click **OK**.

The instance group is created.

----End

Modifying Properties of an Instance Group

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 Click the **Instance Groups** tab. On the **Instance Groups** tab page, locate the row that contains the target instance group.

Click and modify parameters as prompted.

Step 5 Click **OK** to save the modifications.

The default instance group cannot be modified.

----End

Deleting an Instance Group

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 Click the **Instance Groups** tab. On the **Instance Groups** tab page, locate the row that contains the target instance group.

Step 5 Click .

Step 6 In the displayed dialog box, click **OK**.

The default instance group cannot be deleted.

----End

8.2.3.5.2 Viewing Information About an Instance Group

Scenario

The cluster administrator can view the instance group of a specified service on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 On the displayed page, click the **Instance Groups** tab.

Step 5 In the navigation tree, select a role. On the **Basic** tab page, view all instances in the instance group.

NOTE

To move an instance from an instance group to another, perform the following operations:

1. Select the instance to be moved and click **Move**.
2. In the displayed dialog box, select an instance group to which the instance to be moved.

   During the migration, the configuration of the new instance group is automatically inherited. If the instance configuration is modified before the migration, the configuration of the instance prevails.
3. Click **OK**.

   Restart the expired service or instance for the configuration to take effect.

----End

8.2.3.5.3 Configuring Instantiation Group Parameters

Scenario

In a large cluster, users can configure parameters for multiple instances in batches by configuring the related instance groups on FusionInsight Manager, reducing redundant instance configuration items and improving system performance.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 3 Click the specified service name on the service management page.

Step 4 On the displayed page, click the **Instance Groups** tab.

Step 5 In the navigation tree, select the instance group name of a role, and switch to the **Configuration** tab page. Adjust parameters to be modified, and click **Save**. The configuration takes effect for all instances in the instance group.

----End

8.3 Hosts

8.3.1 Host Management Page

8.3.1.1 Viewing the Host List

Overview

Log in to FusionInsight Manager, click **Hosts**, and the host list is displayed on the host management page. You can view the host list and basic information of each host.

You can switch view types and set search criteria to filter and search for hosts.

Host View

You can click **Role View** to view the roles deployed on each host. If the role supports the active/standby mode, the role name is displayed in bold.

Host List

The host list on the host management page contains all hosts in the cluster, and O&M operations can be performed on these hosts.

On the host management page, you can filter hosts by node type or cluster. The rules for filtering host types are as follows:

- A management node is the node where OMS is deployed. Additionally, control roles and data roles may also be deployed on management nodes.

- A control node is the node where control roles are deployed. Additionally, data roles may also be deployed on control nodes.
- A Data Node is the node where only data roles are deployed.

If you select the **Host View**, the IP address, rack planning, AZ name, running status, cluster name, and hardware resource usage of each host are displayed.

Table 8-13 Host running status

Status	Description
Normal	Indicates that the host is in the normal state.
Faulty	Indicates that the host is abnormal.
Unknown	Indicates that the initial status of the host cannot be detected.
Isolated	Indicates that the host is isolated.
Suspended	Indicates that the host is stopped.

8.3.1.2 Viewing the Host Dashboard

Overview

Log in to FusionInsight Manager, click **Hosts**, and click a host name in the host list. The host details page contains the basic information area, disk status area, role list area, and monitoring chart.

Basic Information Area

The basic information area contains the key information about the host, such as the management IP address, service IP address, host type, rack, firewall, number of CPU cores, and OS.

Disk Status Area

The disk status area contains all disk partitions configured for the cluster on the host and the usage of each disk partition.

Instance List Area

The instance list area displays all role instances installed on the host and the status of each role instance. You can click the log file next to a role instance name to view the log file content of the instance online.

Alarm and Event History

The alarm and event history area displays the key alarms and events reported by the current host. The system can display a maximum of 20 historical records.

Chart

The monitoring chart area is displayed on the right of the host details page, and contains the key monitoring metrics of the host.

You can choose ▼ > **Customize** in the upper right corner to customize the monitoring reports to be displayed in the chart area. Select a time range and choose ▼ > **Export** to export detailed monitoring metric data within the specified time range.

You can click  next to the title of a monitoring indicator to open the description of the monitoring indicator.

Click the **Chart** tab of the host to view the full monitoring chart information about the host.

8.3.1.3 Checking Host Processes and Resources

Overview

Log in to FusionInsight Manager, click **Hosts**, and click the specified host name in the host list. On the host details page, click the **Process** and **Resource** tabs.

Host Process

On the **Process** tab page, the information about the role processes of the deployed service instances on the current host is displayed, including the process status, PID, and process running time. You can directly view the log files of each process online.

Host Resource

On the **Resource** tab page, the detailed resource usage of deployed service instances on the current host is displayed, including the CPU, memory, disk, and port usage.

8.3.2 Host Maintenance Operations

8.3.2.1 Starting and Stopping All Instances on a Host

Scenario

If a host is faulty, you may need to stop all the roles on the host and perform maintenance check on the host. After the host fault is rectified, start all roles running on the host to recover host services. You can start or stop all instances on a host on the host management page or host details page on FusionInsight Manager. The following describes how to perform such operations on the host management page.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Hosts**.

Step 3 Select the check box of the target host.

Step 4 Select **Start All Instances** or **Stop All Instances** from the **More** drop-down list to start or stop all role instances.

----End

8.3.2.2 Performing a Host Health Check

Scenario

If the running status of a host is not **Normal**, you can perform health checks on the host to check whether some basic functions are abnormal. During routine O&M, you can perform host health checks to ensure that the configuration parameters and monitoring of each role instance on the host are normal and can run stably for a long time.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Hosts**.

Step 3 Select the check box of the target host.

Step 4 Select **Health Check** from the **More** drop-down list to start the health check.

To export the result of the health check, click **Export Report** in the upper left corner. If any problem is detected, click **Help**.

----End

8.3.2.3 Configuring Racks for Hosts

Scenario

All hosts in a large cluster are usually deployed on multiple racks. Hosts on different racks communicate with each other through switches. The network bandwidth between different hosts on the same rack is much greater than that on different racks. In this case, plan the network topology based on the following requirements:

- To improve the communication speed, it is recommended that data be exchanged between hosts on the same rack.
- To improve the fault tolerance capability, distribute processes or data of distributed services on different hosts of multiple racks as dispersedly as possible.

Hadoop uses a file directory structure to represent hosts.

The HDFS cannot automatically determine the network topology of each DataNode in the cluster. You need to set the rack name to identify the rack where the host is located so that the NameNode can draw the network topology of the required DataNodes and back up data of the DataNodes to different racks. Similarly, YARN needs to obtain rack information and allocate tasks to different NodeManagers as required.

If the cluster network topology changes, you need to reallocate racks for hosts on FusionInsight Manager so that related services can be automatically adjusted.

Impact on the System

If the name of the host rack is changed, storage policy for HDFS replicas, YARN task assignment, and storage location of Kafka partitions will be affected. After the modification, you need to restart the HDFS, YARN, and Kafka for the configuration to take effect.

Improper rack configuration will unbalance loads (including CPU, memory, disk, and network) among nodes in the cluster, which decreases the cluster reliability and stability. Therefore, before allocating racks, take all aspects into consideration and properly set racks.

Rack Allocation Policies

NOTE

Physical rack: indicates the real rack where the host resides.

Logical rack: indicates the rack name of the host on FusionInsight Manager.

Policy 1: Each logical rack has nearly the same number of hosts.

Policy 2: The name of the logical rack of the host must comply with that of the physical rack to which the host belongs.

Policy 3: If there are only few hosts on a physical rack, combine this physical rack and other physical racks with few hosts into a logical rack, which complies with policy 1. Hosts in two equipment rooms cannot be placed in one logical rack. Otherwise, performance problems may be caused.

Policy 4: If there are lots of hosts on a physical rack, divide these hosts into multiple logical racks, which complies with policy 1. Hosts with great differences should not be placed in the same logical rack. Otherwise, the cluster reliability will be decreased.

Policy 5: You are advised to set **default** or other values for logical racks on the first layer, and the values in the same cluster must be consistent.

Policy 6: The number of hosts in each rack cannot be less than 3.

Policy 7: A cluster can contain at most 50 logical racks. If there are too many logical racks in a cluster, the maintenance is difficult.

Best Practices

For example, in a cluster, 100 hosts are located in two equipment rooms A and B. A has 40 hosts and B has 60 hosts. In room A, there are 11 hosts on physical rack Ra1 and 29 hosts on physical rack Ra2. In room B, there are six hosts on physical

rack Rb1, 33 hosts on physical rack Rb2, 18 hosts on physical rack Rb3, and three hosts on physical rack Rb4.

According to the rack allocation policy, each logical rack contains nearly the same number (for example, 20) of hosts. The allocation details are as follows:

- Logical rack /default/racka1: 11 hosts on physical rack Ra1 and nine hosts on physical rack Ra2
- Logical rack /default/racka2: the remaining 20 hosts (except the nine hosts of logical rack /default/racka1) on physical rack Ra2
- Logical rack /default/rackb1: six hosts on physical rack Rb1 and 13 hosts on physical rack Rb2
- Logical rack /default/rackb2: the remaining 20 hosts on physical rack Rb2
- Logical rack /default/rackb3: 18 hosts on physical rack Rb3 and three hosts on physical rack Rb4

Rack allocation example:

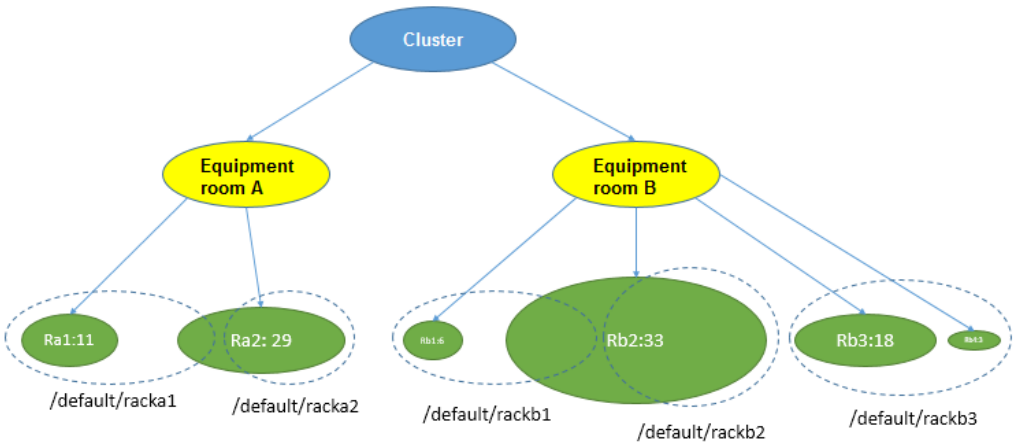

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Hosts**.

Step 3 Select the check box of the target host.

Step 4 Select **Set Rack** from the **More** drop-down list.

- Set rack names in hierarchy based on the actual network topology. Separate racks from different layers using slashes (/).
- Rack naming rules are as follows: */level1/level2/...* The number of levels must be at least 1, and the name cannot be empty. A rack can contain letters, digits, and underscores (_) and cannot exceed 200 characters.

  For example, /default/rack0.
- If the hosts in the rack to be modified contain DataNode instances, ensure that the rack name levels of the hosts where all DataNode instances reside are the same. Otherwise, the configuration fails to be delivered.

Step 5 Click **OK**.

----End

8.3.2.4 Isolating a Host

Scenario

If a host is abnormal or faulty and cannot provide services or affects the cluster performance, you can remove the host from the available node in the cluster temporarily so that the client can access other available nodes.

NOTE

Only non-management nodes can be isolated.

Impact on the System

- After a host is isolated, all role instances on the host will be stopped, and you cannot start, stop, or configure the host and all instances on the host.
- For some services, after a host is isolated, some instances on other nodes do not work, and the service configuration status may expire.
- After a host is isolated, statistics about the monitoring status and indicator data of the host hardware and instances on the host cannot be collected or displayed.
- Retain the default SSH port (22) of the target node. Otherwise, the task described in this section will fail.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Hosts**.

Step 3 Select the check box of the host to be isolated.

Step 4 Select **Isolate** from the **More** drop-down list.

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

Step 5 In the displayed confirmation dialog box, select "I confirm to isolate the selected hosts and accept possible consequences of service faults." Click **OK**.

Wait until the message "Operation succeeded" is displayed, and click **Finish**.

The host is successfully isolated and **Running Status** is **Isolated**.

Step 6 Log in to the isolated host as user **root** and run the **pkill -9 -u omm** command to stop the processes of user **omm** on the node. Then run the **ps -ef | grep 'container' | grep '${BIGDATA_HOME}' | awk '{print $2}' | xargs -I '{}' kill -9 '{}'** command to find and stop the container process.

Step 7 Cancel the isolation status of the host before using the host if you have rectified the host exception or fault.

On the **Hosts** page, select the isolated host and choose **More** > **Cancel Isolation**.

NOTE

After the isolation is canceled, all role instances on the host are not started by default. To start role instances on the host, select the target host on the Hosts page and choose **More** > **Start All Instances**.

----End

8.3.2.5 Exporting Host Information

Scenario

Administrators can export information about all hosts on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Hosts**.

Step 3 Specify the status of required hosts in the drop-down list box on the upper right corner, or click **Advanced Search** to specify hosts.

Step 4 Click **Export All**, select **TXT** or **CSV** for **Save As**, and click **OK**.

----End

8.3.3 Resource Overview

8.3.3.1 Distribution

Log in to FusionInsight Manager and choose **Hosts** > **Resource Overview**. On the **Resource Overview** page that is displayed, click the **Distribution** tab to view resource distribution of each cluster. By default, the monitoring data of the past one hour (**1h**) is displayed. You can click to customize a time range. Time range options are **1h**, **2h**, **6h**, **12h**, **1d**, **1w**, and **1m**.

Figure 8-25 Distribution tab

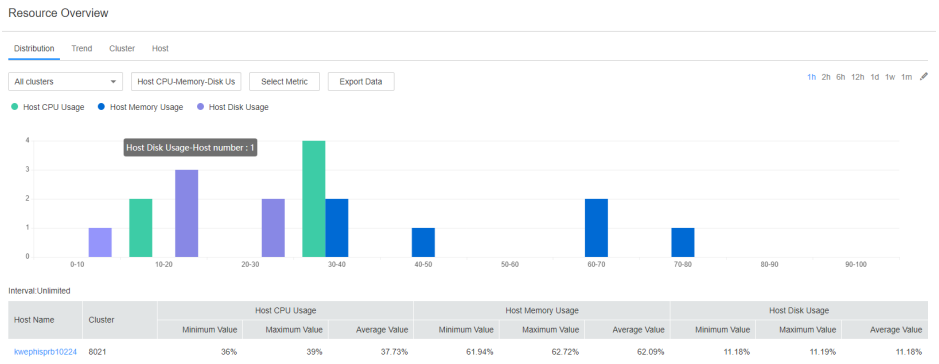

- You can click **Select Metric** to customize the metric to monitor. **Table 8-14** describes all the metrics that you can select. After you select a metric, the host distribution in each range of the metric is displayed.

- When you hover your cursor over a color column, the number of hosts in the current metric range is displayed. See **Figure 8-25**. You can click a color column to view the list of hosts in the metric range.
  - You can click a host name in the **Host Name** column to access the host details page.
  - You can click **View Trends** in the **Operation** column of a host to view the maximum, minimum, and average values of the current metric in the cluster as well as the value of the current host. In the current cluster, if you have selected **Host CPU-Memory-Disk Usage**, **View Trends** is unavailable.
- You can click **Export Data** to export the maximum, minimum, and average values of the current metric of all nodes in the cluster within the time range you have specified.

Table 8-14 Metrics

Category	Metric
Process	• Number of Running Processes • Total Number of Processes • Total Number of omm Processes • Uninterruptible Sleep Process
Network Status	• Host Network Packet Collisions • Number of LAST_ACK States • Number of CLOSING States • Number of LISTENING States • Number of CLOSED States • Number of ESTABLISHED States • Number of SYN_RECV States • Number of TIME_WAITING States • Number of FIN_WAIT2 States • Number of FIN_WAIT1 States • Number of CLOSE_WAIT States • DNS Name Resolution Duration • TCP Ephemeral Port Usage • Host Network Packet Frame Errors
Network Reading	• Host Network Read Packets • Host Network Read Dropped Packets • Host Network Read Error Packets • Host Network Rx Speed

Category	Metric
Disk	• Host Disk Write Speed • Host Used Disk • Host Free Disk • Host Disk Read Speed • Host Disk Usage
Memory	• Free Memory • Cache Memory Size • Total Kernel Cache Memory Size • Shared Memory Size • Host Memory Usage • Used Memory
Network Writing	• Host Network Write Packets • Host Network Write Error Packets • Host Network Tx Speed • Host Network Write Dropped Packets
CPU	• CPU Usage of Processes Whose Priorities Have Been Changed • CPU Usage of User Space Processes • CPU Usage of Kernel Space Processes • Host CPU Usage • CPU Total Time • CPU Idle Time
Host Status	• Host File Handle Usage • Average OS Load in 1 Minute • Average OS Load in 5 Minutes • Average OS Load in 15 Minutes • Host PID Usage

8.3.3.2 Trend

Log in to FusionInsight and choose **Hosts** > **Resource Overview**. On the **Resource Overview** page that is displayed, click the **Trend** tab to view resource trends of all clusters or a single cluster. By default, the monitoring data of the past one hour (**1h**) is displayed. You can click to customize a time range. Time range options are **1h**, **2h**, **6h**, **12h**, **1d**, **1w**, and **1m**. By default, the trend chart of each metric displays the maximum, minimum, and average values of the entire cluster.

Figure 8-26 Trend tab

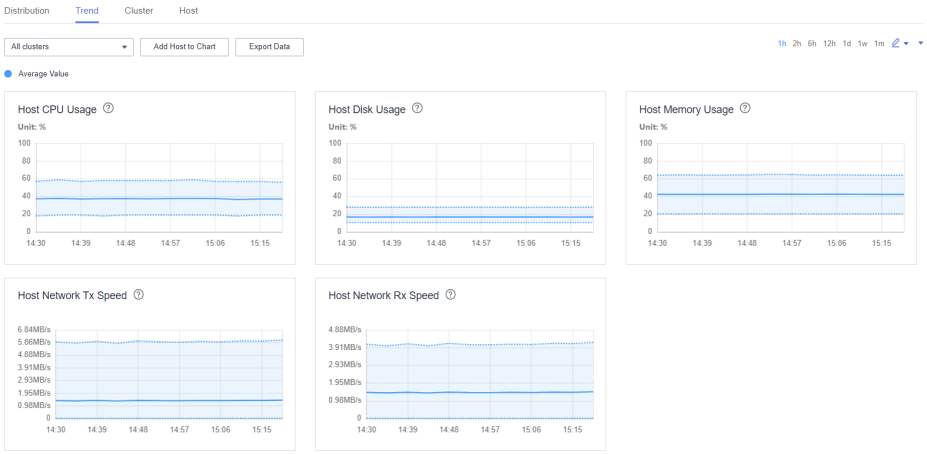

- You can click **Add Host to Chart** to add trend lines of up to 12 hosts to the trend charts.
- You can choose ▼ > **Customize** to customize the metrics to display on the tab page. For details about the metrics, see **Table 8-14** in **Distribution**.
- You can click **Export Data** to export the maximum, minimum, and average values of all nodes in the cluster for all selected metrics within the time range you have specified.

8.3.3.3 Cluster

Log in to FusionInsight Manager and choose **Hosts** > **Resource Overview**. On the **Resource Overview** page that is displayed, click the **Cluster** tab to view resource monitoring of all clusters.

By default, the monitoring data of the past one hour (**1h**) is displayed. You can click ✎ to customize a time range. Time range options are **1h**, **2h**, **6h**, **12h**, **1d**, **1w**, and **1m**.

Figure 8-27 Cluster tab

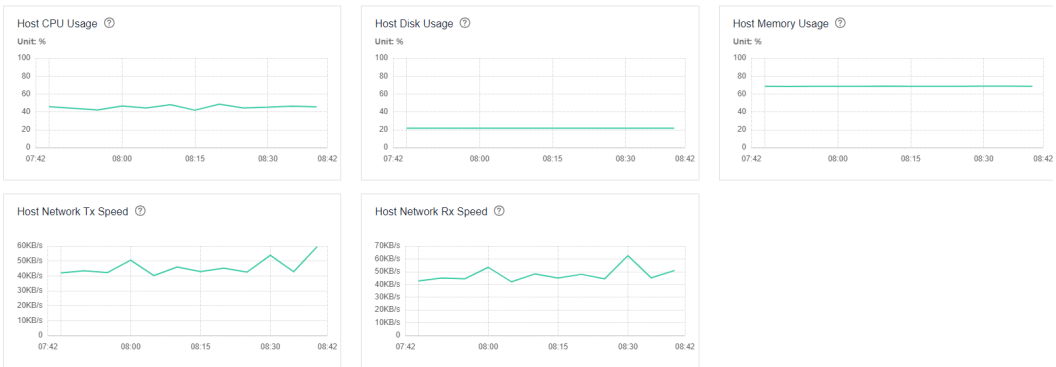

- You can choose ▼ > **Customize** to customize the metrics to display on the tab page. For details about the metrics, see **Table 8-14** in **Distribution**.
- You can click **Export Data** to export the metric values of each cluster within the time range you have specified.

8.3.3.4 Host

Log in to FusionInsight Manager and choose **Hosts** > **Resource Overview**. On the **Resource Overview** page that is displayed, click the **Host** tab to view host resource overview, including basic configurations (CPU/memory) and disk configurations.

You can click **Export Data** to export the configuration list of all hosts in the cluster, including the host name, management IP address, host type, number of cores, CPU architecture, memory capacity, and disk size.

Figure 8-28 Host tab

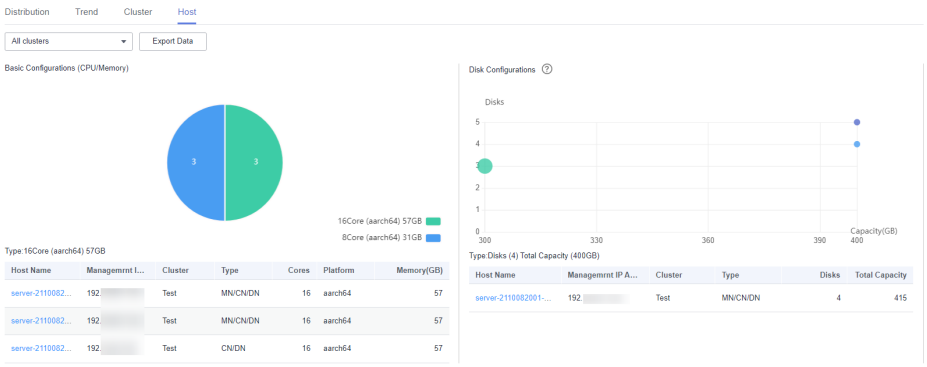

Basic Configurations (CPU/Memory)

You can hover your cursor over the pie chart to view the number of hosts of each hardware configuration in the cluster. The information is displayed in the format of *Number of cores (CPU architecture) Memory size*.

You can click a slice on the pie chart to view the list of hosts.

Disk Configurations

The horizontal axis indicates the total disk capacity (including the OS disk) of a node, and the vertical axis indicates the number of logical disks (including the OS disk).

You can hover your cursor over a dot to view information about disks of the current configuration, including the quantity of disks, total capacity, and number of hosts.

You can click a dot on the chart to view the list of hosts.

8.4 O&M

8.4.1 Alarms

8.4.1.1 Overview of Alarms and Events

Alarms

Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. You can view information about alarms reported by all clusters on the page shown in **Figure 8-29**, including the alarm name, ID, severity, and generation time. By default, the latest 10 alarms are displayed on each page.

Figure 8-29 Alarms

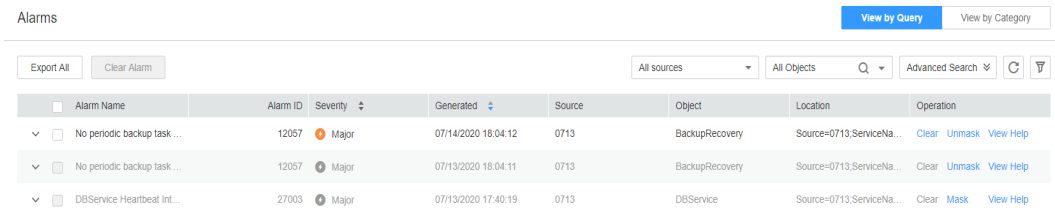

You can click ⌄ on the left of an alarm to view detailed alarm parameters. **Table 8-15** describes the parameters.

Table 8-15 Alarm parameters

Parameter	Description
Alarm ID	Alarm ID
Alarm Name	Alarm name
Severity	Alarm severity. Value options are **Critical**, **Major**, **Minor**, and **Suggestion**.
Generated	Time when an alarm is generated
Cleared	Time when an alarm is cleared. If the alarm is not cleared, -- is displayed.
Source	Cluster name
Object	Service, process, or module that triggers the alarm
Automatically Cleared	Whether the alarm can be automatically cleared after the fault is rectified
Alarm Status	Current status of the alarm. Value options are **Auto**, **Manual**, and **Uncleared**.
Alarm Cause	Indicates the possible cause of an alarm.
Serial Number	Indicates the number of alarms generated by the system.

Parameter	Description
Additional Information	Indicates the error information. MRS 3.3.0 and later versions: You can view the monitoring metric values in **Additional Information** if thresholds are set for the metrics to generate alarms.
Location	Detailed information for locating the alarm, which includes the following: • **Source**: cluster for which the alarm is generated • **ServiceName**: service for which the alarm is generated • **RoleName**: role for which the alarm is generated • **HostName**: host for which the alarm is generated

Manage alarms.

- Click **Export All** to export all alarm details.
- If multiple alarms have been handled, you can select one or more alarms to be cleared and click **Clear Alarm** to clear the alarms in batches. A maximum of 300 alarms can be cleared in each batch.
- You can click to manually refresh the current page and click to filter columns to display.
- You can filter alarms by object or severity.
- You can click **Advanced Search** to search for alarms by alarm ID, name, type, start time, or end time. Click **Search** to filter alarms that meet the search criteria. Click **Advanced Search** again to view the number of search criteria that you have configured.
- You can click **Clear**, **Mask**, or **View Help** to perform corresponding operations on an alarm.
- If there are a large number of alarms, you can click **View by Category** to sort uncleared alarms by alarm ID. After alarms are classified, click the number of uncleared alarms to view alarm details.

Events

Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Events**. On the **Events** page that is displayed, you can view information about all events in the cluster, including the event name, ID, severity, generation time, object, and location. By default, the latest 10 events are displayed on each page.

Figure 8-30 Events page

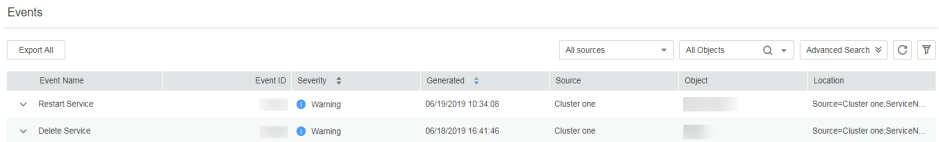

You can click ⌄ on the left of an event to view detailed event parameters. **Table 8-16** describes the parameters.

Table 8-16 Event parameters

Parameter	Description
Event ID	Event ID
Event Name	Event name
Severity	Event severity. Value options are **Critical**, **Major**, **Minor**, and **Suggestion**.
Generated	Time when an event is generated
Object	Object for which the event may be generated
Serial Number	Number of the event generated by the system
Location	Detailed information for locating the event, which includes the following: • **Source**: cluster for which the event is generated • **ServiceName**: service for which the event is generated • **RoleName**: role for which the event is generated • **HostName**: host for which the event is generated
Additional Information	Indicates the error information.
Event Cause	Indicates the possible cause of an event.
Source	Cluster name

Manage events.

- Click **Export All** to export all event details.
- You can click ↻ to manually refresh the current page and click ▽ to filter columns to display.
- You can filter events by object or cluster.
- You can click **Advanced Search** to search for events by event ID, name, severity, start time, or end time.

8.4.1.2 Alarm Threshold

Scenario

You can configure monitoring indicator thresholds to monitor the health status of indicators on FusionInsight Manager. If abnormal data occurs and the preset conditions are met, the system triggers an alarm and displays the alarm information on the alarm page.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Alarm** > **Thresholds**.

Step 3 Select a monitoring metric for a host or service in the cluster.

Figure 8-31 Configuring the threshold for a metric

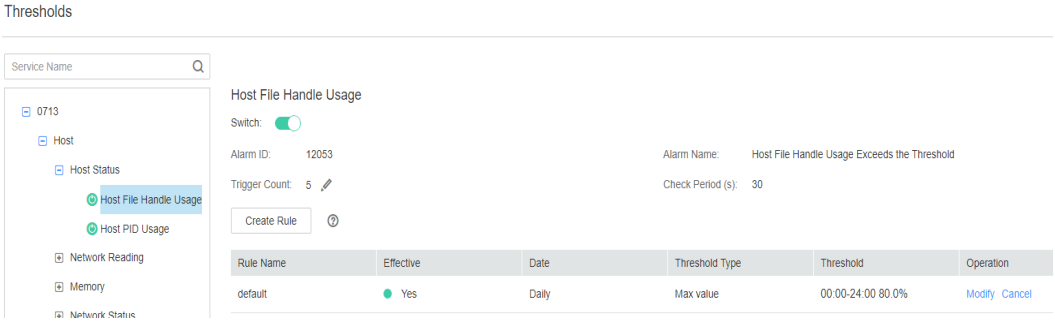

For example, after selecting **Host Memory Usage**, the information about this indicator threshold is displayed.

- If the alarm sending switch is turned on, an alarm will be triggered if the threshold is reached.
- When **Alarm Severity** is on, hierarchical alarms are enabled. The system dynamically reports alarms at each severity based on the real-time metric values and hierarchical thresholds set for the severity. MRS 3.3.0 or later supports this function.
- **Alarm ID** and **Alarm Name**: alarm information triggered against the threshold
- **Trigger Count**: FusionInsight Manager checks whether the value of a monitoring metric reaches the threshold. If the number of consecutive checks reaches the value of **Trigger Count**, an alarm is generated. **Trigger Count** is configurable.
- **Check Period (s)**: interval for the system to check the monitoring metric.
- The rules in the rule list are used to trigger alarms.

Step 4 Click **Create Rule** to add rules used for monitoring indicators.

Table 8-17 Monitoring indicator rule parameters

Parameter	Description	Example Value
Rule Name	Name of a rule.	CPU_MAX
Severity	Select an alarm severity. After **Alarm Severity** is on, you need to configure the alarm severity in **Thresholds**.	• Critical • Major • Minor • Warning

Parameter	Description	Example Value
Threshold Type	You can use the maximum or minimum value of an indicator as the alarm triggering threshold. If **Threshold Type** is set to **Max value**, the system generates an alarm when the value of the specified indicator is greater than the threshold. If **Threshold Type** is set to **Min value**, the system generates an alarm when the value of the specified indicator is less than the threshold.	• Max value • Min value
Date	This parameter is used to set the date when the rule takes effect. If **Alarm Severity** is on, only **Daily** is supported.	• Daily • Weekly • Others
Add Date	This parameter is available only when **Date** is set to **Others**. You can set the date when the rule takes effect. Multiple options are available.	09-30
Thresholds	This parameter is used to set the time range when the rule takes effect. If **Alarm Severity** is on, you cannot set the start time and end time. The default start time and end time are 00:00-23:59.	Start and End Time: 00:00–08:30
	Threshold of the rule monitoring metric After **Alarm Severity** is on, different alarm severities can be set for a cluster based on different thresholds.	Threshold: 10

NOTE

You can click  to set multiple time ranges for the threshold or click  to delete one.

Step 5 Click **OK** to save the rules.

Step 6 Locate the row that contains an added rule, and click **Apply** in the **Operation** column. The value of **Effective** for this rule changes to **Yes**.

A new rule can be applied only after you click **Cancel** for an existing rule.

----End

Monitoring Metric Reference

FusionInsight Manager alarm monitoring metrics are classified as node information metrics and cluster service metrics. **Table 8-18** lists the metrics whose thresholds can be configured for a node, and **Table 8-19** lists metrics whose thresholds can be configured for a component.

NOTE

On FusionInsight Manager of MRS 3.3.0 or later, alarms of some components can be reported by severity. Each alarm severity has a threshold. You can view them on the FusionInsight Manager configuration page.

Table 8-18 Node monitoring metrics

Metric Group	Metric	Description	Default Threshold
CPU	Host CPU Usage	This indicator reflects the computing and control capabilities of the current cluster in a measurement period. By observing the indicator value, you can better understand the overall resource usage of the cluster.	90.0%
Disk	Disk Usage	Indicates the disk usage of a host.	90.0%

Metric Group	Metric	Description	Default Threshold
	Disk Inode Usage	Indicates the disk inode usage in a measurement period.	80.0%
Memory	Host Memory Usage	Indicates the average memory usage at the current time.	90.0%
Host Status	Host File Handle Usage	Indicates the usage of file handles of the host in a measurement period.	80.0%
	Host PID Usage	Indicates the PID usage of a host.	90%
Network Status	TCP Ephemeral Port Usage	Indicates the usage of temporary TCP ports of the host in a measurement period.	80.0%
Network Reading	Read Packet Error Rate	Indicates the read packet error rate of the network interface on the host in a measurement period.	0.5%
	Read Packet Dropped Rate	Indicates the read packet dropped rate of the network interface on the host in a measurement period.	0.5%
	Read Throughput Rate	Indicates the average read throughput (at MAC layer) of the network interface in a measurement period.	80%

Metric Group	Metric	Description	Default Threshold
Network Writing	Write Packet Error Rate	Indicates the write packet error rate of the network interface on the host in a measurement period.	0.5%
	Write Packet Dropped Rate	Indicates the write packet dropped rate of the network interface on the host in a measurement period.	0.5%
	Write Throughput Rate	Indicates the average write throughput (at MAC layer) of the network interface in a measurement period.	80%
Process	Uninterruptible Sleep Process	Number of D state processes on the host in a measurement period	0
	omm Process Usage	omm process usage in a measurement period	90

Table 8-19 Cluster service indicators

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
DBService	Database	Usage of the Number of Database Connections	Usage of database connections	90%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		Disk Space Usage of the Data Directory	Disk space usage of the data directory	80%
Flume	Agent	Heap Memory Usage Calculate	Flume heap memory usage	95.0%
		Flume Direct Memory Usage Statistics	Flume direct memory usage	80.0%
		Flume Non-heap Memory Usage	Flume non-heap memory usage	80.0%
		Total GC duration of Flume process	Flume total GC time	12000 ms
HBase	GC	GC time for old generation	Total GC time of RegionServer	5000 ms
		GC time for old generation	Total GC time of HMaster	5000 ms
	CPU & memory	RegionServer Direct Memory Usage Statistics	RegionServer direct memory usage	90%
		RegionServer Heap Memory Usage Statistics	RegionServer heap memory usage	90%
		HMaster Direct Memory Usage	HMaster direct memory usage	90%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		HMaster Heap Memory Usage Statistics	HMaster heap memory usage	90%
	Service	Number of Online Regions of a RegionServer	Number of regions of a RegionServer	2000
		Region in transaction count over threshold	Number of regions that are in the RIT state and reach the threshold duration	1
	Replication	Replication sync failed times (RegionServer)	Number of times that DR data fails to be synchronized	1
		Number of Log Files to Be Synchronized in the Active Cluster	Number of log files to be synchronized in the active cluster	128
		Number of HFiles to Be Synchronized in the Active Cluster	Number of HFiles to be synchronized in the active cluster	128
	Queue	Compaction Queue Size	Size of the Compaction queue	100
HDFS	File and Block	Lost Blocks	Number of block copies that the HDFS lacks of	0

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		Blocks Under Replicated	Total number of blocks that need to be replicated by the NameNode	1000
	RPC	Average Time of Active NameNode RPC Processing	Average NameNode RPC processing time	100 ms
		Average Time of Active NameNode RPC Queuing	Average NameNode RPC queuing time	200 ms
	Disk	HDFS Disk Usage	HDFS disk usage	80%
		DataNode Disk Usage	Disk usage of DataNodes in the HDFS	80%
		Percentage of Reserved Space for Replicas of Unused Space	Percentage of the reserved disk space of all the copies to the total unused disk space of DataNodes.	90%
	Resource	Faulty DataNodes	Indicates the number of faulty DataNodes.	3
		NameNode Non-Heap Memory Usage Statistics	Indicates the percentage of NameNode non-heap memory usage.	90%
		NameNode Direct Memory Usage Statistics	Indicates the percentage of direct memory used by NameNodes.	90%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		NameNode Heap Memory Usage Statistics	Indicates the percentage of NameNode non-heap memory usage.	95%
		DataNode Direct Memory Usage Statistics	Indicates the percentage of direct memory used by DataNodes.	90%
		DataNode Heap Memory Usage Statistics	DataNode heap memory usage	95%
		DataNode Heap Memory Usage Statistics	Indicates the percentage of DataNode non-heap memory usage.	90%
	Garbage Collection	GC Time (NameNode)/ GC Time (DataNode)	Indicates the Garbage collection (GC) duration of NameNodes per minute.	12000 ms
		GC Time	Indicates the GC duration of DataNodes per minute.	12000 ms
Hive	HQL	Percentage of HQL Statements That Are Executed Successfully by Hive	Indicates the percentage of HQL statements that are executed successfully by Hive.	90.0%
	Background	Background Thread Usage	Background thread usage	90%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
	GC	Total GC time of MetaStore	Indicates the total GC time of MetaStore.	12000 ms
		Total GC Time in Milliseconds	Indicates the total GC time of HiveServer.	12000 ms
	Capacity	Percentage of HDFS Space Used by Hive to the Available Space	Indicates the percentage of HDFS space used by Hive to the available space.	85.0%
	CPU & memory	MetaStore Direct Memory Usage Statistics	MetaStore direct memory usage	95%
		MetaStore Non-Heap Memory Usage Statistics	MetaStore non-heap memory usage	95%
		MetaStore Heap Memory Usage Statistics	MetaStore heap memory usage	95%
		HiveServer Direct Memory Usage Statistics	HiveServer direct memory usage	95%
		HiveServer Non-Heap Memory Usage Statistics	HiveServer non-heap memory usage	95%
		HiveServer Heap Memory Usage Statistics	HiveServer heap memory usage	95%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
	Session	Percentage of Sessions Connected to the HiveServer to Maximum Number of Sessions Allowed by the HiveServer	Indicates the percentage of the number of sessions connected to the HiveServer to the maximum number of sessions allowed by the HiveServer.	90.0%
Kafka	Partition	Percentage of Partitions That Are Not Completely Synchronized	Indicates the percentage of partitions that are not completely synchronized to total partitions.	50%
	Others	Unavailable Partition Percentage	Percentage of unavailable partitions of each Kafka topic	40%
		User Connection Usage on Broker	Usage of user connections on Broker	80%
	Disk	Broker Disk Usage	Indicates the disk usage of the disk where the Broker data directory is located.	80.0%
		Disk I/O Rate of a Broker	I/O usage of the disk where the Broker data directory is located	80%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
	Process	Broker GC Duration per Minute	Indicates the GC duration of the Broker process per minute.	12000 ms
		Heap Memory Usage of Kafka	Indicates the Kafka heap memory usage.	95%
		Kafka Direct Memory Usage	Indicates the Kafka direct memory usage.	95%
Loader	Memory	Heap Memory Usage Calculate	Indicates the Loader heap memory usage.	95%
		Direct Memory Usage of Loader	Indicates the Loader direct memory usage.	80.0%
		Non-heap Memory Usage of Loader	Indicates the Loader non-heap memory usage.	80%
	GC	Total GC time of Loader	Indicates the total GC time of Loader.	12000 ms
MapReduce	Garbage Collection	GC Time	Indicates the GC time.	12000 ms
	Resource	JobHistoryServer Direct Memory Usage Statistics	Indicates the JobHistoryServer direct memory usage.	90%
		JobHistoryServer Non-Heap Memory Usage Statistics	Indicates the JobHistoryServer non-heap memory usage.	90%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		JobHistoryServer Heap Memory Usage Statistics	Indicates the JobHistoryServer non-heap memory usage.	95%
Oozie	Memory	Oozie Heap Memory Usage Calculate	Indicates the Oozie heap memory usage.	95.0%
		Oozie Direct Memory Usage	Indicates the Oozie direct memory usage.	80.0%
		Oozie Non-heap Memory Usage	Indicates the Oozie non-heap memory usage.	80%
	GC	Total GC duration of Oozie	Indicates the Oozie total GC time.	12000 ms
Spark/ Spark2x	Memory	JDBCServer2x Heap Memory Usage Statistics	JDBCServer2x heap memory usage	95%
		JDBCServer2x Direct Memory Usage Statistics	JDBCServer2x direct memory usage	95%
		JDBCServer2x Non-Heap Memory Usage Statistics	JDBCServer2x non-heap memory usage	95%
		JobHistory2x Direct Memory Usage Statistics	JobHistory2x direct memory usage	95%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		JobHistory2x Non-Heap Memory Usage Statistics	JobHistory2x non-heap memory usage	95%
		JobHistory2x Heap Memory Usage Statistics	JobHistory2x heap memory usage	95%
		IndexServer2x Direct Memory Usage Statistics	IndexServer2x direct memory usage	95%
		IndexServer2x Heap Memory Usage Statistics	IndexServer2x heap memory usage	95%
		IndexServer2x Non-Heap Memory Usage Statistics	IndexServer2x non-heap memory usage	95%
	GC Count	Full GC Number of JDBCServer2x	Full GC times of JDBCServer2x	12
		Full GC Number of JobHistory2x	Full GC times of JobHistory2x	12
		Full GC Number of IndexServer2x	Full GC times of IndexServer2x	12
	GC Time	Total GC Time in Milliseconds	Total GC time of JDBCServer2x	12000 ms
		Total GC Time in Milliseconds	Total GC time of JobHistory2x	12000 ms

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		Total GC Time in Milliseconds	Total GC time of IndexServer2x	12000 ms
Storm	Cluster	Number of Available Supervisors	Indicates the number of available Supervisor processes in the cluster in a measurement period.	1
		Slot Usage	Indicates the slot usage in the cluster in a measurement period.	80.0%
	Nimbus	Nimbus Heap Memory Usage Calculate	Indicates the Nimbus heap memory usage.	80%
Yarn	Resources	NodeManager Direct Memory Usage Statistics	Indicates the percentage of direct memory used by NodeManagers.	90%
		NodeManager Heap Memory Usage Statistics	Indicates the percentage of NodeManager heap memory usage.	95%
		NodeManager Non-Heap Memory Usage Statistics	Indicates the percentage of NodeManager non-heap memory usage.	90%
		ResourceManager Direct Memory Usage Statistics	Indicates the Kafka direct memory usage.	90%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		ResourceManager Heap Memory Usage Statistics	Indicates the ResourceManager heap memory usage.	95%
		ResourceManager Non-Heap Memory Usage Statistics	Indicates the ResourceManager non-heap memory usage.	90%
	Garbage collection	GC Time	Indicates the GC duration of NodeManager per minute.	12000 ms
		GC Time	Indicates the GC duration of ResourceManager per minute.	12000 ms
	Others	Failed Applications of root queue	Number of failed tasks in the root queue	50
		Terminated Applications of root queue	Number of killed tasks in the root queue	50
	CPU & memory	Pending Memory	Pending memory capacity	83886080MB
	Application	Pending Applications	Pending tasks	60
ZooKeeper	Connection	ZooKeeper Connections Usage	Indicates the percentage of the used connections to the total connections of ZooKeeper.	80%
	CPU & memory	Directmemory Usage Calculate	Indicates the ZooKeeper heap memory usage.	95%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		Heap Memory Usage Calculate	Indicates the ZooKeeper direct memory usage.	80%
	GC	ZooKeeper GC Duration per Minute	Indicates the GC time of ZooKeeper every minute.	12000 ms
Ranger	GC	UserSync GC Duration	UserSync garbage collection (GC) duration	12000 ms
		RangerAdmin GC Duration	RangerAdmin GC duration	12000 ms
		TagSync GC Duration	TagSync GC duration	12000 ms
	CPU & memory	UserSync Non-Heap Memory Usage	UserSync non-heap memory usage	80.0%
		UserSync Direct Memory Usage	UserSync direct memory usage	80.0%
		UserSync Heap Memory Usage	UserSync heap memory usage	95.0%
		RangerAdmin Non-Heap Memory Usage	RangerAdmin non-heap memory usage	80.0%
		RangerAdmin Heap Memory Usage	RangerAdmin heap memory usage	95.0%
		RangerAdmin Direct Memory Usage	RangerAdmin direct memory usage	80.0%

Service	Monitoring Indicator Group Name	Indicator Name	Description	Default Threshold
		TagSync Direct Memory Usage	TagSync direct memory usage	80.0%
		TagSync Non-Heap Memory Usage	TagSync non-heap memory usage	80.0%
		TagSync Heap Memory Usage	TagSync heap memory usage	95.0%
ClickHouse	Cluster Quota	Clickhouse service quantity quota usage in ZooKeeper	Quota of the ZooKeeper nodes used by a ClickHouse service	90%
		Capacity quota usage of the Clickhouse service in ZooKeeper	Capacity quota of ZooKeeper directory used by the ClickHouse service	90%
IoTDB	GC	IoTDBServer GC Duration	IoTDBServer garbage collection (GC) duration	12000 ms
	CPU & memory	IoTDBServer Heap Memory Usage	IoTDBServer heap memory usage	90%
		IoTDBServer Direct Memory Usage	IoTDBServer direct memory usage	90%

8.4.1.3 Configuring the Alarm Masking Status

Scenario

If you do not want FusionInsight Manager to report specified alarms in the following scenarios, you can manually mask the alarms.

- Some unimportant alarms and minor alarms need to be masked.
- When a third-party product is integrated with MRS, some alarms of the product are duplicated with the alarms of MRS and need to be masked.
- When the deployment environment is special, certain alarms may be falsely reported and need to be masked.

After an alarm is masked, new alarms with the same ID as the alarm are neither displayed on the **Alarm** page nor counted. The reported alarms are still displayed.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Alarm** > **Masking Setting**.

Step 3 In the **Masking Setting** area, select the specified service or module.

Step 4 Select an alarm from the alarm list.

Figure 8-32 Masking an alarm

The information about the alarm is displayed, including the alarm name, ID, severity, masking status, and operations can be performed on the alarm.

- The masking status includes **Display** and **Masking**.
- Operations include **Masking** and **Help**.

NOTE

You can filter specified alarms based on the masking status and alarm severity.

Step 5 Set the masking status for an alarm:

- Click **Masking**. In the displayed dialog box, click **OK** to change the alarm masking status to **Masking**.
- Click **Cancel Masking**. In the dialog box that is displayed, click **OK** to change the masking status of the alarm to **Display**.

----End

8.4.2 Log

8.4.2.1 Log Online Search

Scenario

FusionInsight Manager allows you to search for logs online and view the log content of components to locate faults.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Log** > **Online Search**.

Figure 8-33 Online Search

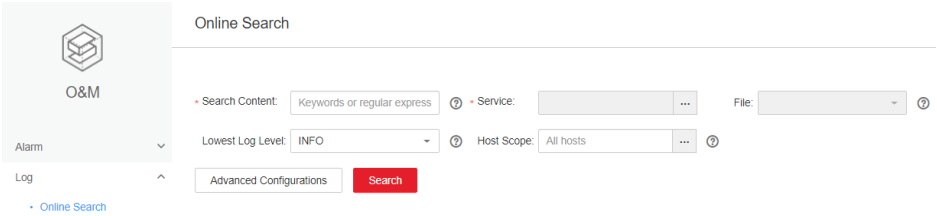

Step 3 Configure the parameters listed in **Table 8-20** to search for the logs you need. You can select a default log search duration (including **0.5h**, **1h**, **2h**, **6h**, **12h**, **1d**, **1w**, and **1m**), or click to customize **Start Data** and **End Data**.

Table 8-20 Log search parameters

Parameter	Description
Search Content	Keywords or regular expression to be searched for
Service	Service or module for which you want to query logs
File	Log files to be searched for when only one role is selected
Lowest Log Level	Lowest level of logs to be queried. After you select a level, the logs of this level and higher levels are displayed. The levels in ascending order are as follows: TRACE < DEBUG < INFO < WARN < ERROR < FATAL
Host Scope	• You can click ... to select hosts. • Enter the host name of the node for which you want to query logs or the IP address of the management plane. • Use commas (,) to separate IP addresses, for example, **192.168.10.10**,**192.168.10.11**. • Use hyphens (-) to indicate an IP address segment if the IP addresses are consecutive, for example, **192.168.10.[10-20]**. • Use hyphens (-) to indicate an IP address segment if the IP addresses are consecutive, and use commas (,) to separate IP address segments, for example, **192.168.10.[10-20,30-40]**. **NOTE** – If this parameter is not specified, all hosts are selected by default. – A maximum of 10 expressions can be entered at a time. – A maximum of 2,000 hosts can be matched for all entered expressions at a time.

Parameter	Description
Advanced Configurations	• **Max Quantity**: maximum number of logs that can be displayed at a time. If the number of queried logs exceeds the value of this parameter, the earliest logs will be ignored. If this parameter is not set, the maximum number of logs that can be displayed at a time is not limited. • **Timeout Duration**: log query timeout duration. This parameter is used to limit the maximum log query time on each node. When the query times out, the query is stopped and the logs that have been searched for are still displayed.

Step 4 Click **Search**. Table 8-21 describes the fields in search results.

Table 8-21 Parameters in search results

Parameter	Description
Time	Time when a line of log is generated
Host Name	Host name of the node where the log file recording the line of log is located
Location	Path of the log file recording the line of log Click the location information to go to the online log browsing page. By default, 100 lines of logs before and 100 lines after the line of log are displayed. You can click **Load More** on the top or bottom of the page to view more logs. Click **Download** to download the log file to the local PC.
Line No.	Line number of a line of log in the log file
Level	Level of the line of log
Log	Log content

NOTE

You can click **Stop** to forcibly stop the search. You can view the search results in the list.

Step 5 Click **Filter** to filter the logs to display on the page. Table 8-22 lists the fields that you can use to filter logs. After you configure these parameters, click **Filter** to search for logs meeting the search criteria. You can click **Reset** to clear the information that you have filled in.

Table 8-22 Parameters for filtering logs

Parameter	Description
Keywords	Keywords of the losg to be searched for
Host Name	Name of the host to be searched for
Location	Path of the log file to be searched for
Started	Start time for logs to be searched for
Completed	End time for logs to be searched for

----End

8.4.2.2 Log Download

Scenario

FusionInsight Manager allows you to batch export logs generated on all instances of each service.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Log** > **Download**.

Step 3 Select a log download range:

1. **Service**: Click ... and select a service.
2. **Host**: Enter the IP address of the host where the service is deployed. You can also click ... to select the required host.
3. **Max. Concurrent Nodes**: Set the maximum number of concurrent nodes for log collection as you need. MRS 3.3.0 or later supports this parameter.
4. Click in the upper right corner and configure **Start Time** and **End Time**.

Step 4 Click **Download**.

The downloaded log package contains the topology information of the start time and end time, helping you quickly find the log you need.

The topology file is named in the format of **topo_<***Topology structure change time***>.txt**. The file contains the node IP address, host name, and service instances that reside on the node. (OMS nodes are identified by **Manager:Manager**.)

Example:


```
192.168.204.124|suse-124|
DBService:DBServer;KrbClient:KerberosClient;LdapClient:SlapdClient;LdapServer:SlapdServer;Manager:Manager;meta:meta
```

----End

8.4.3 Perform a Health Check

8.4.3.1 Viewing a Health Check Task

Scenario

Administrators can view all health check tasks in the health check management center to check whether the cluster is affected after the modification.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Health Check**.

Figure 8-34 Health Check

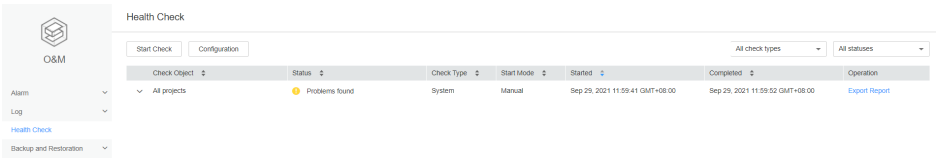

By default, all saved health check reports are listed. The parameters for a health check report are as follows:

Table 8-23 Parametes for a health check report

Parameter	Description
Check Object	Object to be checked. You can expand the list to view its details.
Status	Check result status. Value options are **No problems found**, **Problems found**, and **Checking**.
Check Type	Entity on which the check is to be performed. Value options are **System**, **Cluster**, **Host**, **Service**, and **OMS**. If you select **Cluster**, all items are checked by default.
Start Mode	Whether the health check is automatically or manually performed
Started	Start time of the check
Completed	End time of the check
Operation	Operations you can perform. Value options are **Export Report** and **View Help**.

NOTE

- In the upper right corner of the check list, you can filter health checks by check type or status.
- If **Check Type** is **Cluster**, **View Help** is displayed in the **Check Object** drop-down list.
- During a health check, the system determines whether check objects are healthy based on their historical monitoring metric data.

----End

8.4.3.2 Managing Health Check Reports

Scenario

FusionInsight Manager allows you to download and delete health check reports.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Health Check**.

Step 3 Locate the row containing the target health check report and click **Export Report** in the **Operation** to download the report.

----End

8.4.3.3 Modifying Health Check Configuration

Scenario

Administrators can enable automatic health check to reduce manual operation time. By default, the automatic health check checks the entire cluster.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Health Check** > **Configuration**.

Periodic Health Check indicates whether to enable automatic health check. Selecting **Enable** to enable the automatic health check, and selecting **Disable** to disable the function.

Set the health check period to **Daily**, **Weekly**, or **Monthly** as required.

Step 3 Click **OK** to save the configurations.

----End

8.4.4 Configuring Backup and Backup Restoration

8.4.4.1 Creating a Backup Task

Scenario

You can create backup tasks on FusionInsight Manager. Executing backup tasks backs up related data.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Backup and Restoration** > **Backup Management**. On the page that is displayed, click **Create**.

Figure 8-35 Creating a backup task

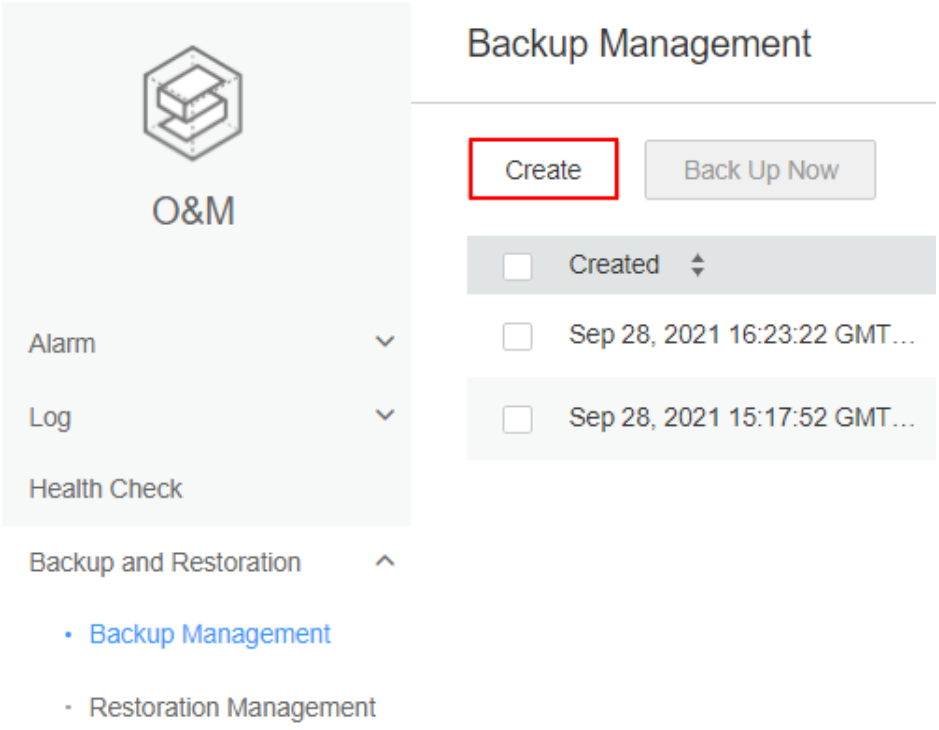

Step 3 Set **Backup Object** to **OMS** or the cluster whose data you want to back up.

Step 4 Enter a task name in the **Name** text box.

Step 5 Set **Mode** to **Periodic** or **Manual** as required.

Table 8-24 Backup types

Type	Parameter	Description
Periodic backup	Start Time	Indicates the time when a periodic backup task is started for the first time.
	Period	Task execution interval. Value options are **Hours** and **Days**.

Type	Parameter	Description
	Backup Policy	The following policies can be selected: • Full backup at the first time and subsequent incremental backup • Full backup every time • Full backup once every n times
Manual backup	N/A	You need to manually execute the task to back up data.

Step 6 Set required parameters in the **Configuration** area.

- Metadata and service data can be backed up.
- For details about how to back up data of different components, see **Backup and Recovery Management**.

Step 7 Click **OK** to save the configurations.

Step 8 In the backup task list, you can view the created backup task.

Locate the row that contains the target backup task, choose **More** > **Back Up Now** in the **Operation** column to execute the task immediately.

----End

8.4.4.2 Creating a Backup Restoration Task

Scenario

You can create a backup restoration task on FusionInsight Manager. After the restoration task is executed, the specified backup data is restored to the cluster.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Backup and Restoration** > **Restoration Management**. On the page that is displayed, click **Create**.

Step 3 Configure **Task Name**.

Step 4 Set **Recovery Object** to **OMS** or the cluster whose data you want to restore.

Step 5 Set the required parameters in the **Recovery Configuration** area.

- Metadata and service data can be restored.
- For details about how to restore data of different components, see **Backup and Recovery Management**.

Step 6 Click **OK** to save the configurations.

Step 7 In the restoration task list, you can view the created restoration tasks.

Locate the row containing the target restoration task, click **Start** in the **Operation** column to execute the restoration task immediately.

----End

8.4.4.3 Managing Backup and Backup Restoration Tasks

Scenario

You can also maintain and manage backup restoration tasks on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **O&M** > **Backup and Restoration** > **Backup Management** or **Restoration Management**.

Step 3 In the **Operation** column of the specified task in the task list, select the operation to be performed.

Table 8-25 Maintenance and management operations

Operation Entry	Description
Config	Modify parameters for the backup task.
Recover	After some service data is successfully backed up, you can use this function to quickly restore data.
More > **Back Up Now**	Perform this operation to execute the backup task immediately.
More > **Stop**	Perform this operation to stop a running task.
More > **Delete** or **Delete**	This operation is used to delete tasks.
More > **Suspend**	Perform this operation to disable the automatic backup task function.
More > **Resume**	Perform this operation to enable the automatic backup task function.
More > **View History** or **View History**	Perform this operation to switch to the task run log page to view the task running details and backup path.
View	Perform this operation to check the parameter settings of the restoration task.
Start	Perform this operation to run the restoration task.

----End

8.5 Audit

8.5.1 Overview

Scenario

The **Audit** page displays the user operations on Manager. On this page, administrators can view historical user operations on Manager. For details about the audit information, see **Audit Logs**.

Overview

Log in to FusionInsight Manager and choose **Audit**. The **Audit** page displays the operation type, risk level, start time, end time, user, host name, service, instance, and operation result.

Figure 8-36 Audit information list

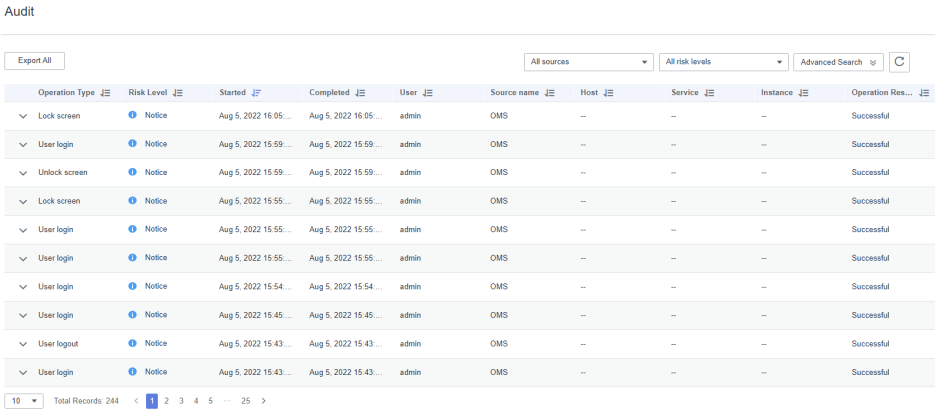

- You can select audit logs at the **Critical**, **Major**, **Minor**, or **Notice** level from the **All risk levels** drop-down list.
- In **Advanced Search**, you can set filter criteria to query audit logs.
  a. You can query audit logs by user management, cluster, service, and health in the **Operation Type** column.
  b. In the **Service** column, you can select a service to query corresponding audit logs.

     NOTE

     You can select **--** to search for audit logs using all other search criteria except services.

  c. You can query audit logs by operation result. Value options are **All**, **Successful**, **Failed**, and **Unknown**.
- You can click to manually refresh the current page or click to filter the columns displayed in the page.

- Click **Export All** to export all audit information at a time. The audit information can be exported in **TXT** or **CSV** format.

8.5.2 Configuring Audit Log Dumping

Scenario

The audit logs of FusionInsight Manager are stored in the database by default. If the audit logs are retained for a long time, the disk space of the data directory may be insufficient. To store audit logs to another archive server, administrators can set the required dump parameters to automatically dump these logs. This facilitates the management of audit logs.

If you do not configure the audit log dumping, the system automatically saves the audit logs to a file when the number of audit logs reaches 100,000 pieces. The path is **${BIGDATA_DATA_HOME}/dbdata_om/dumpData/iam/operatelog** on the active management node. The file name format is **OperateLog_store**_*YY_MM_DD_HH_MM_SS*.**csv**. A maximum of 50 historical audit log files can be saved.

NOTE

Archived audit logs will not be displayed on FusionInsight Manager. Only new audit logs generated after the old logs are automatically saved are displayed.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Audit** > **Configuration**.

Step 3 Click the switch on the right of **Audit Log Dumping Flag**.

Audit Log Dump is disabled by default. If  is displayed, **Audit Log Dump** is enabled.

Step 4 Set the dump parameters based on information provided in **Table 8-26**

Figure 8-37 Dump parameters

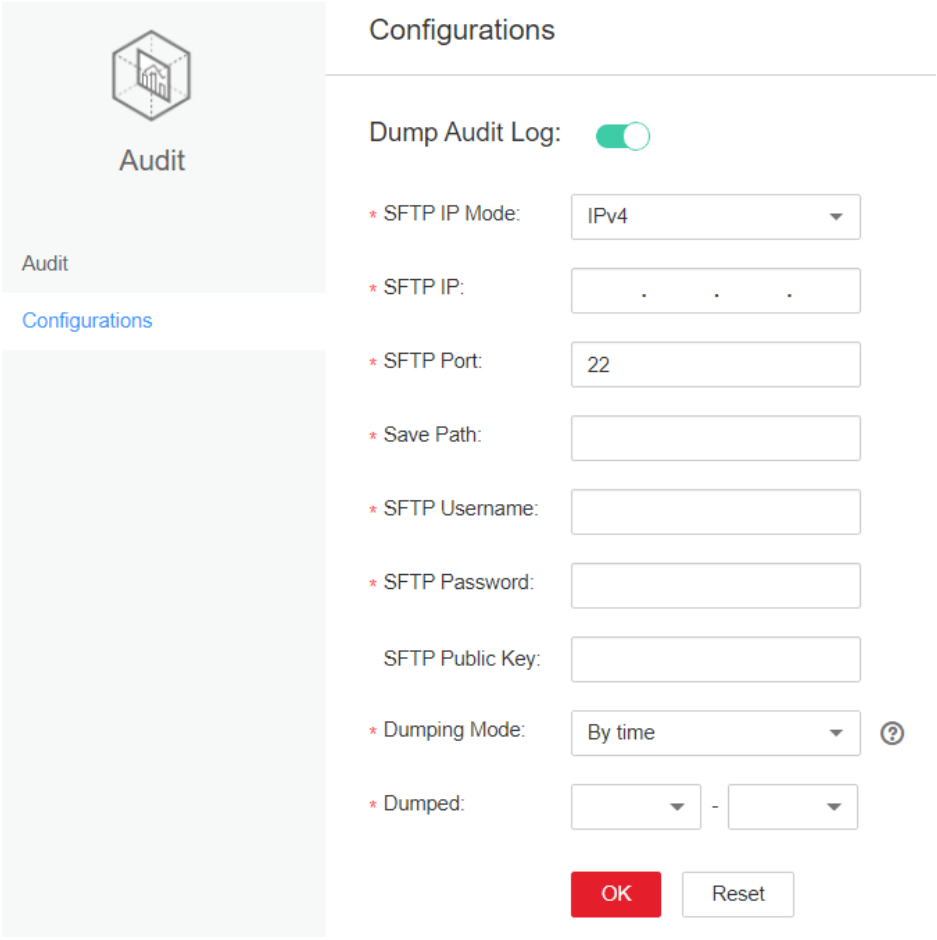

Table 8-26 Audit log dump parameters

Parameter	Description	Value
SFTP IP Mode	Mode of the destination IP address. The value can be **IPv4** or **IPv6**.	IPv4
SFTP IP	SFTP server for storing dumped audit logs. You are advised to use the SFTP service based on SSH v2 to prevent security risks.	**192.168.10.51** (example value)
SFTP Port	Connection port of the SFTP server for storing dumped audit logs	**22** (example value)
Save Path	Path for storing audit logs on the SFTP server	**/opt/omm/oms/auditLog** (example value)

Parameter	Description	Value
SFTP Username	User name for logging in to the SFTP server	**root** (example value)
SFTP Password	Password for logging in to the SFTP server	*Password for logging in to the SFTP server*
SFTP Public key	Specifies the public key of the SFTP server. This parameter is optional. You are advised to set the public key of the SFTP server. Otherwise, security risks may exist.	-
Dumping Mode	Dump mode. Value options are as follows: • **By Quantity**: If the number of pieces of logs reaches the value of this parameter (**100000** by default), the logs are dumped. • **By Time**: specifies the date when logs are dumped. The dumping frequency is once a year.	• By Quantity • By Time
Dumping Date	This parameter is available only when **Dumping Mode** is set to **By time**. After you select a dump date, the system starts dumping on this date. The logs to be dumped include all the audit logs generated before January 1 00:00 of the current year.	11-06

NOTE

If the SFTP public key is empty, the system displays a security risk warning. Evaluate the security risk and then save the configuration.

Step 5 Click **OK** to complete the settings.

NOTE

Key fields in the audit log dump file are as follows:

- **USERTYPE** indicates the user type. Value **0** indicates a human-machine user, and value **1** indicates a machine-machine user.
- **LOGLEVEL** indicates the security level. Value **0** indicates Critical, value **1** indicates Major, value **2** indicates Minor, and value **3** indicates Warning.
- **OPERATERESULT** indicates the operation result. Value **0** indicates that the operation is successful, and value **1** indicates that the operation is failed.

----End

8.6 Tenant Resources

8.6.1 Multi-Tenancy

8.6.1.1 Overview

Definition

Multi-tenancy refers to multiple resource sets (a resource set is a tenant) in the MRS big data cluster and is able to allocate and schedule resources. The resources include computing resources and storage resources.

Context

Modern enterprises' data clusters are becoming more and more centralized and cloud-based. Enterprise-class big data clusters must meet the following requirements:

- Carry data of different types and formats and run jobs and applications of different types (such analysis, query, and stream processing).
- Isolate data of a user from that of another user who has demanding requirements on data security, such as a bank or government institute.

The preceding requirements bring the following challenges to the big data clusters:

- Proper allocation and scheduling of resources to ensure stable operating of applications and jobs.
- Strict access control to ensure data and service security.

Multi-tenancy isolates the resources of a big data cluster into resource sets. Users can lease desired resource sets to run applications and jobs and store data. In a big data cluster, multiple resource sets can be deployed to meet diverse requirements of multiple users.

The MRS big data cluster provides a complete enterprise-class big data multi-tenant solution.

Highlights

- Proper resource configuration and isolation

  The resources of a tenant are isolated from those of another tenant. The resource use of a tenant does not affect other tenants. This mechanism ensures that each tenant can configure resources based on service requirements, improving resource utilization.
- Resource consumption measurement and statistics

  Tenants are system resource applicants and consumers. System resources are planned and allocated based on tenants. Resource consumption by tenants can be measured and collected.

- Assured data security and access security

  In multi-tenant scenarios, the data of each tenant is stored separately to ensure data security. The access to tenants' resources is controlled to ensure access security.

8.6.1.2 Technical Principles

8.6.1.2.1 Multi-Tenant Management

Unified Multi-Tenant Management

Log in to FusionInsight Manager and choose **Tenant Resources** > **Tenant Resources Management**. On the page that is displayed, you can find that FusionInsight Manager is a unified multi-tenant management platform that integrates multiple functions such as tenant lifecycle management, tenant resource configuration, tenant service association, and tenant resource usage statistics, delivering a mature multi-tenant management model and achieving centralized tenant and service management.

Graphical User Interface

FusionInsight Manager provides the graphical multi-tenant management interface and manages and operates multiple levels of tenants using the tree structure. Additionally, FusionInsight Manager integrates the basic information and resource quota of the current tenant in one interface to facilitate O&M and management, as shown in **Figure 8-38**.

Figure 8-38 Tenant management page of FusionInsight Manager

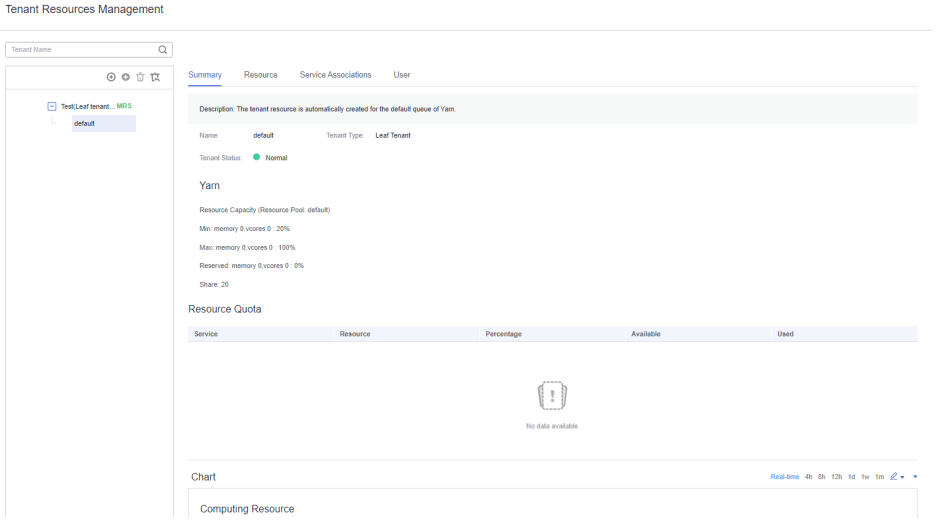

Hierarchical Tenant Management

FusionInsight Manager supports a hierarchical tenant management model in which you can add sub-tenants to an existing tenant to re-configure resources. Sub-tenants of level-1 tenants are level-2 tenants. So on and so forth. FusionInsight Manager provides enterprises with a field-tested multi-tenant management model, enabling centralized tenant and service management.

Simplified Permission Management

FusionInsight Manager hides internal permission management details from common users and simplifies permission management operations for administrators, improving usability and user experience of tenant permission management.

- FusionInsight Manager employs role-based access control (RBAC) to configure different permissions for users based on service scenarios during multi-tenant management.
- The administrator of tenants has tenant management permissions, including viewing resources and services of the current tenant, adding or deleting sub-tenants of the current tenant, and managing permissions of sub-tenants' resources. FusionInsight Manager supports setting of the administrator for a single tenant so that the management over this tenant can be delegated to a user who is not the system administrator.
- Roles of a tenant have all permissions on the computing resources and storage resources of the tenant. When a tenant is created, the system automatically creates roles for this tenant. You can add a user and bind the user to the tenant roles so that the user can use the resources of the tenant.

Clear Resource Management

- **Self-Service Resource Configuration**

  In FusionInsight Manager, you can configure the computing resources and storage resources during the creation of a tenant and add, modify, or delete the resources of the tenant.

  Permissions of the roles that are associated with a tenant are updated automatically when you modify the computing or storage resources of the tenant.

- **Resource Usage Statistics**

  Resource usage statistics are critical for administrators to determine O&M activities based on the status of cluster applications and services, improving the cluster O&M efficiency. FusionInsight Manager displays the resource statistics of tenants in **Resource Quota**, including the vCores, memory, and HDFS storage resources.

  NOTE

  - **Resource Quota** dynamically calculates the resource usage of tenants.

Service	Resource	Percentage	Available	Used
HDFS	Space	0.00%	20.00 GB	0 MB
Yarn	Memory	0.00%	8.00 GB	0 MB
Yarn	CPU	0.00%	4 vCores	0 vCores

    The available resources of the Superior scheduler are calculated as follows:

    - Superior

      The available Yarn resources (memory and CPU) are allocated in proportion based on the queue weight.

  - When the tenant administrator is bound to a tenant role, the tenant administrator has the permissions to manage the tenant and use all resources of the tenant.

- **Graphical Resource Monitoring**

  Graphical resource monitoring supports the graphical display of monitoring metrics listed in **Table 8-27**, as shown in **Figure 8-39**.

Figure 8-39 Refined monitoring

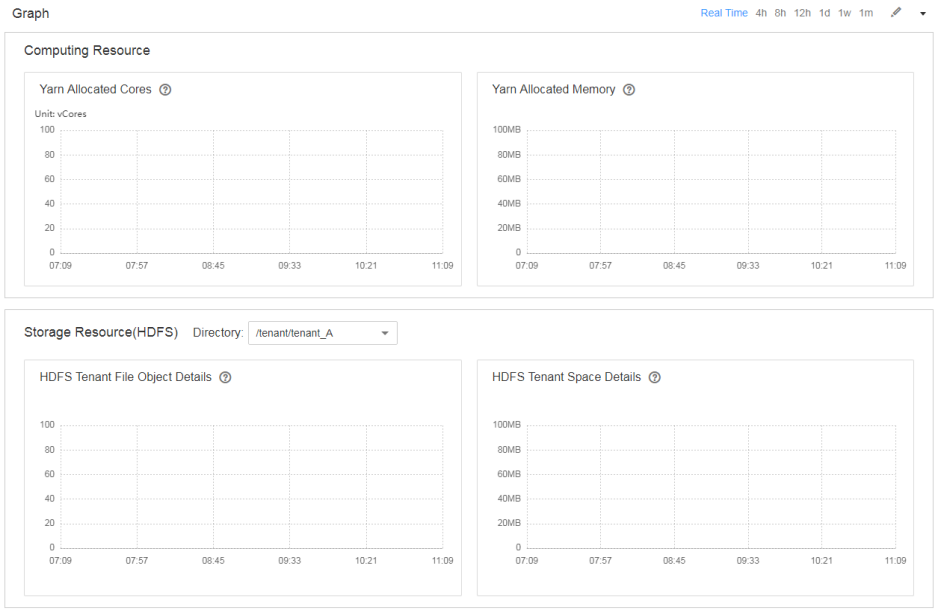

By default, the real-time monitoring data is displayed. You can click to customize a time range. The default time ranges include 4 hours, 8 hours, 12 hours, 1 day, 1 week, and 1 month. Click and select **Export** to export the monitoring metric information.

Table 8-27 Monitoring metrics

Service	Metric Item	Description
HDFS	HDFS Tenant Space Details • Allocated Space • Used Space	HDFS can monitor a specified storage directory. The storage directory is the same as the directory added by the current tenant in **Resource**.
	HDFS Tenant File Object Details • Number of Used File Objects	
Yarn	Yarn Allocated Cores • Maximum Number of CPU Cores in an AM • Allocated Cores • Number of Used CPU Cores in an AM	Monitoring information of the current tenant is displayed. If no sub-item is configured for a tenant, this information is not displayed. The monitoring data is obtained from **Scheduler** > **Application Queues** > **Queue:** *Tenant name* on the native web UI of Yarn.

Service	Metric Item	Description
	Yarn Allocated Memory • Allocated Maximum AM Memory • Allocated Memory • Used AM Memory	

8.6.1.2.2 Multi-Tenant Model

Related Model

The following figure shows a multi-tenant model.

Figure 8-40 Multi-tenant model

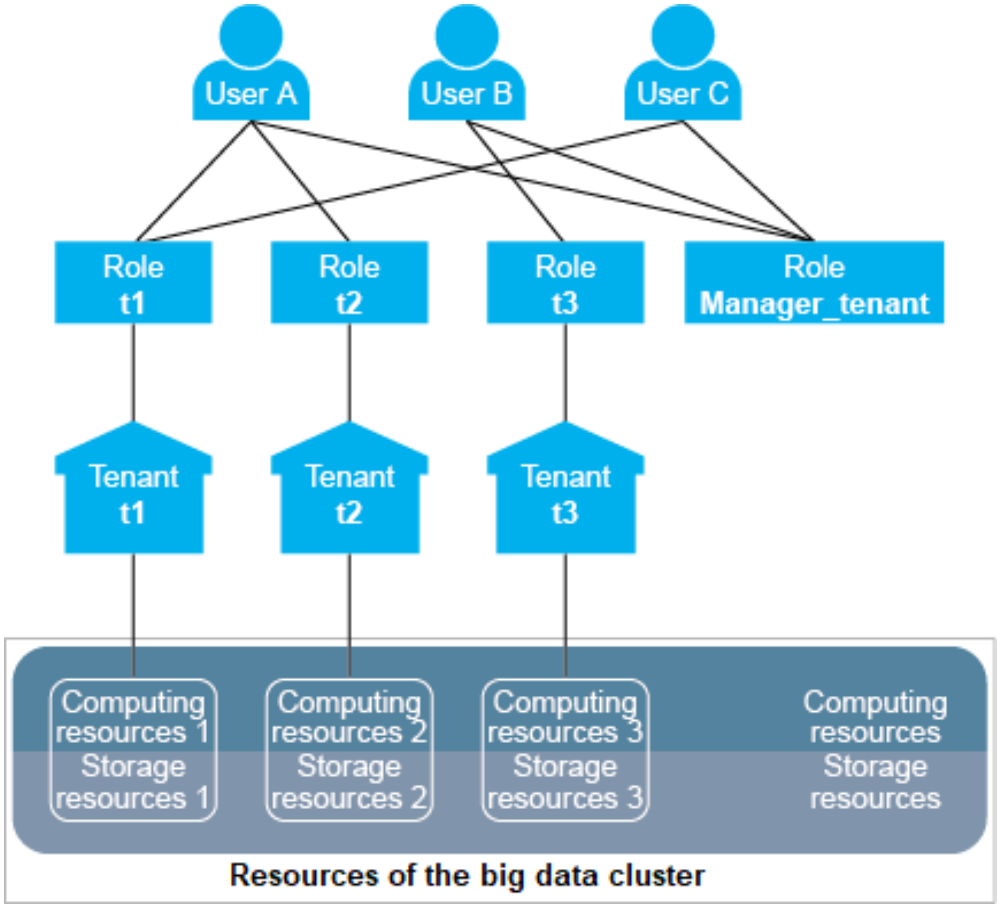

Table 8-28 describes the concepts involved in **Figure 8-40**.

Table 8-28 Concepts in the model

Concept	Description
User	A natural person who has a username and password and uses the big data cluster. There are three different users in **Figure 8-40**: user A, user B, and user C.
Role	A role is a carrier of one or more permissions. Permissions are assigned to specific objects, for example, access permissions for the **/tenant** directory in HDFS. **Figure 8-40** shows four roles: **t1**, **t2**, **t3**, and **Manager_tenant**. • Roles **t1**, **t2**, and **t3** are automatically generated when tenants are created. The role names are the same as the tenant names. That is, roles **t1**, **t2**, and **t3** map to tenants **t1**, **t2**, and **t3**. Role names and tenant names need to be used in pair. • Role **Manager_tenant** is defaulted in the cluster and cannot be used separately.
Tenant	A tenant is a resource set in a big data cluster. Multiple tenants are referred to as multi-tenancy. The resource sets further divided under a tenant are called sub-tenants. **Figure 8-40** shows three tenants: **t1**, **t2**, and **t3**.
Resource	• Computing resources include CPUs and memory. The computing resources of a tenant are allocated from the total computing resources in the cluster. One tenant cannot occupy the computing resources of another tenant. In **Figure 8-40**, computing resources 1, 2, and 3 are allocated for tenants **t1**, **t2**, and **t3** respectively from the cluster's computing resources. • Storage resources include disks and third-party storage systems. The storage resources of a tenant are allocated from the total storage resources in the cluster. One tenant cannot occupy the storage resources of another tenant. In **Figure 8-40**, storage resources 1, 2, and 3 are allocated for tenants **t1**, **t2**, and **t3** respectively from the cluster's storage resources.

If a user wants to use a tenant's resources or add or delete a sub-tenant of a tenant, the user needs to be bound to both the tenant role and role **Manager_tenant**. **Table 8-29** lists the roles bound to each user in **Figure 8-40**.

Table 8-29 Roles bound to each user

User	Role	Permission
User A	• Role **t1** • Role **t2** • Role **Manager_tenant**	• Uses the resources of tenants **t1** and **t2**. • Adds or deletes sub-tenants of tenants **t1** and **t2**.
User B	• Role **t3** • Role **Manager_tenant**	• Uses the resources of tenant **t3**. • Adds or deletes sub-tenants of tenant **t3**.
User C	• Role **t1** • Role **Manager_tenant**	• Uses the resources of tenant **t1**. • Adds or deletes sub-tenants of tenant **t1**.

A user can be bound to multiple roles, and one role can also be bound to multiple users. Users are associated with tenants after being bound to the tenant roles. Therefore, tenants and users form a many-to-many relationship. One user can use the resources of multiple tenants, and multiple users can use the resources of the same tenant. For example, in **Figure 8-40**, user A uses the resources of tenants **t1** and **t2**, and users A and C uses the resources of tenant **t1**.

NOTE

The concepts of a parent tenant, sub-tenant, level-1 tenant, and level-2 tenant are all designed for the multi-tenant service scenarios. Pay attention to the differences these concepts and the concepts of a leaf tenant resource and non-leaf tenant resource on FusionInsight Manager.

- Level-1 tenant: determined based on the tenant's level. For example, the first created tenant is a level-1 tenant and its sub-tenant is a level-2 tenant.
- Parent tenant and sub-tenant: indicates the hierarchical relationship between tenants.
- Non-leaf tenant resource: indicates the tenant type selected during tenant creation. This tenant type can be used to create sub-tenants.
- Leaf tenant resource: indicates the tenant type selected during tenant creation. This tenant type cannot be used to create sub-tenants.

Multi-Tenant Platform

Tenant is a core concept of the FusionInsight big data platform. It plays an important role in big data platforms' transformation from user-centered to multi-tenant to keep up with enterprises' multi-tenant application environments. **Figure 8-41** shows the transformation of big data platforms.

Figure 8-41 Platform transformation from user-centered to multi-tenant

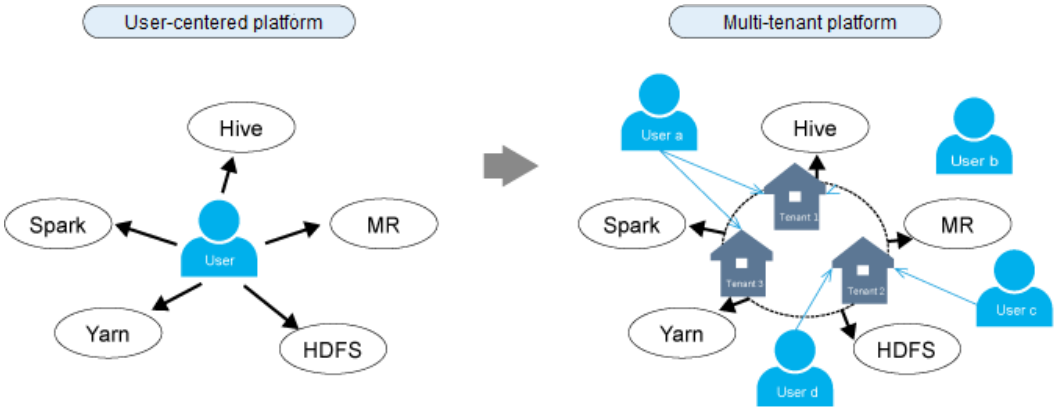

On a user-centered big data platform, users can directly access and use all resources and services.

- However, user applications may use only partial cluster resources, resulting in low resource utilization.
- The data of different users may be stored together, decreasing data security.

On a multi-tenant big data platform, users use required resources and services by accessing the tenants.

- Resources are allocated and scheduled based on application requirements and used based on tenants, increasing resource utilization.
- Users can access the resources of tenants only after being associated with tenant roles, enhancing access security.
- The data of tenants is isolated, ensuring data security.

8.6.1.2.3 Resource Overview

MRS cluster resources are classified into computing resources and storage resources. The multi-tenant architecture implements resource isolation.

- **Computing resources**

  Computing resources include CPUs and memory. One tenant cannot occupy the computing resources of another tenant.
- **Storage resources**

  Storage resources include disks and third-party storage systems. One tenant cannot access the data of another tenant.

Computing Resources

Computing resources are divided into static service resources and dynamic resources.

- **Static Service Resources**

  Static service resources are computing resources allocated to each service and are not shared between services. The total computing resources of each service are fixed. These services include Flume, HBase, HDFS, and Yarn.
- **Dynamic Resources**

Dynamic resources are computing resources dynamically scheduled to a job queue by the distributed resource management service Yarn. Yarn dynamically schedules resources for the job queues of MapReduce, Spark2x, Flink, and Hive.

NOTE

The resources allocated to Yarn in a big data cluster are static service resources but can be dynamically allocated to job queues by Yarn.

Storage Resources

Storage resources are data storage resources that can be allocated by the distributed file storage service HDFS. Directory is the basic unit of allocating HDFS storage resources. Tenants can obtain storage resources from the specified directories in the HDFS file system.

8.6.1.2.4 Dynamic Resources

Overview

Yarn provides distributed resource management for a big data cluster. The total volume of resources allocated to Yarn can be configured. Then Yarn allocates and schedules computing resources for job queues. The computing resources of MapReduce, Spark, Flink, and Hive job queues are allocated and scheduled by Yarn.

Yarn queues are fundamental units of scheduling computing resources.

The resources obtained by tenants using Yarn queues are dynamic resources. Users can dynamically create and modify the queue quotas and view the status and statistics of the queues.

Resource Pools

Nowadays, enterprise IT systems often face complex cluster environments and diverse upper-layer requirements. For example:

- Heterogeneous cluster: The computing speed, storage capacity, and network performance of each node in the cluster are different. All the tasks of complex applications need to be properly allocated to each compute node in the cluster based on service requirements.
- Computing isolation: Data must be shared among multiple departments but computing resources must be distributed onto different compute nodes.

These require that the compute nodes be further partitioned.

Resource pools are used to specify the configuration of dynamic resources. Yarn queues are associated with resource pools for resource allocation and scheduling.

One tenant can have only one default resource pool. Users can be bound to the role of a tenant to use the resources in the resource pool of the tenant. To use resources in multiple resource pools, a user can be bound to roles of multiple tenants.

Scheduling Mechanism

Yarn dynamic resources support label-based scheduling. This policy creates labels for compute nodes (Yarn NodeManagers) and adds the compute nodes with the same label into the same resource pool. Then Yarn dynamically associates the queues with resource pools based on the resource requirements of the queues.

For example, a cluster has more than 40 nodes which are labeled by **Normal**, **HighCPU**, **HighMEM**, or **HighIO** based on their hardware and network configurations and added into four resource pools, respectively. **Table 8-30** describes the performance of each node in the resource pool.

Table 8-30 Performance of each node in a resource pool

Label	Number of Nodes	Hardware and Network Configuration	Added To	Associated With
Normal	10	General	Resource pool A	Common queue
HighCPU	10	High-performance CPU	Resource pool B	Computing-intensive queue
HighMEM	10	Large memory	Resource pool C	Memory-intensive queue
HighIO	10	High-performance network	Resource pool D	I/O-intensive queue

A queue can use only the compute nodes in its associated resource pool.

- A common queue is associated with resource pool A and uses **Normal** nodes with general hardware and network configurations.
- A computing-intensive queue is associated with resource pool B and uses **HighCPU** nodes with high-performance CPUs.
- A memory-intensive queue is associated with resource pool C and uses **HighMEM** nodes with large memory.
- An I/O-intensive queue is associated with resource pool C and uses **HighIO** nodes with high-performance network.

Yarn queues are associated with specified resource pools to efficiently utilize resources in resource pools and maximize node performance.

FusionInsight Manager supports a maximum of 50 resource pools. The system has a default resource pool.

Schedulers

By default, the Superior scheduler is enabled for the MRS cluster.

- The Superior scheduler is an enhanced version and named after the Lake Superior, indicating that the scheduler can manage a large amount of data.

To meet enterprise requirements and tackle scheduling challenges faced by the Yarn community, the Superior scheduler makes the following enhancements:

- Enhanced resource sharing policy

  The Superior scheduler supports queue hierarchy. It integrates the functions of open-source schedulers and shares resources based on configurable policies. In terms of instances, administrators can use the Superior scheduler to configure an absolute value or percentage policy for queue resources. The resource sharing policy of the Superior scheduler enhances label-based scheduling of Yarn as a resource pool feature. The nodes in the Yarn cluster can be grouped based on the capacity or service type to ensure that queues can more efficiently utilize resources.

- Tenant-based resource reservation policy

  Some tenants may run critical tasks at some time, and their resource requirements must be preferentially addressed. The Superior scheduler builds a mechanism to support the resource reservation policy. Reserved resources can be allocated to the critical tasks running in the specified tenant queues in a timely manner to ensure proper task execution.

- Fair sharing among tenants and resource pool users

  The Superior scheduler allows shared resources to be configured for users in a queue. Each tenant may have users with different weights. Heavily weighted users may require more shared resources.

- Ensured scheduling performance in a big cluster

  The Superior scheduler receives heartbeats from each NodeManager and saves resource information in memory, which enables the scheduler to control cluster resource usage globally. The Superior scheduler uses the push scheduling model, which makes the scheduling more precise and efficient and remarkably improves cluster resource utilization. Additionally, the Superior scheduler delivers excellent performance when the interval between NodeManager heartbeats is long and prevents heartbeat storms in big clusters.

- Priority policy

  If the minimum resource requirement of a service cannot be met after the service obtains all available resources, a preemption occurs. The preemption function is disabled by default.

8.6.1.2.5 Storage Resources

Overview

As a distributed file storage service in a big data cluster, HDFS stores all the user data of the upper-layer applications in the big data cluster, including the data written to HBase tables or Hive tables.

A directory is the basic unit of allocating HDFS storage resources. HDFS supports the conventional hierarchical file structure. Users or applications can create directories and create, delete, move, or rename files in directories. Tenants can obtain storage resources from specified directories in the HDFS file system.

Scheduling Mechanism

HDFS directories can be stored on nodes with specified labels or disks of specified hardware types. For example:

- When both real-time query and data analysis tasks are running in the same cluster, the real-time query tasks need to be deployed only on certain nodes, and the task data must also be stored on these nodes.
- Based on actual service requirements, key data needs to be stored on highly reliable nodes.

Administrators can flexibly configure HDFS data storage policies based on actual service requirements and data features to store data on specified nodes.

For tenants, storage resources refer to the HDFS resources they use. Data of specified directories can be stored to the tenant-specified storage paths, thereby implementing storage resource scheduling and ensuring data isolation between tenants.

Users can add or delete HDFS storage directories of tenants and set the file quantity quota and storage capacity quota of directories to manage storage resources.

8.6.1.3 Multi-Tenancy Usage

8.6.1.3.1 Overview

Tenants are used in resource control and service isolation scenarios. Administrators need to determine the service scenarios of cluster resources and then plan tenants.

NOTE

- Yarn in a new cluster uses the Superior scheduler by default. For details, see **Using the Superior Scheduler**.

Multi-tenancy involves three types of operations: creating a tenant, managing tenants, and managing resources. **Table 8-31** describes these operations.

Table 8-31 Multi-tenant operations

Operation	Action	Description
Creating a tenant	• Add a tenant. • Add a sub-tenant. • Create a user and bind the user to the role of a tenant.	During the creation of a tenant, you can configure its computing resources, storage resources, and associated services based on service requirements. In addition, you can add users to the tenant and bind necessary roles to these users. A user to create a level-1 tenant needs to be bound to the **Manager_administrator** or **System_administrator** role. A user to create a sub-tenant needs to be bound to the role of the parent tenant at least.

Operation	Action	Description
Managing tenants	• Manage the tenant directory. • Restore tenant data. • Clear non-associated queues of a tenant. • Delete a tenant.	You can edit tenants as services change. A user to manage or delete a level-1 tenant or restore tenant data needs to be bound to the **Manager_administrator** or **System_administrator** role. A user to manage or delete a sub-tenant needs to be bound to the role of the parent tenant at least.
Managing resources	• Create a resource pool. • Modify a resource pool. • Delete a resource pool. • Configure a queue. • Configure the queue capacity policy of a resource pool. • Clear configurations of a queue.	You can reconfigure resources for tenants as the services change. A user to manage resources needs to be bound to the **Manager_administrator** or **System_administrator** role.

8.6.1.3.2 Process Overview

Administrators need to determine the service scenarios of cluster resources and then plan tenants. After that, administrators add tenants and configure dynamic resources, storage resources, and associated services for the tenants on FusionInsight Manager.

Process Overview shows the process for creating a tenant.

Figure 8-42 Creating a tenant

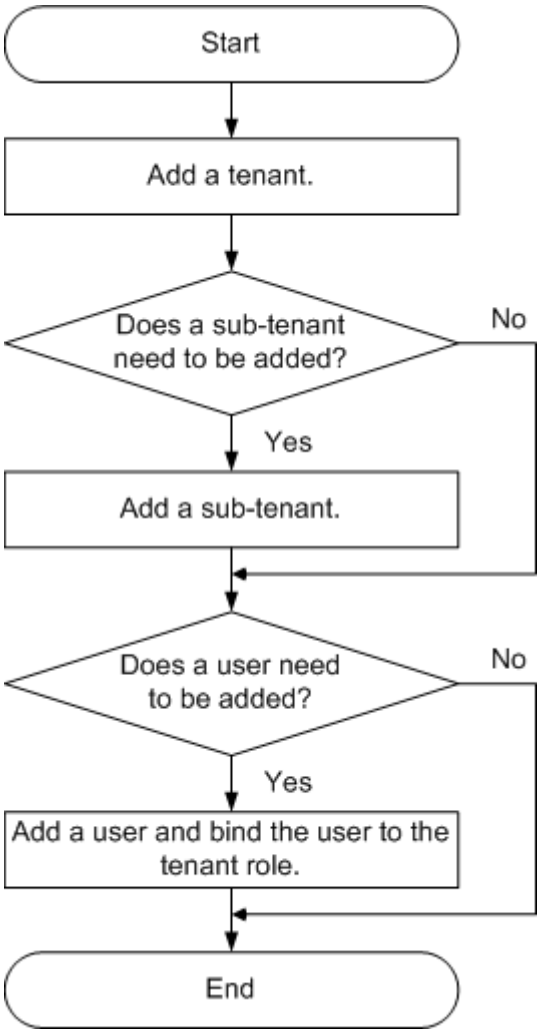

Table 8-32 describes the operations for creating a tenant.

Table 8-32 Operations for creating a tenant

Operation	Description
Add a tenant.	You can configure the computing resources, storage resources, and associated services of the tenant.
Add a sub-tenant.	You can configure the computing resources, storage resources, and associated services of the sub-tenant.
Add a user and bind the user to the tenant role.	If a user wants to use the resources of tenant **tenant1** or add or delete sub-tenants for **tenant1**, the user must be bound to both the **Manager_tenant** and **tenant1**_*Cluster ID* roles.

8.6.2 Using the Superior Scheduler

8.6.2.1 Creating Tenants

8.6.2.1.1 Adding a Tenant

Scenario

You can create tenants on FusionInsight Manager based on the resource consumption and isolation planning and requirements of services.

Prerequisites

- A tenant name has been planned 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.
- Resources to be allocated to the current tenant have been planned to ensure that the sum of resources of direct sub-tenants at each level does not exceed the resources of the current tenant.

Procedure

Step 1 Log in to FusionInsight Manager and choose **Tenant Resources**.

Step 2 Click . On the page that is displayed, configure tenant attributes according to **Table 8-33**.

Table 8-33 Tenant parameters

Parameter	Description
Cluster	Indicates the cluster for which you want to create a tenant. (This parameter is unavailable for clusters of MRS 3.3.0 or later.)
Name	• Indicates the name of the current tenant. The value consists of 3 to 50 characters, including digits, letters, and underscores (_). • Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or Yarn queue that exists in the current cluster.

Parameter	Description
Tenant Resource Type	Specifies whether the tenant is a leaf tenant. ● When **Leaf Tenant Resource** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. ● When **Non-leaf Tenant Resource** is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant. **NOTE** MRS 3.2.0 or later: If you select **ClickHouse** for **Service**, this parameter can only be set to **Leaf Tenant**.
Computing Resource	Specifies the dynamic computing resources for the current tenant. ● When **Yarn** is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name. – A leaf tenant can directly submit jobs to the queue. – A non-leaf tenant cannot directly submit jobs to the queue. However, Yarn adds an extra queue (hidden) named **default** for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue. ● If **Yarn** is not selected, the system does not automatically create a queue.
Configuration Mode	Indicates the configuration mode of computing resource parameters. ● If you select **Basic**, you only need to set **Default Resource Pool Capacity (%)**. ● If you select **Advanced**, you can manually configure the resource allocation weight and the minimum, maximum, and reserved resources of the tenant.
Default Resource Pool Capacity (%)	Indicates the percentage of computing resources used by the current tenant in the default resource pool. The value ranges from **0** to **100%**.
Weight	Indicates the resource allocation weight. The value ranges from **0** to **100**.

Parameter	Description
Minimum Resource	Indicates the resources guaranteed for the tenant (preemption supported). The value can be a percentage or an absolute value of the parent tenant's resources. When a tenant has a light workload, the resources of the tenant are automatically allocated to other tenants. When the available tenant resources are less than the value of **Minimum Resource**, the tenant can preempt the resources that have been lent to other tenants.
Maximum Resource	Indicates the maximum resources that can be used by the tenant. The tenant cannot obtain more resources than the value configured. The value can be a percentage or an absolute value of the parent tenant's resources.
Reserved Resource	Indicates the resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is running in the current tenant resources. The value can be a percentage or an absolute value of the parent tenant's resources.
Storage Resource	Specifies storage resources for the current tenant. ● When **HDFS** is selected, the system automatically allocates storage resources. ● When **HDFS** is not selected, the system does not automatically allocate storage resources.
Quota	Indicates the quota for files and directories.
Space Quota	Indicates the quota for the HDFS storage space used by the current tenant. ● If the unit is set to **MB**, the value ranges from **1** to **8796093022208**. If the unit is set to **GB**, the value ranges from **1** to **8589934592**. ● This parameter indicates the maximum HDFS storage space that can be used by the tenant, but not the actual space used. ● 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.
Storage Path	Indicates an HDFS directory for tenant resource data. ● The system automatically creates a folder named after the tenant name in the **/tenant** directory by default. For example, the default HDFS storage directory for tenant **ta1** is **/tenant/ta1**. ● When a tenant is created for the first time, the system creates the **/tenant** directory in the HDFS root directory. The storage path is customizable.

Parameter	Description
Service	Specifies whether to associate resources of other services. For details, see **Step 4**.
Description	Indicates the description of the current tenant.

NOTE

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

- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System** > **Permission** > **Role**. The role name is in the format of *Tenant name_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and bind the user to the role of the tenant. For details, see **Adding a User and Binding the User to a Tenant Role**.
- During the tenant creation, the system automatically creates a Yarn queue named after the tenant. If the queue name already exists, the new queue is named **Tenant name-***N*. *N* indicates a natural number starting from **1**. When a same name exists, the value *N* increases automatically to differentiate the queue from others. For example, **saletenant**, **saletenant-1**, and **saletenant-2**.

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, and click **OK**.

- Set **Service** to **HBase** and **Association Type** to **Exclusive** or **Shared**.

  NOTE

  - **Exclusive** indicates that the service resources are used by the tenant exclusively and cannot be associated with other tenants.
  - **Shared** indicates that the service resources can be shared with other tenants.

- MRS 3.2.0 or later: Select **ClickHouse** for **Service**.
  - **Association Type**: When **Service** is set to **ClickHouse**, **Association Type** can only be set to **Shared**. MRS 3.3.0 and later versions support **Exclusive** and **Shared**.
  - **Associate Logical Cluster**: If the logical cluster function is not enabled for ClickHouse, **default_cluster** is selected by default. If the function is enabled, select the logical cluster to which you want to associate.
  - **CPU Priority**: The CPU priority ranges from –20 to 19. This value is associated with the NICE value of the OS. A smaller value indicates a higher CPU priority.
  - **Memory**: The maximum value of this parameter is **100**, in percentage. For example, if this parameter is set to **80**, the total memory that can be used by the current tenant is calculated as follows: Available memory x 80%.

  - **Concurrency**: The maximum number of concurrent resources available for all the users bound to the tenant. This parameter is required for clusters of MRS 3.3.0 or later.

NOTE

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

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

For clusters of MRS 3.3.0 or later, click to export basic information about all tenants in the current cluster.

----End

8.6.2.1.2 Adding a Sub-Tenant

Scenario

You can create sub-tenants on FusionInsight Manager and allocate resources of the current tenant to the sub-tenants based on the resource consumption and isolation planning and requirements of services.

Prerequisites

- A parent non-leaf tenant has been added.
- A sub-tenant name has been planned 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.
- Resources to be allocated to the current tenant have been planned to ensure that the sum of resources of direct sub-tenants at each level does not exceed the resources of the current tenant.

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 attributes for the sub-tenant according to **Table 8-34**.

Table 8-34 Sub-tenant parameters

Parameter	Description
Cluster	Indicates the cluster to which the parent tenant belongs.
Parent Tenant Resource	Indicates the name of the parent tenant.
Name	• Indicates the name of the current tenant. The value consists of 3 to 50 characters, including digits, letters, and underscores (_). • 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.
Tenant Resource Type	Specifies whether the tenant is a leaf tenant. • When **Leaf Tenant Resource** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. • When **Non-leaf Tenant Resource** 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.
Computing Resource	Specifies the dynamic computing resources for the current tenant. • When **Yarn** is selected, the system automatically creates a queue in Yarn and the queue is named the same as the sub-tenant name. – A leaf tenant can directly submit jobs to the queue. – A non-leaf tenant cannot directly submit jobs to the queue. However, Yarn adds an extra queue (hidden) named **default** for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue. • If **Yarn** is not selected, the system does not automatically create a queue.
Configuration Mode	Indicates the configuration mode of computing resource parameters. • If you select **Basic**, you only need to set **Default Resource Pool Capacity (%)**. • If you select **Advanced**, you can manually configure the resource allocation weight and the minimum, maximum, and reserved resources of the tenant.

Parameter	Description
Default Resource Pool Capacity (%)	Indicates the percentage of computing resources used by the current tenant. The base value is the total resources of the parent tenant.
Weight	Indicates the resource allocation weight. The value ranges from **0** to **100**.
Minimum Resource	Indicates the resources guaranteed for the tenant (preemption supported). The value can be a percentage or an absolute value of the parent tenant's resources. When a tenant has a light workload, the resources of the tenant are automatically allocated to other tenants. When the available tenant resources are less than the value of **Minimum Resource**, the tenant can preempt the resources that have been lent to other tenants.
Maximum Resource	Indicates the maximum resources that can be used by the tenant. The tenant cannot obtain more resources than the value configured. The value can be a percentage or an absolute value of the parent tenant's resources.
Reserved Resource	Indicates the resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is running in the current tenant resources. The value can be a percentage or an absolute value of the parent tenant's resources.
Storage Resource	Specifies storage resources for the current tenant. • When **HDFS** is selected, the system automatically creates a folder named after the sub-tenant in the HDFS parent tenant directory. • When **HDFS** is not selected, the system does not automatically allocate storage resources.
Quota	Indicates the quota for files and directories.

Parameter	Description
Space Quota	Indicates the quota for the HDFS storage space used by the current tenant. • If the unit is set to **MB**, the value ranges from **1** to **8796093022208**. If the unit is set to **GB**, the value ranges from **1** to **8589934592**. • This parameter indicates the maximum HDFS storage space that can be used by the tenant, but not the actual space used. • 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. • 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.
Storage Path	Indicates the HDFS storage directory for the tenant. • 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 **ta1s** and the parent directory is **/tenant/ta1**, the storage path for the sub-tenant is then **/tenant/ta1/ta1s**. • The storage path is customizable in the parent directory.
Service	Specifies whether to associate resources of other services. For details, see **Step 4**.
Description	Indicates the description of the current tenant.

NOTE

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

- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System** > **Permission** > **Role**. The role name is in the format of *Tenant name_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and bind the user to the role of the tenant. For details, see **Adding a User and Binding the User to a Tenant Role**.
- 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

8.6.2.1.3 Adding a User and Binding the User to a Tenant Role

Scenario

A newly created tenant cannot directly log in to the cluster to access resources. You need to add a user for the tenant on FusionInsight Manager and bind the user to the role of the tenant to assign operation permissions to the user.

Prerequisites

You have clarified service requirements and created a tenant.

Procedure

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

Step 2 If you want to add a user to the system, click **Create**.

Figure 8-43 Adding a user

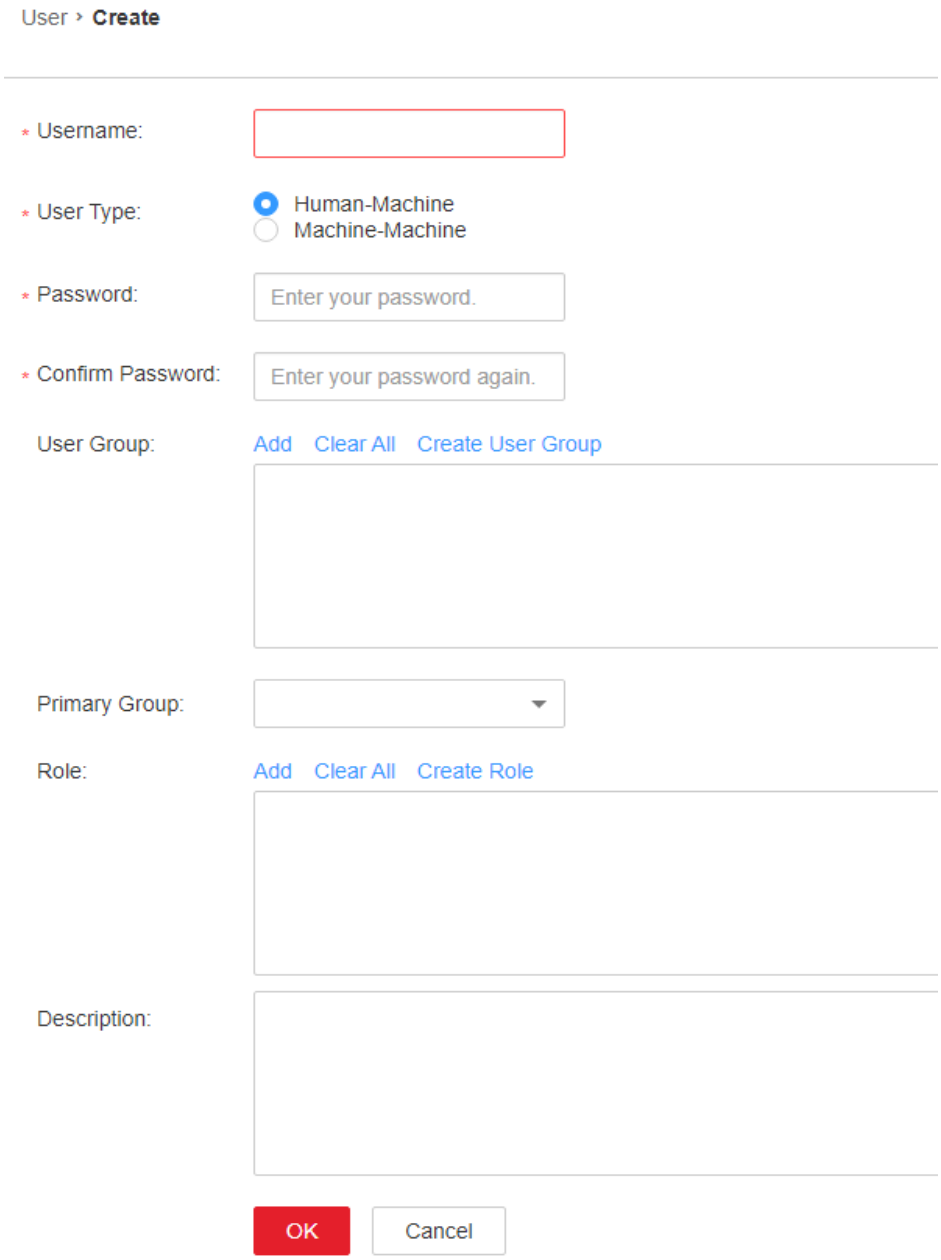

If you want to bind tenant roles to an existing user in the system, locate the row of the user and click **Modify** in the **Operation** column.

Set user attributes according to **Table 8-35**.

Table 8-35 User parameters

Parameter	Description
Username	Indicates the current username. The value contains 3 to 32 characters, including digits, letters, underscores (_), hyphens (-), and spaces. • The username cannot be the same as the OS username of any node in the cluster. Otherwise, the user cannot be used. • A username that differs only in alphabetic case from an existing username is not allowed. For example, if **User1** has been created, you cannot create **user1**. Enter the correct username when using **User1**.
User Type	The options are **Human-Machine** and **Machine-Machine**. • **Human-Machine** user: used for FusionInsight Manager O&M and component client operations. If you select this option, set both **Password** and **Confirm Password** accordingly. • **Machine-Machine** user: used for application development. If you select this option, the password is randomly generated.
Password	This parameter is mandatory if **User Type** is set to **Human-Machine**. The password must contain 8 to 64 characters of at least four types of the following: uppercase letters, lowercase letters, digits, special characters, and spaces. The password cannot be the username or the username spelled backwards.
Confirm Password	Enter the password again.
User Group	In the **User Group** area, click **Add** and select user groups to add the user to the groups. • If roles have been added to the user groups, the user can be granted the permissions of the roles. • For example, add the user to the Hive user group to assign Hive permissions to the user.
Primary Group	Select a group as the primary group for the user to create directories and files. The drop-down list contains all groups selected in **User Group**.

Parameter	Description
Role	Click **Add** to bind a tenant role to the user. **NOTE** • If a user wants to use the resources of tenant **tenant1** and to add or delete sub-tenants for **tenant1**, the user must be bound to both the **Manager_tenant** and **tenant1**_*Cluster ID* roles. • If the tenant has been associated with the HBase service and Ranger authentication is enabled for the cluster, you need to configure the HBase execution permissions on the Ranger page.
Description	Indicates the description of the current user.

Step 3 Click **OK**.

----End

8.6.2.2 Managing Tenants

8.6.2.2.1 Managing Tenant Directories

Scenario

You can manage the HDFS storage directories used by specified tenants based on service requirements on FusionInsight Manager, such as adding tenant directories, changing the quotas for directories and files and for storage space, and deleting directories.

Prerequisites

A tenant with HDFS storage resources has been added.

Viewing a Tenant Directory

Step 1 Log in to FusionInsight Manager and choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 View the **HDFS Storage** table.

- The **File Number Threshold** column provides the quota for files and directories of the tenant directory.
- The **Space Quota** column provides the storage space size of the tenant directory.

----End

Adding a Tenant Directory

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 In the **HDFS Storage** area, click **Create Directory**.

- **Parent Directory**: indicates the storage directory used by the parent tenant of the current tenant.

  NOTE

  This parameter is not displayed if the current tenant is not a sub-tenant.

- Set **Path** to a tenant directory path.

  NOTE

  If the current tenant is not a sub-tenant, the new path is created in the HDFS root directory.

- Set **Quota** to the quota for files and directories.
- **File Number Threshold (%)** is valid only when **Quota** is set. If the ratio of the number of used files to the value of **Quota** exceeds the value of this parameter, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.

  NOTE

  The number of used files is collected every hour. Therefore, the alarm indicating that the ratio of used files exceeds the threshold is delayed.

- Set **Space Quota** to the storage space size of the tenant directory.
- If the ratio of used storage space to the value of **Space Quota** exceeds the **Storage Space Threshold (%)** value, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.

  NOTE

  The used storage space is collected every hour. Therefore, the alarm indicating that the ratio of used storage space exceeds the threshold is delayed.

Step 5 Click **OK**.

----End

Modifying a Tenant Directory

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 In the **HDFS Storage** table, click **Modify** in the **Operation** column of the specified tenant directory.

- Set **Quota** to the quota for files and directories.
- **File Number Threshold (%)** is valid only when **Quota** is set. If the ratio of the number of used files to the value of **Quota** exceeds the value of this

parameter, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.

- Set **Space Quota** to the storage space size of the tenant directory.
- If the ratio of used storage space to the value of **Space Quota** exceeds the **Storage Space Threshold (%)** value, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.

Step 5 Click **OK**.

----End

Deleting a Tenant Directory

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 In the **HDFS Storage** table, click **Delete** in the **Operation** column of the specified tenant directory.

NOTE

The tenant directory that is created by the system during tenant creation cannot be deleted.

Step 5 Click **OK**.

----End

8.6.2.2.2 Restoring Tenant Data

Scenario

Tenant data is stored on FusionInsight Manager and cluster components. When components are recovered from failures or reinstalled, some configuration data of all tenants may become abnormal. In this case, you need to manually restore the configuration data on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager and choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Check the tenant data status.

1. On the **Summary** page, check **Tenant Status**. A green icon indicates that the tenant is available and gray indicates that the tenant is unavailable.
2. Click **Resource** and check the icons on the left of **Yarn** and **HDFS Storage**. A green icon indicates that the resource is available, and gray indicates that the resource is unavailable.
3. Click **Service Associations** and check the **Status** column of the associated services. **Normal** indicates that the component can provide services for the associated tenant. **Not Available** indicates that the component cannot provide services for the tenant.

4. If any of the preceding check items is abnormal, go to **Step 4** to restore tenant data.

Step 4 Click . In the displayed dialog box, enter the password of the current login user and click **OK**.

Step 5 In the **Restore Tenant Resource Data** window, select one or more components to restore data, and click **OK**. The system automatically restores the tenant data.

----End

8.6.2.2.3 Deleting a Tenant

Scenario

You can delete tenants that are no longer used on FusionInsight Manager based on service requirements to release resources occupied by the tenants.

Prerequisites

- A tenant has been added.
- The tenant has no sub-tenants. If the tenant has sub-tenants, delete them; otherwise, the tenant cannot be deleted.
- The role of the tenant is not associated with any user or user group.

Procedure

Step 1 Log in to FusionInsight Manager and choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant and click .

NOTE

- If you want to retain the tenant data, select **Reserve the data of this tenant resource**. Otherwise, the storage space of the tenant will be deleted.

Step 3 In the **Delete Tenant** dialog box, enter **DELETE** in the confirmation text box. Click **OK**.

It takes a few minutes to save the configuration. After the tenant is deleted, the role and storage space of the tenant are also deleted.

NOTE

After the tenant is deleted, the queue of the tenant still exists in Yarn. The queue of the tenant is not displayed on the role management page in Yarn.

----End

8.6.2.3 Managing Resources

8.6.2.3.1 Adding a Resource Pool

Scenario

In a cluster, you can logically group Yarn NodeManagers into Yarn resource pools. Each NodeManager belongs to only one resource pool. You can create a custom

resource pool on FusionInsight Manager and add the hosts that have not been added to any custom resource pools to this resource pool so that specified queues can use the computing resources provided by these hosts.

The system contains a **default** resource pool by default. All NodeManagers that are not added to custom resource pools belong to this resource pool.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Resource Pool**.

Step 3 Click **Add Resource Pool**.

Step 4 Set resource pool attributes.

Figure 8-44 Adding resources

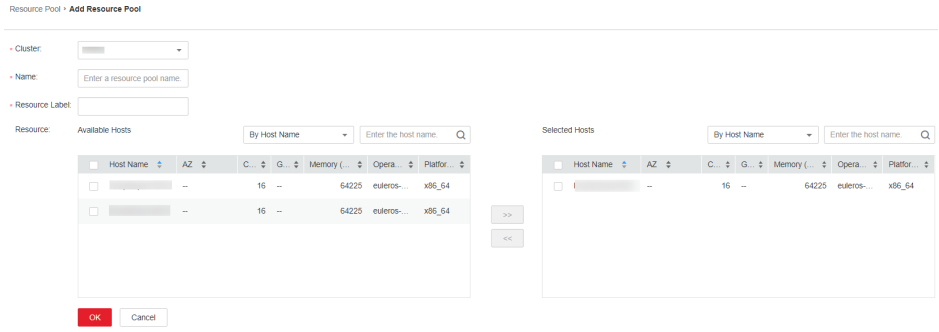

- **Cluster**: Select the cluster to which the resource pool is to be added. (This parameter is unavailable for clusters of MRS 3.3.0 or later.)
- **Name**: Enter the name of the resource pool. The name contains 1 to 50 characters, including digits, letters, and underscores (_), and cannot start with an underscore (_).
- **Resource Label**: Enter the resource label of the resource pool. The value can contain 1 to 50 characters, including digits, letters, underscores (_), and hyphens (-), and must start with a digit or letter.
- **Resource**: In the **Available Hosts** area, select specified hosts and click >> to add the hosts to the **Selected Hosts** area. Only hosts in the cluster can be selected. The host list in the resource pool can be left blank.

  NOTE

  You can filter hosts by host name, number of CPU cores, memory, operating system, or platform type based on service requirements.

Step 5 Click **OK**.

After the resource pool is created, you can view its name, members, and mode in the resource pool list. Hosts that are added to the custom resource pool are no longer members of the **default** resource pool.

----End

8.6.2.3.2 Modifying a Resource Pool

Scenario

When hosts in a resource pool need to be adjusted based on service requirements, you can modify members in the resource pool on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Resource Pool**.

Step 3 In the resource pool list, locate the row that contains the specified resource pool, and click **Edit** in the **Operation** column (for MRS 3.3.0 or later, click **Modify**).

Step 4 In the **Resource** area, modify hosts.

- Adding hosts: Select desired hosts in **Available Hosts** and click >> to add them to the resource pool.
- Deleting hosts: Select desired hosts in **Selected Hosts** and click << to remove them from the resource pool. The host list in the resource pool can be left blank.

Step 5 Click **OK**.

----End

8.6.2.3.3 Deleting a Resource Pool

Scenario

If a resource pool is no longer used based on service requirements, you can delete it on FusionInsight Manager.

Prerequisites

- Any queue in the cluster does not use the resource pool to be deleted as the default resource pool. Before deleting the resource pool, cancel the default resource pool. For details, see **Modifying Queue Resources**.
- Resource distribution policies of all queues have been cleared from the resource pool to be deleted. For details, see **Clearing Queue Configurations**.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Resource Pool**.

Step 3 Locate the row that contains the specified resource pool, and click **Delete** in the **Operation** column.

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

----End

8.6.2.3.4 Modifying Queue Resources

Scenario

If resources need to be adjusted based on service requirements, you can scale up or down a queue by modifying the queue configuration of a specified tenant on FusionInsight Manager. Yarn queues are associated with resource pools for resource allocation and scheduling.

Prerequisites

A tenant who uses the Superior scheduler has been added.

Procedure

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 Choose **Dynamic Resource Plan**.

Step 3 Click the **Queue Configurations** tab.

Step 4 In **All tenants resources** area, locate the row that contains the target tenant resource and click **Modify** in the **Operation** column.

NOTE

- You can also access the **Modify Queue Configuration** page as follows: In the tenant list on the **Tenant Resources Management** page, click the target tenant, click the **Resource** tab, and click next to **Queue Configurations (*Queue name*)**.
- A queue can be bound to only one non-default resource pool.
- For parameters such as **Max Allocated vCores**, **Max Allocated Memory(MB)**, **Max Running Apps**, **Max Running Apps per User**, and **Max Pending Apps**, if the value of a sub-tenant is **-1**, the value of the parent tenant can be set to a specific limit. If the parent tenant value is a specific limit, the sub-tenant value can be set to **-1**.
- **Max Allocated vCores** and **Max Allocated Memory(MB)** must be both changed to values other than **-1**.
- For queues with cross-resource-pool scheduling enabled, existing resource pools cannot be deleted during job running. Otherwise, running jobs may be continuously blocked because they cannot obtain resources. Similarly, if a new resource pool is configured for a queue during job running, the queue in the running state may not immediately use the resources in the new resource pool. The new resources are available only for jobs submitted after modification.

Table 8-36 Queue configuration parameters

Parameter	Description
Max Master Shares(%)	Indicates the maximum percentage of resources occupied by all ApplicationMasters in the current queue.
Max Allocated vCores	Indicates the maximum number of cores that can be allocated to a single Yarn container in the current queue. The default value is **-1**, indicating that the number of cores is not limited within the value range.

Parameter	Description
Max Allocated Memory(MB)	Indicates the maximum memory that can be allocated to a single Yarn container in the current queue. The default value is **-1**, indicating that the memory is not limited within the value range.
Max Running Apps	Indicates the maximum number of tasks that can be executed at the same time in the current queue. The default value is **-1**, indicating that the number is not limited within the value range (the meaning is the same if the value is empty). Value **0** indicates that tasks cannot be executed. The value ranges from **-1** to **2147483647**.
Max Running Apps per User	Indicates the maximum number of tasks that can be executed by each user in the current queue at the same time. The default value is **-1**, indicating that the number is not limited within the value range (the meaning is the same if the value is empty). Value **0** indicates that tasks cannot be executed. The value ranges from **-1** to **2147483647**.
Max Pending Apps	Indicates the maximum number of tasks that can be suspended at the same time in the current queue. The default value is **-1**, indicating that the number is not limited within the value range (the meaning is the same if the value is empty). Value **0** indicates that tasks cannot be suspended. The value ranges from **-1** to **2147483647**.
Resource Allocation Rule	Indicates the rule for allocating resources to different tasks of a user. The rule can be **FIFO** or **FAIR**. If a user submits multiple tasks in the current queue and the rule is **FIFO**, the tasks are executed one by one in sequential order; If the rule is **FAIR**, resources are evenly allocated to all tasks.
Default Resource Label	Indicates that tasks are executed on a node with a specified resource label.
Cross-Pool Scheduling	Indicates whether containers in the current queue support cross-pool scheduling. This parameter is available for clusters of MRS 3.3.0 or later. This function cannot be enabled for the default queue.
Cross-Pool AM Scheduling	Indicates whether ApplicationMasters in the current queue support cross-pool scheduling. This parameter is available for clusters of MRS 3.3.0 or later. This function cannot be enabled for the default queue.

Parameter	Description
Active	• **ACTIVE**: indicates that the current queue can receive and execute tasks. • **INACTIVE**: indicates that the current queue can receive but cannot execute tasks. Tasks submitted to the queue are suspended.
Open	• **OPEN**: indicates that the current queue is opened. • **CLOSED**: indicates that the current queue is closed. Tasks submitted to the queue are rejected.
Migrate Queue Upon Fault	If cross-AZ HA is enabled for a cluster and an AZ is faulty, set **Migrate Queue Upon Fault** to **TRUE** to migrate running queues of the tenant to other AZs.

Step 5 Click **OK**.

----End

8.6.2.3.5 Configuring the Queue Capacity Policy of a Resource Pool

Scenario

After a resource pool is added, you can configure the capacity policy of available resources for Yarn queues so that jobs in the queues can be properly executed in the resource pool.

This section describes how to configure the queue policy on FusionInsight Manager. Tenant queues equipped with the Superior scheduler can use resources in different resource pools.

Prerequisites

- You have logged in to FusionInsight Manager.
- A resource pool has been added.
- The target queue is not associated with the resource pools of other queues except the default resource pool.

Procedure

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 Choose **Dynamic Resource Plan**.

Step 3 Click the **Resource Distribution Policy** tab.

Step 4 Select the name of the target cluster from **Cluster** and select a resource pool from **Resource Pool**.

Step 5 Locate the row that contains the target queue in the **Resource Allocation** area, and click **Modify** in the **Operation** column.

Step 6 On the **Resource Configuration Policy** tab of the **Modify Resource Allocation** window, set the resource configuration policy of the queue in the resource pool.

Figure 8-45 Resource configuration policy

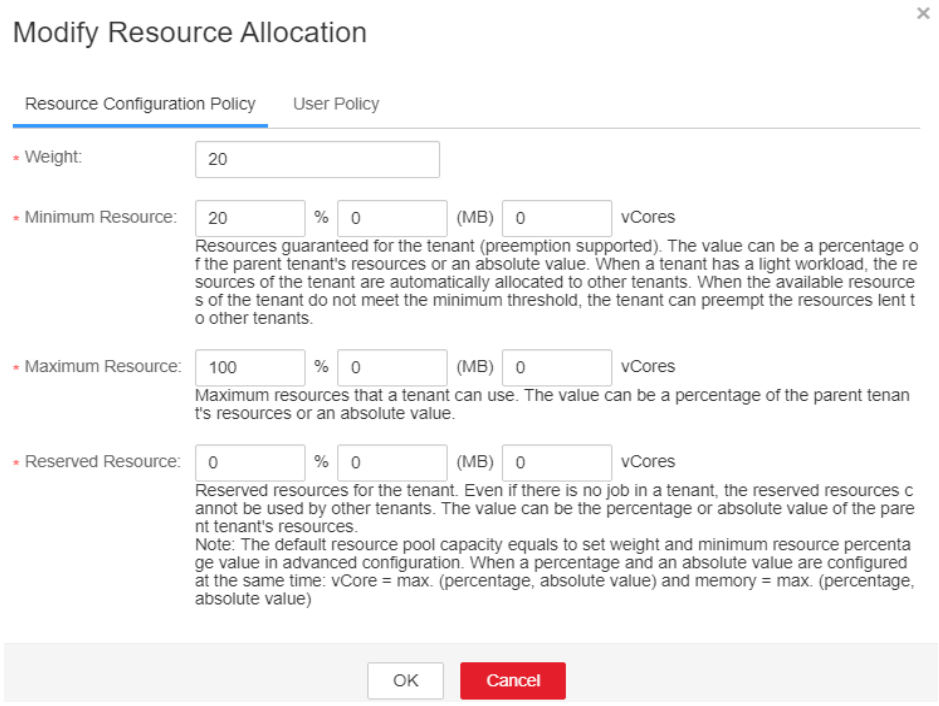

- **Weight**: The task queue with a larger weight preempts resources first when resources are insufficient. Its initial value is the same as the minimum resource percentage.
- **Minimum Resource**: indicates the minimum resources that a tenant can obtain.
- **Maximum Resource**: indicates the maximum resources that a tenant can obtain.
- **Reserved Resource**: indicates the resources that are reserved for the tenant's queues and cannot be lent to other tenants' queues.

Step 7 Click the **User Policy** tab in the **Modify Resource Allocation** window and set the user policy.

Figure 8-46 User policy

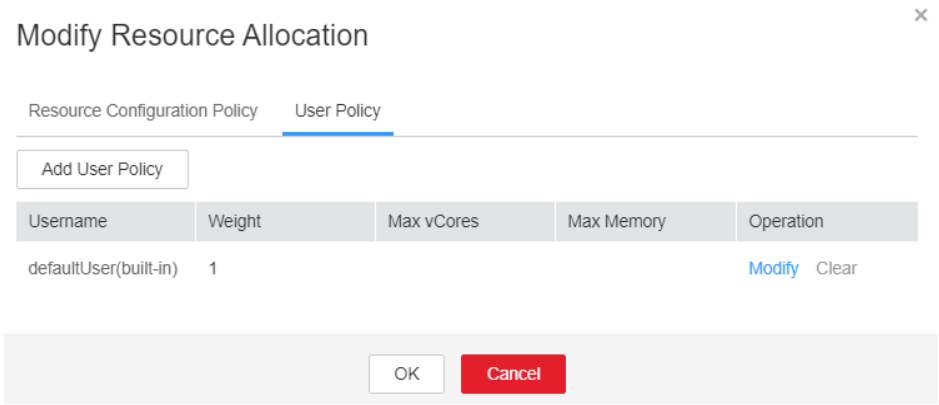

NOTE

defaultUser(built-in) indicates that the policy specified for **defaultUser** is used if a user does not have a policy. The default policy cannot be deleted.

- Click **Add User Policy** to add a user policy.

Figure 8-47 Adding a user policy

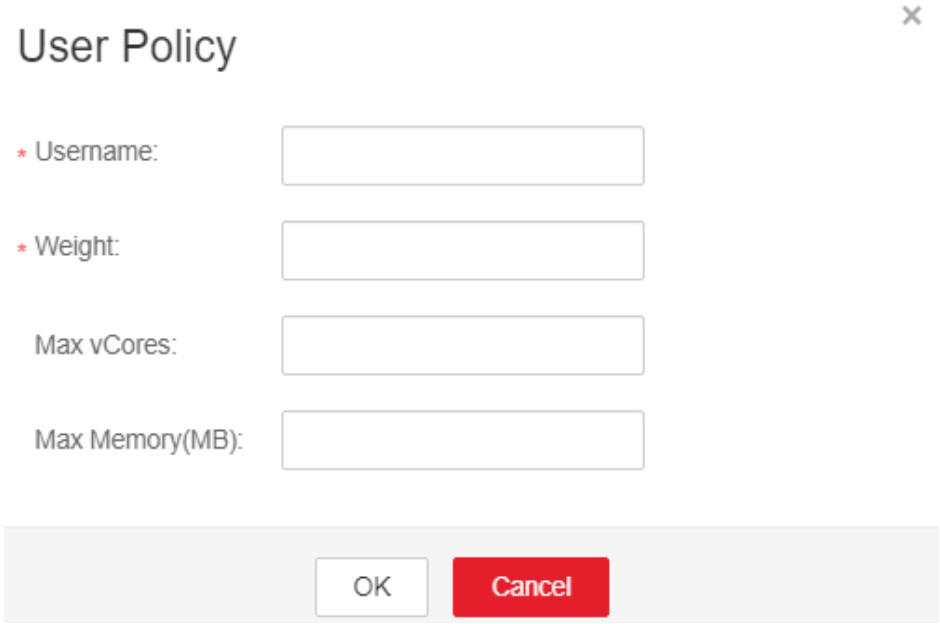

  - **Username**: indicates the name of a user.
  - **Weight**: The task queue with a larger weight preempts resources first when resources are insufficient.
  - **Max vCores**: indicates the maximum number of virtual cores that the user can obtain.
  - **Max Memory(MB)**: indicates the maximum memory that the user can obtain.
- Click **Modify** in the **Operation** column to modify an existing user policy.
- Click **Clear** in the **Operation** column to delete an existing user policy.

Step 8 Click **OK**.

----End

8.6.2.3.6 Clearing Queue Configurations

Scenario

You can clear the configurations of a queue on FusionInsight MRS Manager when the queue does not need resources of a resource pool or the resource pool needs to be disassociated from the queue. Clearing queue configurations cancels the resource capacity policy of the queue in the resource pool.

Prerequisites

You have changed the default resource pool of the queue to another one. If a queue is to be disassociated from a resource pool, this resource pool cannot serve as the default resource pool of the queue. For details, see **Modifying Queue Resources**.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Dynamic Resource Plan**.

Step 3 In **Resource Pools**, select the target resource pool.

Step 4 Locate the row that contains the target resource name in the **Resource Allocation** area, and click **Clear** in the **Operation** column.

Step 5 In the displayed dialog box, click **OK** to clear the queue configurations from the current resource pool.

----End

8.6.2.4 Managing Global User Policies

Scenario

If a tenant uses a Superior scheduler, you can configure the global policy for users to use the resource scheduler, including:

- Maximum running apps
- Maximum pending apps
- Default queue

Procedure

- Add a policy.
  a. On FusionInsight Manager, choose **Tenant Resources**.
  b. Choose **Dynamic Resource Plan**.
  c. Click the **Global User Policy** tab.

     NOTE

     defaults(default setting) indicates that the policy specified for **defaults** is used if a user does not have a global policy. The default policy cannot be deleted.
  d. Click **Create Global User Policy**. In the displayed dialog box, set the following parameters:

Figure 8-48 Creating a global user policy

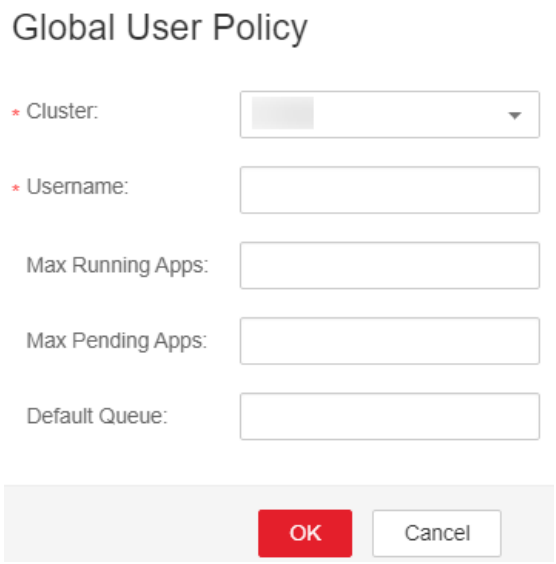

  - **Cluster**: Select the target cluster. (This parameter is unavailable for clusters of MRS 3.3.0 or later.)
  - **Username**: indicates the user for whom resource scheduling is controlled. Enter an existing username in the current cluster.
  - **Max Running Apps**: indicates the maximum number of tasks that the user can run in the current cluster.
  - **Max Pending Apps**: indicates the maximum number of tasks that the user can suspend in the current cluster.
  - **Default Queue**: indicates the queue of the user. Enter the name of an existing queue in the current cluster.

- Modify a policy.
  a. On FusionInsight Manager, choose **Tenant Resources**.
  b. Choose **Dynamic Resource Plan**.
  c. Click the **Global User Policy** tab.
  d. In the row that contains the desired user policy, click **Modify** in the **Operation** column.
  e. In the displayed dialog box, modify parameters and click **OK**.
- Delete a policy.
  a. On FusionInsight Manager, choose **Tenant Resources**.
  b. Choose **Dynamic Resource Plan**.
  c. Click the **Global User Policy** tab.
  d. In the row that contains the desired user policy, click **Delete** in the **Operation** column.
  In the displayed dialog box, click **OK**.

8.6.3 Using the Capacity Scheduler

8.6.3.1 Creating Tenants

8.6.3.1.1 Adding a Tenant

Scenario

You can create tenants on FusionInsight Manager based on the resource consumption and isolation planning and requirements of services.

Prerequisites

- A tenant name has been planned 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.
- Resources to be allocated to the current tenant have been planned to ensure that the sum of resources of direct sub-tenants at each level does not exceed the resources of the current tenant.

Procedure

Step 1 Log in to FusionInsight Manager and choose **Tenant Resources**.

Step 2 Click . On the page that is displayed, configure tenant attributes according to **Table 8-37**.

Table 8-37 Tenant parameters

Parameter	Description
Cluster	Indicates the cluster for which you want to create a tenant. (This parameter is unavailable for clusters of MRS 3.3.0 or later.)
Name	• Indicates the name of the current tenant. The value consists of 3 to 50 characters, including digits, letters, and underscores (_). • Plan a tenant name based on service requirements. The name cannot be the same as that of a role, HDFS directory, or Yarn queue that exists in the current cluster.

Parameter	Description
Tenant Resource Type	Specifies whether the tenant is a leaf tenant. ● When **Leaf Tenant Resource** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. ● When **Non-leaf Tenant Resource** is selected, the current tenant is not a leaf tenant and sub-tenants can be added to the current tenant. **NOTE** MRS 3.2.0 or later: If you select **ClickHouse** for **Service**, this parameter can only be set to **Leaf Tenant**.
Computing Resource	Specifies the dynamic computing resources for the current tenant. ● When **Yarn** is selected, the system automatically creates a queue in Yarn and the queue is named the same as the tenant name. – A leaf tenant can directly submit jobs to the queue. – A non-leaf tenant cannot directly submit jobs to the queue. However, Yarn adds an extra queue (hidden) named **default** for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue. ● If **Yarn** is not selected, the system does not automatically create a queue.
Configuration Mode	Indicates the configuration mode of computing resource parameters. ● If you select **Basic**, you only need to set **Default Resource Pool Capacity (%)**. ● If you select **Advanced**, you can manually configure the resource allocation weight and the minimum, maximum, and reserved resources of the tenant.
Default Resource Pool Capacity (%)	Indicates the percentage of computing resources used by the current tenant in the default resource pool. The value ranges from **0** to **100%**.
Weight	Indicates the resource allocation weight. The value ranges from **0** to **100**.

Parameter	Description
Minimum Resource	Indicates the resources guaranteed for the tenant (preemption supported). The value can be a percentage or an absolute value of the parent tenant's resources. When a tenant has a light workload, the resources of the tenant are automatically allocated to other tenants. When the available tenant resources are less than the value of **Minimum Resource**, the tenant can preempt the resources that have been lent to other tenants.
Maximum Resource	Indicates the maximum resources that can be used by the tenant. The tenant cannot obtain more resources than the value configured. The value can be a percentage or an absolute value of the parent tenant's resources.
Reserved Resource	Indicates the resources reserved for the tenant. The reserved resources cannot be used by other tenants even if no job is running in the current tenant resources. The value can be a percentage or an absolute value of the parent tenant's resources.
Storage Resource	Specifies storage resources for the current tenant. ● When **HDFS** is selected, the system automatically allocates storage resources. ● When **HDFS** is not selected, the system does not automatically allocate storage resources.
Quota	Indicates the quota for files and directories.
Space Quota	Indicates the quota for the HDFS storage space used by the current tenant. ● If the unit is set to **MB**, the value ranges from **1** to **8796093022208**. If the unit is set to **GB**, the value ranges from **1** to **8589934592**. ● This parameter indicates the maximum HDFS storage space that can be used by the tenant, but not the actual space used. ● 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.
Storage Path	Indicates an HDFS directory for tenant resource data. ● The system automatically creates a folder named after the tenant name in the **/tenant** directory by default. For example, the default HDFS storage directory for tenant **ta1** is **/tenant/ta1**. ● When a tenant is created for the first time, the system creates the **/tenant** directory in the HDFS root directory. The storage path is customizable.

Parameter	Description
Service	Specifies whether to associate resources of other services. For details, see **Step 4**.
Description	Indicates the description of the current tenant.

NOTE

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

- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System** > **Permission** > **Role**. The role name is in the format of *Tenant name_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and bind the user to the role of the tenant. For details, see **Adding a User and Binding the User to a Tenant Role**.
- During the tenant creation, the system automatically creates a Yarn queue named after the tenant. If the queue name already exists, the new queue is named **Tenant name-***N*. *N* indicates a natural number starting from **1**. When a same name exists, the value *N* increases automatically to differentiate the queue from others. For example, **saletenant**, **saletenant-1**, and **saletenant-2**.

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.

- Set **Service** to **HBase** and **Association Type** to **Exclusive** or **Shared**.

  NOTE

  - **Exclusive** indicates that the service resources are used by the tenant exclusively and cannot be associated with other tenants.
  - **Shared** indicates that the service resources can be shared with other tenants.

- MRS 3.2.0 or later: Select **ClickHouse** for **Service**.
  - **Association Type**: When **Service** is set to **ClickHouse**, **Association Type** can only be set to **Shared**.
  - **Associate Logical Cluster**: If the logical cluster function is not enabled for ClickHouse, **default_cluster** is selected by default. If the function is enabled, select the logical cluster to which you want to associate.
  - **CPU Priority**: The CPU priority ranges from –20 to 19. This value is associated with the NICE value of the OS. A smaller value indicates a higher CPU priority.
  - **Memory**: The maximum value of this parameter is **100**, in percentage. For example, if this parameter is set to **80**, the total memory that can be used by the current tenant is calculated as follows: Available memory x 80%.

NOTE

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

1. Click **OK**.

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

----End

8.6.3.1.2 Adding a Sub-Tenant

Scenario

You can create sub-tenants on FusionInsight Manager and allocate resources of the current tenant to the sub-tenants based on the resource consumption and isolation planning and requirements of services.

Prerequisites

- A parent non-leaf tenant has been added.
- A tenant name has been planned 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.
- Resources to be allocated to the current tenant have been planned to ensure that the sum of resources of direct sub-tenants at each level does not exceed the resources of the current tenant.

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 attributes for the sub-tenant according to **Table 8-38**.

Table 8-38 Sub-tenant parameters

Parameter	Description
Cluster	Indicates the cluster to which the parent tenant belongs.
Parent Tenant Resource	Indicates the name of the parent tenant.

Parameter	Description
Name	• Indicates the name of the current tenant. The value consists of 3 to 50 characters, including digits, letters, and underscores (_). • 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.
Tenant Type	Specifies whether the tenant is a leaf tenant. • When **Leaf Tenant** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. • When **Non-leaf Tenant** 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.
Computing Resource	Specifies the dynamic computing resources for the current tenant. • When **Yarn** is selected, the system automatically creates a queue in Yarn and the queue is named the same as the sub-tenant name. – A leaf tenant can directly submit jobs to the queue. – A non-leaf tenant cannot directly submit jobs to the queue. However, Yarn adds an extra queue (hidden) named **default** for the non-leaf tenant to record the remaining resource capacity of the tenant. Actual jobs do not run in this queue. • If **Yarn** is not selected, the system does not automatically create a queue.
Default Resource Pool Capacity (%)	Indicates the 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 (%)	Indicates the maximum percentage of computing resources used by the current tenant. The base value is the total resources of the parent tenant.
Storage Resource	Specifies storage resources for the current tenant. • When **HDFS** is selected, the system automatically creates a folder named after the sub-tenant in the HDFS parent tenant directory. • When **HDFS** is not selected, the system does not automatically allocate storage resources.
Quota	Indicates the quota for files and directories.

Parameter	Description
Space Quota	Indicates the quota for the HDFS storage space used by the current tenant. • If the unit is set to **MB**, the value ranges from **1** to **8796093022208**. If the unit is set to **GB**, the value ranges from **1** to **8589934592**. • This parameter indicates the maximum HDFS storage space that can be used by the tenant, but not the actual space used. • 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. • 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.
Storage Path	Indicates the HDFS storage directory for the tenant. • 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 **ta1s** and the parent directory is **/tenant/ta1**, the storage path for the sub-tenant is then **/tenant/ta1/ta1s**. • The storage path is customizable in the parent directory.
Description	Indicates the description of the current tenant.

NOTE

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

- The new role has permissions on the computing and storage resources. This role and its permissions are automatically controlled by the system and cannot be manually managed by choosing **System** > **Permission** > **Role**. The role name is in the format of *Tenant name_Cluster ID*. The ID of the first cluster is not displayed by default.
- When using this tenant, create a system user and bind the user to the role of the tenant. For details, see **Adding a User and Binding the User to a Tenant Role**.
- 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

8.6.3.1.3 Adding a User and Binding the User to a Tenant Role

Scenario

A newly created tenant cannot directly log in to the cluster to access resources. You need to add a user for the tenant on FusionInsight Manager and bind the user to the role of the tenant to assign operation permissions to the user.

Prerequisites

You have clarified service requirements and created a tenant.

Procedure

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

Step 2 If you want to add a user to the system, click **Create**.

If you want to bind tenant roles to an existing user in the system, locate the row of the user and click **Modify** in the **Operation** column.

Set user attributes according to **Table 8-39**.

Table 8-39 User parameters

Parameter	Description
Username	Specifies the current user name. The value can contain 3 to 32 characters, including digits, letters, underscores (_), hyphens (-), and spaces. • The username cannot be the same as the OS username of any node in the cluster. Otherwise, the user cannot be used. • A username that differs only in alphabetic case from an existing username is not allowed. For example, if **User1** has been created, you cannot create **user1**. Enter the correct username when using **User1**.
User Type	The options are **Human-Machine** and **Machine-Machine**. • **Human-Machine** user: used for FusionInsight Manager O&M and component client operations. If you select this option, set both **Password** and **Confirm Password** accordingly. • **Machine-Machine** user: used for application development. If you select this option, the password is randomly generated.
Password	This parameter is mandatory if **User Type** is set to **Human-Machine**. The password must contain 8 to 64 characters of at least four types of the following: uppercase letters, lowercase letters, digits, special characters, and spaces. The password cannot be the username or the username spelled backwards.
Confirm Password	Enter the password again.
User Group	In the **User Group** area, click **Add** and select user groups to add the user to the groups. • If roles have been added to the user groups, the user can be granted the permissions of the roles. • For example, add the user to the Hive user group to assign Hive permissions to the user.
Primary Group	Select a group as the primary group for the user to create directories and files. The drop-down list contains all groups selected in **User Group**.
Role	Click **Add** to bind a tenant role to the user. **NOTE** If a user wants to use the resources of tenant **tenant1** and to add or delete sub-tenants for **tenant1**, the user must be bound to both the **Manager_tenant** and **tenant1**_*Cluster ID* roles.

Parameter	Description
Description	Indicates the description of the current user.

Step 3 Click **OK**.

----End

8.6.3.2 Managing Tenants

8.6.3.2.1 Managing Tenant Directories

Scenario

You can manage the HDFS storage directories used by specified tenants based on service requirements on FusionInsight Manager, such as adding tenant directories, changing the quotas for directories and files and for storage space, and deleting directories.

Prerequisites

A tenant with HDFS storage resources has been added.

Viewing a Tenant Directory

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 View the **HDFS Storage** table.

- The **File Number Threshold** column provides the quota for files and directories of the tenant directory.
- The **Space Quota** column provides the storage space size of the tenant directory.

----End

Adding a Tenant Directory

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 In the **HDFS Storage** area, click **Create Directory**.

Figure 8-49 Creating a directory

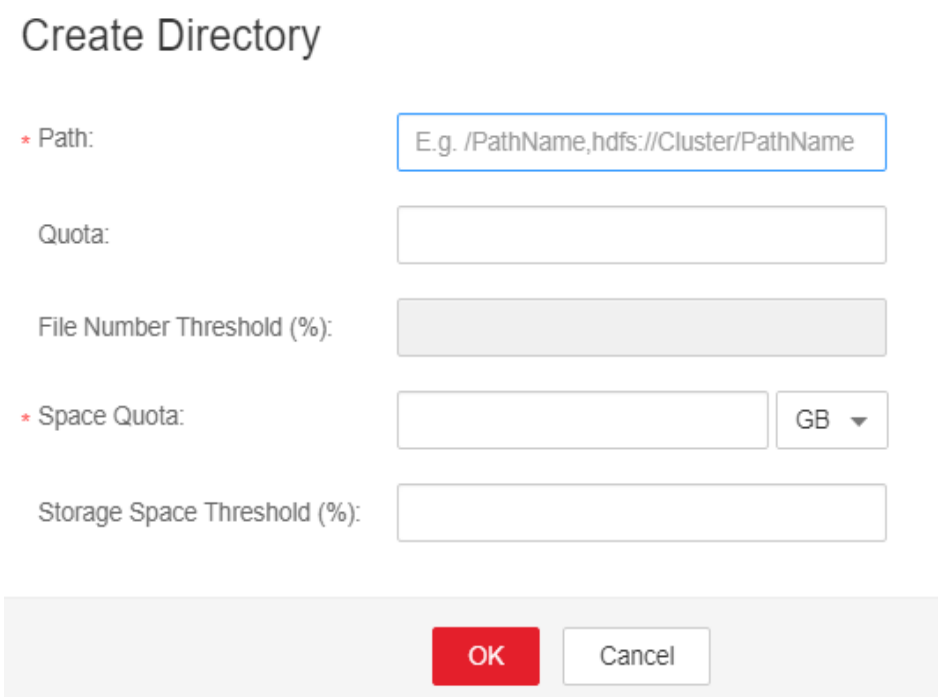

- **Parent Directory**: indicates the storage directory used by the parent tenant of the current tenant.

  NOTE

  This parameter is not displayed if the current tenant is not a sub-tenant.

- Set **Path** to a tenant directory path.

  NOTE

  If the current tenant is not a sub-tenant, the new path is created in the HDFS root directory.

- Set **Quota** to the quota for files and directories.
- **File Number Threshold (%)** is valid only when **Quota** is set. If the ratio of the number of used files to the value of **Quota** exceeds the value of this parameter, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.

  NOTE

  The number of used files is collected every hour. Therefore, the alarm indicating that the ratio of used files exceeds the threshold is delayed.

- Set **Space Quota** to the storage space size of the tenant directory.
- If the ratio of used storage space to the value of **Space Quota** exceeds the **Storage Space Threshold (%)** value, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.

  NOTE

  The used storage space is collected every hour. Therefore, the alarm indicating that the ratio of used storage space exceeds the threshold is delayed.

Step 5 Click **OK**.

----End

Modifying a Tenant Directory

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 In the **HDFS Storage** table, click **Modify** in the **Operation** column of the specified tenant directory.

- Set **Quota** to the quota for files and directories.
- **File Number Threshold (%)** is valid only when **Quota** is set. If the ratio of the number of used files to the value of **Quota** exceeds the value of this parameter, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.
- Set **Space Quota** to the storage space size of the tenant directory.
- If the ratio of used storage space to the value of **Space Quota** exceeds the **Storage Space Threshold (%)** value, an alarm is generated. If this parameter is not specified, no alarm is reported in this scenario.

Step 5 Click **OK**.

----End

Deleting a Tenant Directory

Step 1 On FusionInsight Manager, choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Click the **Resource** tab.

Step 4 In the **HDFS Storage** table, click **Delete** in the **Operation** column of the specified tenant directory.

NOTE

The tenant directory that is created by the system during tenant creation cannot be deleted.

Step 5 Click **OK**.

----End

8.6.3.2.2 Restoring Tenant Data

Scenario

Tenant data is stored on FusionInsight Manager and cluster components. When components are recovered from failures or reinstalled, some configuration data of all tenants may become abnormal. In this case, you need to manually restore the configuration data on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager and choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant.

Step 3 Check the tenant data status.

1. On the **Summary** page, check **Tenant Status**. A green icon indicates that the tenant is available and gray indicates that the tenant is unavailable.
2. Click **Resource** and check the icons on the left of **Yarn** and **HDFS Storage**. A green icon indicates that the resource is available, and gray indicates that the resource is unavailable.
3. Click **Service Associations** and check the **Status** column of the associated services. **Normal** indicates that the component can provide services for the associated tenant. **Not Available** indicates that the component cannot provide services for the tenant.
4. If any of the preceding check items is abnormal, go to **Step 4** to restore tenant data.

Step 4 Click . In the displayed dialog box, enter the password of the current login user and click **OK**.

Step 5 In the **Restore Tenant Resource Data** window, select one or more components to restore data, and click **OK**. The system automatically restores the tenant data.

----End

8.6.3.2.3 Deleting a Tenant

Scenario

You can delete tenants that are no longer used on FusionInsight Manager based on service requirements to release resources occupied by the tenants.

Prerequisites

- A tenant has been added.
- The tenant has no sub-tenants. If the tenant has sub-tenants, delete them; otherwise, the tenant cannot be deleted.
- The role of the tenant is not associated with any user or user group.

Procedure

Step 1 Log in to FusionInsight Manager and choose **Tenant Resources**.

Step 2 In the tenant list on the left, click the target tenant and click .

NOTE

- If you want to retain the tenant data, select **Reserve the data of this tenant resource**. Otherwise, the storage space of the tenant will be deleted.
- To delete a tenant without retaining the tenant data as a user who does not belong to the supergroup, you should first log in to the HDFS client as a user who belongs to the supergroup and then manually clear the storage space of that tenant to avoid residual data.

Step 3 In the **Delete Tenant** dialog box, enter **DELETE** in the confirmation text box. Click **OK**.

It takes a few minutes to save the configuration. After the tenant is deleted, the role and storage space of the tenant are also deleted.

NOTE

After the tenant is deleted, the queue of the tenant still exists in Yarn. The queue of the tenant is not displayed on the role management page in Yarn.

----End

8.6.3.2.4 Clearing Non-associated Queues of a Tenant

Scenario

If Yarn uses the Capacity scheduler, deleting a tenant only sets the queue capacity of the tenant to **0** and the tenant status to **STOPPED** but does not clear the queues of the tenant in Yarn. Limited by the Yarn mechanism, queues cannot be dynamically deleted. You can run commands to manually delete residual queues.

Impact on the System

- During the script execution, the Controller service is restarted, Yarn configurations are synchronized, and the active and standby ResourceManagers are restarted.
- FusionInsight Manager becomes inaccessible during the restart of the Controller service.
- After the active and standby ResourceManagers are restarted, an alarm is generated indicating that Yarn and components that depend on Yarn are temporarily unavailable.

Prerequisites

Queues of a deleted tenant still exist.

Procedure

Step 1 Check that queues of the deleted tenant still exist.

1. On FusionInsight Manager, choose **Cluster**, click the name of the target cluster, and choose **Services** > **Yarn**. Click the link of the active ResourceManager in **ResourceManager WebUI** to go to the ResourceManager web UI.
2. Click **Scheduler** in the navigation tree on the left. In the right pane, you can view that queues of the tenant still exist in the **STOPPED** state and their **Configured Capacity** is **0**.

Step 2 Log in to the active management node as user **omm**.

Step 3 Switch the directory and execute the **cleanQueuesAndRestartRM.sh** script.

cd ${BIGDATA_HOME}/om-server/om/sbin

./cleanQueuesAndRestartRM.sh -c *Cluster ID*

NOTE

You can choose **Cluster**, click the cluster name, and choose **Cluster Properties** on FusionInsight Manager to view the cluster ID.

During the script execution, you need to enter **yes** and the password.

```
Running the script will restart Controller and restart ResourceManager.
Are you sure you want to continue connecting (yes/no)?yes
Please input admin password:
Begin to backup queues ...
...
```

Step 4 After the script is executed successfully, log in to FusionInsight Manager, choose **Cluster**, click the cluster name, and choose **Services** > **Yarn**. Click the link of the active ResourceManager in **ResourceManager WebUI** to go to the ResourceManager web UI.

Step 5 Click **Scheduler** in the navigation tree on the left. In the right pane, you can view that queues of the tenant have been cleared.

----End

8.6.3.3 Managing Resources

8.6.3.3.1 Adding a Resource Pool

Scenario

In a cluster, you can logically group Yarn NodeManagers into Yarn resource pools. Each NodeManager belongs to only one resource pool. You can create a custom resource pool on FusionInsight Manager and add the hosts that have not been added to any custom resource pools to this resource pool so that specified queues can use the computing resources provided by these hosts.

The system contains a **default** resource pool by default. All NodeManagers that are not added to custom resource pools belong to this resource pool.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Resource Pool**.

Step 3 Click **Add Resource Pool**.

Step 4 Set resource pool attributes.

- **Cluster**: Select the cluster to which the resource pool is to be added. (This parameter is unavailable for clusters of MRS 3.3.0 or later.)
- **Name**: Enter the name of the resource pool. The name contains 1 to 50 characters, including digits, letters, and underscores (_), and cannot start with an underscore (_).
- **Resource Label**: Enter the resource label of the resource pool. The value can contain 1 to 50 characters, including digits, letters, underscores (_), and hyphens (-), and must start with a digit or letter.

- **Resource**: In the **Available Hosts** area, select specified hosts and click >> to add the hosts to the **Selected Hosts** area. Only hosts in the cluster can be selected. The host list in the resource pool can be left blank.

  NOTE

  You can filter hosts by host name, number of CPU cores, memory, operating system, or platform type based on service requirements.

Step 5 Click **OK**.

After the resource pool is created, you can view its name, members, and mode in the resource pool list. Hosts that are added to the custom resource pool are no longer members of the **default** resource pool.

----End

8.6.3.3.2 Modifying a Resource Pool

Scenario

When hosts in a resource pool need to be adjusted based on service requirements, you can modify members in the resource pool on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Resource Pool**.

Step 3 In the resource pool list, locate the row that contains the specified resource pool, and click **Edit** in the **Operation** column (for MRS 3.3.0 or later, click **Modify**).

Step 4 In the **Resource** area, modify hosts.

- Adding hosts: Select desired hosts in **Available Hosts** and click >> to add them to the resource pool.
- Deleting hosts: Select desired hosts in **Selected Hosts** and click << to remove them from the resource pool. The host list in the resource pool can be left blank.

Step 5 Click **OK**.

----End

8.6.3.3.3 Deleting a Resource Pool

Scenario

If a resource pool is no longer used based on service requirements, you can delete it on FusionInsight Manager.

Prerequisites

- Any queue in the cluster does not use the resource pool to be deleted as the default resource pool. Before deleting the resource pool, cancel the default resource pool. For details, see **Modifying Queue Resources**.
- Resource distribution policies of all queues have been cleared from the resource pool to be deleted. For details, see **Clearing Queue Configurations**.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Resource Pool**.

Step 3 Locate the row that contains the specified resource pool, and click **Delete** in the **Operation** column.

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

----End

8.6.3.3.4 Modifying Queue Resources

Scenario

If resources need to be adjusted based on service requirements, you can scale up or down a queue by modifying the queue configuration of a specified tenant on FusionInsight Manager. Yarn queues are associated with resource pools for resource allocation and scheduling.

Prerequisites

A tenant who uses the Capacity scheduler has been added.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Dynamic Resource Plan**.

The **Resource Distribution Policy** page is displayed by default.

Step 3 Click the **Queue Configurations** tab.

Step 4 In **All tenants resources** area, locate the row that contains the target tenant resource and click **Modify** in the **Operation** column.

NOTE

- You can also access the **Modify Queue Configuration** page as follows: In the tenant list on the **Tenant Resources Management** page, click the target tenant, click the **Resource** tab, and click next to **Queue Configurations (*Queue name*)**.
- A queue can be bound to only one non-default resource pool. That is, a newly added resource pool can be bound to only one queue to serve as the default resource pool of the queue.

Table 8-40 Queue configuration parameters

Parameter	Description
Tenant Resources Name (Queue)	Indicates the tenant name and queue name.
Maximum Applications	Indicates the maximum number of applications.
Maximum AM Resource Percent	Indicates the maximum percentage of resources that can be used to run the ApplicationMaster in a cluster.
Minimum User Resource Upper-Limit Percent (%)	Indicates the minimum resource guarantee (percentage) of a user. The resources for each user in a queue are limited at any time. If applications of multiple users are running at the same time in a queue, the resource usage of each user fluctuates between the minimum value and the maximum value. The minimum value is determined by the number of running applications, while the maximum value is determined by this parameter. For example, assume that this parameter is set to **25**. If two users submit applications to the queue, each user can use a maximum of 50% resources; if three users submit applications to the queue, each user can use a maximum of 33% resources; if four users submit applications to the queue, each user can use a maximum of 25% resources.
User Resource Upper-Limit Factor	Indicates the limit factor of the maximum user resource usage. The maximum user resource usage percentage can be obtained by multiplying the limit factor with the percentage of the tenant's actual resource usage in the cluster.
Status	Indicates the status of a resource plan. The value can be **Running** or **Stopped**.
Default Resource Pool	Indicates the resource pool used by the queue. The default value is **default**. If you want to change the resource pool, configure the queue capacity first. For details, see **Configuring the Queue Capacity Policy of a Resource Pool**.

Step 5 Click **OK**.

----End

8.6.3.3.5 Configuring the Queue Capacity Policy of a Resource Pool

Scenario

After a resource pool is added, you can configure the capacity policy of available resources for Yarn queues so that jobs in the queues can be properly executed in

the resource pool. A queue can have the queue capacity policy of only one resource pool.

You can view queues and configure queue capacity policies in any resource pool. After the queue policies are configured, Yarn queues are associated with resource pools.

Prerequisites

A queue has been added, that is, a tenant associated with computing resources has been created.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Dynamic Resource Plan**.

The **Resource Distribution Policy** page is displayed by default.

Step 3 Select the name of the target cluster from **Cluster** and select a resource pool from **Resource Pool**.

Step 4 Locate the row that contains the target resource name in the **Resource Allocation** area, and click **Modify** in the **Operation** column.

Step 5 In the **Modify Resource Allocation** window, configure the resource capacity policy of the queue in the resource pool.

- **Capacity (%)**: indicates the percentage of computing resources used by the current tenant.
- **Maximum Capacity (%)**: indicates the maximum percentage of computing resources used by the current tenant.

Step 6 Click **OK**.

NOTE

After the resource capacity values of a queue are deleted and saved, the resource capacity policy of the queue in the resource pool is canceled, indicating that the queue is disassociated from the resource pool. To achieve this, you need to change the default resource pool of the queue to another one. For details, see **Modifying Queue Resources**.

----End

8.6.3.3.6 Clearing Queue Configurations

Scenario

You can clear the configurations of a queue on FusionInsight MRS Manager when the queue does not need resources of a resource pool or the resource pool needs to be disassociated from the queue. Clearing queue configurations cancels the resource capacity policy of the queue in the resource pool.

Prerequisites

You have changed the default resource pool of the queue to another one. If a queue is to be disassociated from a resource pool, this resource pool cannot serve

as the default resource pool of the queue. For details, see **Modifying Queue Resources**.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Tenant Resources** > **Dynamic Resource Plan**.

Step 3 Select the name of the target cluster from **Cluster** and select a resource pool from **Resource Pool**.

Step 4 Locate the row that contains the target resource name in the **Resource Allocation** area, and click **Clear** in the **Operation** column.

Step 5 In the displayed dialog box, click **OK** to clear the queue configurations from the current resource pool.

----End

8.6.4 Switching the Scheduler

Scenario

The newly installed MRS cluster uses the Superior scheduler by default. If the cluster is upgraded from an earlier version, you can switch the YARN scheduler from the Capacity scheduler to the Superior scheduler with a few clicks.

Prerequisites

- The network connectivity of the cluster is proper and secure, and the YARN service status is normal.
- During scheduler switching, tenants cannot be added, deleted, or modified. In addition, services cannot be started or stopped.

Switching Between the Capacity Scheduler and Superior Scheduler (Available for Clusters of MRS 3.3.0 or Later)

This function is only available for clusters of MRS 3.3.0 or later.

Constraints

- This operation is available for only the scenario where a cluster is newly provisioned and the scheduler needs to be switched.
- During the scheduler switchover, do not perform any operation on the cluster. Otherwise, the operation may fail due to database modification.

Impact on the system

- Because the ResourceManager is restarted during scheduler switching, submitting jobs to Yarn will fail at that time.
- After the scheduler is switched, the parameters of the scheduler that takes over the workload are used.

Procedure

Step 1 Log in to FusionInsight Manager. Choose **Cluster** > **Services** > **Yarn** and check whether the Yarn service status is normal. If the service is abnormal, restore the service.

Step 2 Log in to the active management node as user **omm**.

Step 3 Switch the scheduler.

- Run the following command to switch from the Capacity scheduler to the Superior scheduler:

  sh ${BIGDATA_HOME}/om-server/om/sbin/cleanSwitchScheduler.sh 1

  If information similar to the following is displayed, the switch is successful:

  ```
  Will change scheduler type to SUPERIOR
  Start to delete all tenant resource.
  End to delete all tenant resource.
  Start to delete all resource pool.
  End to delete all resource pool.
  ...
  End to switch scheduler by reset.
  ```

- Run the following command to switch from the Capacity scheduler to the Superior scheduler:

  sh ${BIGDATA_HOME}/om-server/om/sbin/cleanSwitchScheduler.sh 0

  If information similar to the following is displayed, the switch is successful:

  ```
  Will change scheduler type to CAPACITY
  Start to delete all tenant resource.
  End to delete all tenant resource.
  Start to delete all resource pool.
  End to delete all resource pool.
  ...
  End to switch scheduler by reset.
  ```

NOTE

You can query the scheduler switching logs on the active management node.

- ${BIGDATA_LOG_HOME}/controller/aos/clean_switch_scheduler.log
- ${BIGDATA_LOG_HOME}/controller/aos/aos.log
- ${BIGDATA_LOG_HOME}/controller/aos/plugin.log

----End

Switching from the Capacity Scheduler to the Superior Scheduler

Impact on the system

- Because the ResourceManager is restarted during scheduler switching, submitting jobs to YARN will fail at that time.
- During scheduler switching, tasks in a job being executed on YARN will continue, but new tasks cannot be started.
- After scheduler switching is complete, jobs executed on YARN may fail, causing service interruptions.
- After scheduler switching is complete, parameters of the Superior scheduler are used for tenant management.
- After scheduler switching is complete, tenant queues whose capacity is 0 in the Capacity scheduler cannot be allocated resources in the Superior scheduler. As a result, jobs submitted to these tenant queues fail to be

executed. Therefore, you are advised not to set the capacity of a tenant queue to 0 in the Capacity scheduler.

- After scheduler switching is complete, you cannot add or delete resource pools, YARN node labels, or tenants during the observation period. If such an operation is performed, the scheduler cannot be rolled back to the Capacity scheduler.

  NOTE

  The recommended observation period for scheduler switching is one week. If resource pools, YARN node labels, or tenants are added or deleted during this period, the observation period ends immediately.

- Rollback may cause the loss of partial or all Yarn job information.

Procedure

Step 1 Modify YARN service parameters and ensure that the YARN service status is normal.

1. Log in to FusionInsight Manager as an administrator.
2. Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Yarn**. Click **Configurations** then **All Configurations**, search for **yarn.resourcemanager.webapp.pagination.enable**, and check whether the value is **true**.
   - If yes, go to **Step 1.3**.
   - If no, set the parameter to **true** and click **Save** to save the configuration. On the **Dashboard** tab page of YARN, choose **More** > **Restart Service**, verify the identity, and click **OK**. After the service is restarted, go to **Step 1.3**.
3. Choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the YARN service status is normal.

Step 2 Log in to the active management node as user **omm**.

Step 3 Switch the scheduler.

The following switching modes are available:

0: converts the Capacity scheduler configurations into the Superior scheduler configurations and then switches the Capacity scheduler to the Superior scheduler.

1: converts the Capacity scheduler configurations into the Superior scheduler configurations only.

2: switches the Capacity scheduler to the Superior scheduler only.

- Mode **0** is recommended if the cluster environment is simple and the number of tenants is less than 20.

  Run the following command:

  sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c *Cluster ID* **-m 0**

  NOTE

  You can choose **Cluster**, click the cluster name, and choose **Cluster Properties** on FusionInsight Manager to view the cluster ID.


```
Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
Start to convert Capacity scheduler configurations to Superior. Please wait...
Convert configurations successfully.
Start to switch the Yarn scheduler to Superior. Please wait...
Switch the Yarn scheduler to Superior successfully.
```

- If the cluster environment or tenant information is complex and you need to retain the queue configurations of the Capacity scheduler on the Superior scheduler, it is recommended that you use mode **1** first to convert the Capacity scheduler configurations, check the converted configurations, and then use mode **2** to switch the Capacity scheduler to the Superior scheduler.

  a. Run the following command to convert the Capacity scheduler configurations into the Superior scheduler configurations:

     sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c *Cluster ID* **-m 1**

     ```
     Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
     Start to convert Capacity scheduler configurations to Superior. Please wait...
     Convert configurations successfully.
     ```

  b. Run the following command to switch the Capacity scheduler to the Superior scheduler:

     sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c *Cluster ID* **-m 2**

     ```
     Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
     Start to switch the Yarn scheduler to Superior. Please wait...
     Switch the Yarn scheduler to Superior successfully.
     ```

- If you do not need to retain the queue configurations of the Capacity scheduler, use mode **2**.

  a. Log in to FusionInsight Manager and delete all tenants except the default tenant.

  b. On FusionInsight Manager, delete all resource pools except the default resource pool.

     Run the following command to switch the Capacity scheduler to the Superior scheduler:

     sh ${BIGDATA_HOME}/om-server/om/sbin/switchScheduler.sh -c *Cluster ID* **-m 2**

     ```
     Start to convert Capacity scheduler to Superior Scheduler, clusterId=1
     Start to switch the Yarn scheduler to Superior. Please wait...
     Switch the Yarn scheduler to Superior successfully.
     ```

NOTE

You can query the scheduler switching logs on the active management node.

- ${BIGDATA_LOG_HOME}/controller/aos/switch_scheduler.log
- ${BIGDATA_LOG_HOME}/controller/aos/aos.log

----End

Rollback Operations

You can manually switch the Superior scheduler back to the Capacity scheduler. However, this operation is only a workaround and is not allowed in most cases.

If the customer has special requirements for switching back to the Capacity scheduler, the following conditions must be met:

- The observation period has not expired.
- No resource pool, YARN node label, or tenant is added or deleted during the observation period.

NOTICE

If resource pools, YARN node labels, or tenants are added or deleted, resource pools or queues may not exist after the Superior scheduler is switched back to the Capacity scheduler. As a result, the Capacity scheduler cannot run properly.

The procedure is as follows:

Step 1 Change the scheduler to the Capacity scheduler and start YARN.

1. Log in to FusionInsight Manager.
2. Go to the **Configurations** page of YARN and modify the parameters listed in Table 8-41.

Table 8-41 Modifying YARN configuration items

Parameter	Description
yarn.resourcemanager.scheduler.class	org.apache.hadoop.yarn.server.resourcemanager.scheduler.capacity.CapacityScheduler
yarn.http.rmwebapp.external.classes	Left empty
hadoop.http.rmwebapp.scheduler.page.classes	Left empty
yarn.resourcemanager.webapp.pagination.enable	false

3. Click **Save** and then click **OK** in the displayed dialog box.
4. Roll restart the YARN service, enter the password, and click **OK**.

Step 2 Log in to the active management node and restart the AOS service.

1. Log in to the active OMS server as user **omm** using PuTTY.
2. Run the following command to disable logout upon timeout:

   TMOUT=0

   NOTE

   After the operations in this section are complete, run the **TMOUT=***Timeout interval* command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.
3. Run the following command to restart the AOS service:

   ${BIGDATA_HOME}/om-server/om/sbin/aos_cmd.sh restart

----End

8.7 System

8.7.1 Configuring Permissions

8.7.1.1 Managing Users

8.7.1.1.1 Creating a User

Scenario

FusionInsight Manager supports a maximum of 50,000 users (including built-in users). By default, only user **admin** has the highest operation permissions of FusionInsight Manager. You need to create users on FusionInsight Manager and assign operation permissions to the users based on service requirements.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 On the **User** page, click **Create**.

Step 4 Set **Username**. The username can contain digits, letters, underscores (_), hyphens (-), and spaces. It is case-insensitive and cannot be the same as any existing username in the system or OS.

Step 5 Set **User Type** to **Human-Machine** or **Machine-Machine**.

- **Human-Machine** user: used for FusionInsight Manager O&M and component client operations. If you select this option, you also need to select the password policy and set **Password** and **Confirm Password**.
- **Machine-Machine** user: used for component application development. If you select this option, the password is randomly generated.

Step 6 In the **User Group** area, click **Add** to add one or more user groups to the list.

NOTE

- If the selected user group has been bound to a role or a permission policy has been configured in Ranger, the user can obtain the corresponding permissions.
- After FusionInsight Manager is installed, some user groups generated by default have special permissions. Select desired user groups based on the descriptions on the UI.
- If existing user groups cannot meet your requirements, click **Create User Group** to create a user group. For details, see **Creating a User Group**.

Step 7 Select a group from the **Primary Group** drop-down list to create directories and files.

The drop-down list contains all groups selected in **User Group**.

NOTE

A user can belong to multiple groups (including the primary group and secondary groups). The primary group is set to facilitate maintenance and comply with the permission mechanism of the Hadoop community. The primary group has the same permission control functionality as other groups.

Step 8 In the **Role** area, click **Add** to bind roles to the user.

NOTE

- Adding a role when you create a user can specify the user permissions.
- If the permissions granted to the user from the user group cannot meet service requirements, you can bind other created roles to the user. You can also click **Create Role** to create a role first. For details, see **Creating a Role**.

  It takes 3 minutes to make role permission assignment to the user take effect. If the permissions obtained from the user group are enough, you do not need to add a role.
- After Ranger authentication is enabled for a component, you need to configure Ranger policies to assign permissions to the user except the permissions of default user group or role.
- If a user is not added to a user group or assigned a role, the user cannot view information or perform operations after logging in to FusionInsight Manager.

Step 9 Enter information in **Description**.

Step 10 Click **OK**.

After a human-machine user is created, you need to change the initial password as prompted after logging in to FusionInsight Manager.

----End

8.7.1.1.2 Modifying User Information

Scenario

You can modify user information on FusionInsight Manager, including the user group, primary group, role permission assignment, and user description.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 Locate the row that contains the target user and click **Modify** in the **Operation** column.

Modify the parameters based on service requirements.

NOTE

It takes three minutes at most for the change of the user group or role permissions to take effect.

MRS 3.1.2 or later:

- Users (except **admin**) cannot modify their own password policies.
- Locked users cannot modify their password policies.
- After the password policy bound to a user is modified, the modification takes effect when the user changes the password next time.
- After the password policy bound to a user is modified, if the remaining password validity period is greater than the password validity period in the new password policy, the password validity period is set to the validity period in the new password policy. If the remaining password validity period is less than the password validity period in the new password policy, the password validity period remains unchanged.

Step 4 Click **OK**.

----End

8.7.1.1.3 Exporting User Information

Scenario

You can export information about all created users on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 Click **Export All** to export all user information at a time.

The exported user information contains the username, creation time, description, user type (**0** indicates a human-machine account, **1** indicates a machine-machine account), primary group, user group list, and roles bound to the user.

Step 4 Set **Save AS** to **TXT** or **CSV**. Click **OK**.

----End

8.7.1.1.4 Locking a User

Scenario

A user may be suspended for a long period of time due to service changes. For security purposes, you can lock such a user.

You can lock a user in using either of the following methods:

- Automatic locking: You can set **Password Retries** in the password policy to automatically lock the user whose login attempts exceed this parameter value. For details, see **Configuring Password Policies**.
- Manual locking: You manually lock a user.

This section describes how to lock a user manually. Machine-machine users cannot be locked.

Impact on the System

A locked user cannot log in to FusionInsight Manager or perform identity authentication in the cluster. A locked user can be used only after being manually unlocked or the lock time expires.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 Locate the row that contains the target user and click **Lock** in the **Operation** column.

Step 4 In the window that is displayed, select **I have read the information and understand the impact**. Click **OK**.

----End

8.7.1.1.5 Unlocking a User

Scenario

You can unlock a user on FusionInsight Manager if the user has been locked because the number of login attempts exceeds the threshold. Only users created on FusionInsight Manager can be unlocked.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 Locate the row that contains the target user and click **Unlock** in the **Operation** column.

Step 4 In the window that is displayed, select **I have read the information and understand the impact**. Click **OK**.

----End

8.7.1.1.6 Deleting a User

Scenario

Based on service requirements, you can delete system users that are no longer used on FusionInsight Manager.

NOTE

- After a user is deleted, the provisioned ticket granting ticket (TGT) is still valid within 24 hours. The user can use the TGT for security authentication and access the system.
- If a new user has the same name as the deleted user, the new user will inherit all owner permissions of the deleted user. You are advised to determine whether to delete the resources owned by the deleted user based on service requirements, for example, files in HDFS.
- The default user **admin** cannot be deleted.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 Locate the row that contains the target user, click **More**, and select **Delete**.

NOTE

To delete users in batches, select the users at a time and click **Delete**.

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

----End

8.7.1.1.7 Changing a User Password

Scenario

For security purposes, the password of a human-machine user must be changed periodically.

If users have the permission to use FusionInsight Manager, they can change their passwords on FusionInsight Manager.

If users do not have the permission to use FusionInsight Manager, they can change their passwords on the client.

Prerequisites

- You have obtained the current password policy.
- The user has installed the client on any node in the cluster and obtained the IP address of the node. The password of the client installation user can be obtained from the administrator.

Changing the Password on FusionInsight Manager

Step 1 Log in to FusionInsight Manager.

Step 2 Move the cursor to the username in the upper right corner of the page.

On the user account drop-down menu, choose **Change Password**.

Figure 8-50 Changing the password

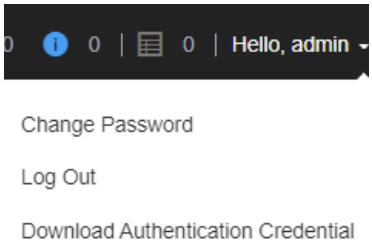

Step 3 On the displayed page, set **Current Password**, **New Password**, and **Confirm Password**, and click **OK**.

The password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~!@#$%^&*()-_=+|[{}];',<.>/\?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be a common easily-cracked password.
- Cannot be the same as the password used in the latest *N* times. *N* indicates the value of **Repetition Rule** configured in **Configuring Password Policies**.

----End

Changing the Password on the Client

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

Step 2 Run the following command to switch to the client directory, for example, **/opt/client**:

cd /opt/client

Step 3 Run the following command to configure environment variables:

source bigdata_env

Step 4 Change the user password. This operation takes effect for all servers.

kpasswd *System username*

For example, if you want to change the password of system user **test1**, run the **kpasswd test1** command.

The password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~!@#$%^&*()-_=+|[{}];',<.>/\?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be a common easily-cracked password.
- Cannot be the same as the password used in the latest *N* times. *N* indicates the value of **Repetition Rule** configured in **Configuring Password Policies**.

NOTE

If an error occurs during the running of the **kpasswd** command, try the following operations:

- Stop the SSH session and start it again.
- Run the **kdestroy** command and then run the **kpasswd** command again.

----End

8.7.1.1.8 Initializing a Password

Scenario

If a user forgets the password or the public account password needs to be changed periodically, you can initialize the password on FusionInsight Manager. After the password is initialized, the system user needs to change the password upon first login.

NOTE

This operation applies only to human-machine users.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 Locate the row that contains the target user, click **More**, and select **Initialize Password**. In the displayed dialog box, enter the password of the current login user and click **OK**. In the **Initialize Password** dialog box, click **OK**.

Step 4 Set **New Password** and **Confirm Password**, and click **OK**.

The password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~!@#$%^&*()-_=+|[{}];',<.>/\?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be a common easily-cracked password.
- Cannot be the same as the password used in the latest *N* times. *N* indicates the value of **Repetition Rule** configured in **Configuring Password Policies**.

----End

8.7.1.1.9 Exporting an Authentication Credential File

Scenario

If a user uses a security mode cluster to develop applications, the keytab file of the user needs to be obtained for security authentication. You can export keytab files on FusionInsight Manager.

NOTE

After a user password is changed, the exported keytab file becomes invalid, and you need to export a keytab file again.

Prerequisites

Before downloading the keytab file of a Human-Machine user, the password of the user must be changed at least once on the Manager portal or a client; otherwise, the downloaded keytab file cannot be used For details, see **Changing a User Password**.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User**.

Step 3 Locate the row that contains the user whose keytab file needs to be exported, choose **More** > **Download Authentication Credential**, specify the save path after the file is automatically generated, and keep the file properly.

The authentication credential includes the **krb5.conf** file of the Kerberos service.

After the authentication credential file is decompressed, you can obtain the following two files:

- The **krb5.conf** file contains the authentication service connection information.
- The **user.keytab** file contains user authentication information.

----End

8.7.1.2 Managing User Groups

Scenario

FusionInsight Manager supports a maximum of 5000 user groups (including built-in user groups). You can create and manage different user groups based on service scenarios on FusionInsight Manager. A user group is bound to a role to obtain operation permissions. After a user is added to a user group, the user can obtain the operation permissions of the user group. A user group can be used to classify users and manage multiple users.

Prerequisites

- You have learned service requirements and created roles required by service scenarios.
- You have logged in to FusionInsight Manager.

Creating a User Group

Step 1 Choose **System** > **Permission** > **User Group**.

Step 2 Above the user group list, click **Create User Group**.

Figure 8-51 Creating a user group

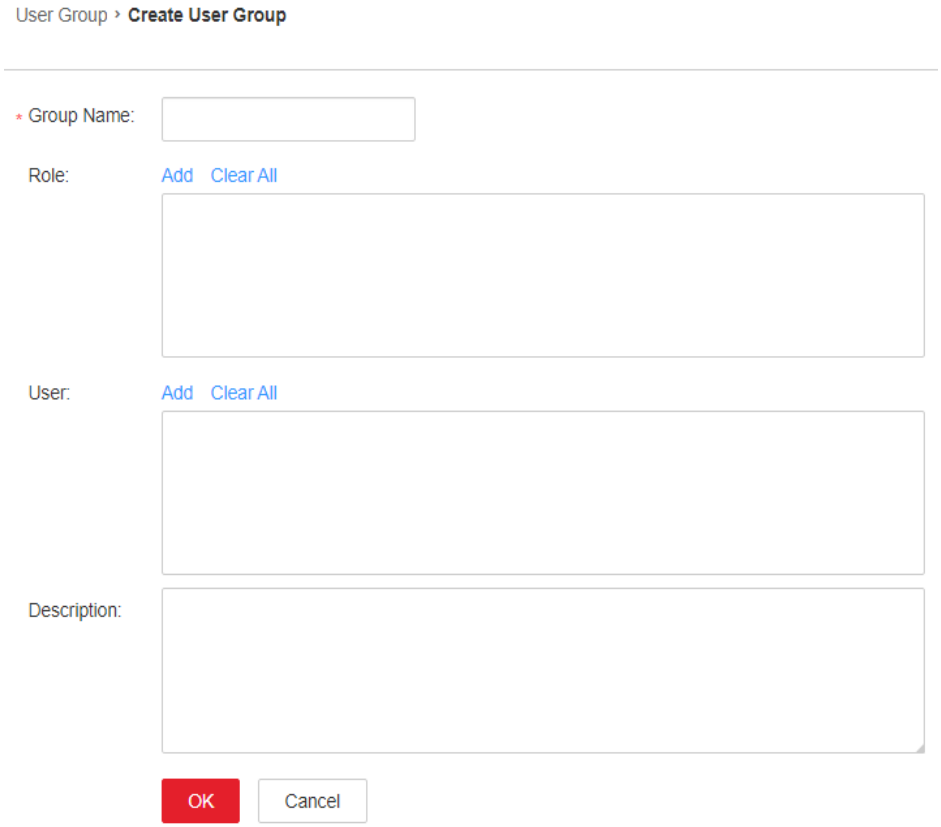

Step 3 Set **Group Name** and **Description**.

The group name contains 1 to 64 characters, including case-insensitive letters, digits, underscores (_), hyphens (-), and spaces. It cannot be the same as an existing user group name in the system.

Step 4 In the **Role** area, click **Add** to select a role and add it.

NOTE

- For components (except HDFS and Yarn) for which Ranger authorization has been enabled, the permissions of non-default roles on Manager do not take effect. You need to configure Ranger policies to assign permissions to user groups.
- If the resource requests of HDFS and Yarn are beyond the Ranger policies, the ACL rules of the components still take effect.

Step 5 In the **User** area, click **Add** to select a user and add it.

Step 6 Click **OK**.

The user group is created.

----End

Viewing User Group Information

By default, all user groups are displayed in the user group list. You can click the arrow on the left of a user group name to view details about the user group, including the user quantity, specific users, and bound roles of the user group.

Modifying Information About a User Group

Locate the row that contains the target user group, and click **Modify** to modify its information.

Exporting Information About a User Group

Click **Export All** to export all user group information at a time in **TXT** or **CSV** format.

The user group information contains the user group name, description, user list, and role list.

Deleting a User Group

Locate the row that contains the target user group, and click **Delete**. To delete multiple user groups in batches, select the target user groups and click **Delete** above the user group list. A user group that contains users cannot be deleted. To delete such a user group, delete all its users by modifying the user group first.

8.7.1.3 Managing Roles

Scenario

FusionInsight Manager supports a maximum of 5000 roles (including system built-in roles but excluding roles automatically created by tenants). Based on different service requirements, you need to create and manage different roles on FusionInsight Manager and perform authorization management for FusionInsight Manager and components using roles.

Prerequisites

- You have learned service requirements.
- You have logged in to FusionInsight Manager.

Creating a Role

Step 1 Choose **System** > **Permission** > **Role**.

Step 2 On the displayed page, click **Create Role** and fill in **Role Name** and **Description**.

The role name consists of 3 to 50 characters, including digits, letters, and underscores (_). It cannot be the same as an existing role name in the system. The role name cannot start with **Manager**, **System**, or **default**. For example, the role name cannot be **Manager_test**.

Figure 8-52 Creating a role

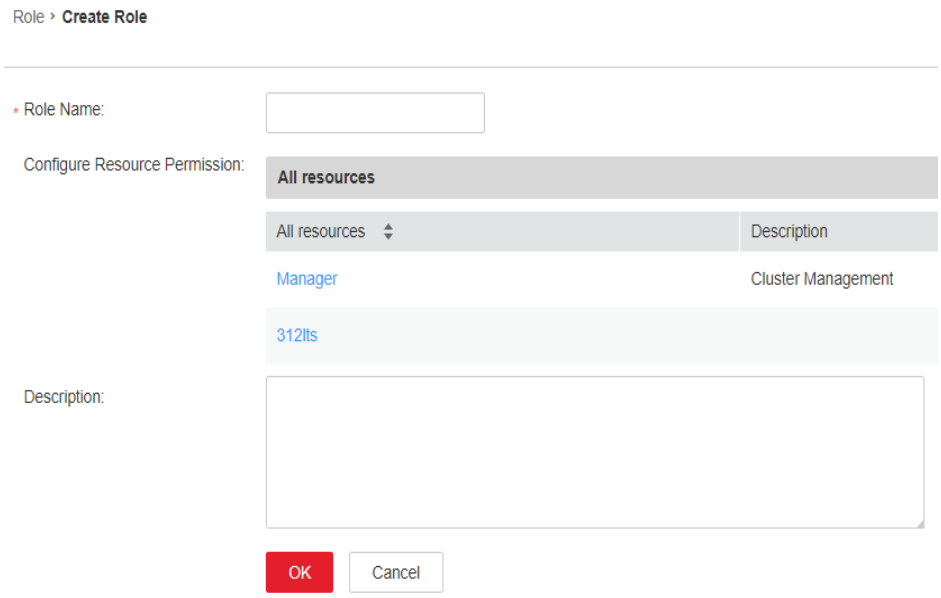

Step 3 In the **Configure Resource Permission** area, click the cluster whose permissions are to be added and select service permissions for the role.

When setting permissions for a component, enter a resource name in the search text box in the upper right corner and click the search icon to view the search result.

The search result contains only directories, but not subdirectories. Search by keyword supports fuzzy match and is case-insensitive.

NOTE

- For components (except HDFS and Yarn) for which Ranger authorization has been enabled, the permissions of non-default roles on Manager do not take effect. You need to configure Ranger policies to assign permissions to user groups.
- If the resource requests of HDFS and Yarn are beyond the Ranger policies, the ACL rules of the components still take effect.
- A maximum of 1000 permissions can be set for a component at a time.

Step 4 Click **OK**.

----End

Modifying Role Information

Locate the row that contains the target role and click **Modify**.

Exporting Role Information

Click **Export All** to export all role information at a time in **TXT** or **CSV** format.

Role information includes the role name and description.

Deleting a Role

Locate the row that contains the target role and click **Delete**. To delete multiple roles in batches, select the target roles and click **Delete** above the role list. A role

bound to a user cannot be deleted. To delete such a role, disassociate the role from the user by modifying the user first.

Task Example (Creating a Manager Role)

Step 1 Choose **System** > **Permission** > **Role**.

Step 2 On the displayed page, click **Create Role** and fill in **Role Name** and **Description**.

Step 3 In the **Configure Resource Permission** area, click **Manager** and set permissions for the role.

Figure 8-53 Setting permissions

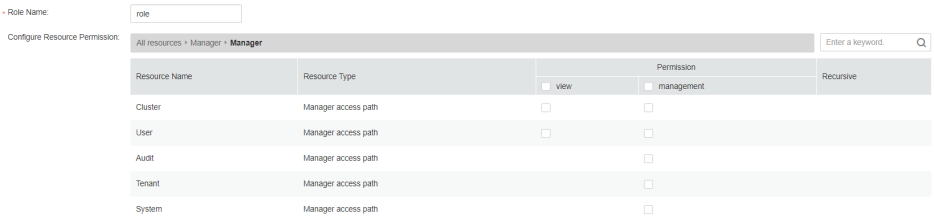

Manager permissions:

- Cluster
  - **view** permission: permission to view information on the **Cluster** page and view alarms and events under **O&M** > **Alarm**.
  - **management** permission: permission for management on the **Cluster** and **O&M** pages.
- User
  - **view** permission: permission to view information on pages under **System** > **Permission**.
  - **management** permission: permission for management on pages under **System** > **Permission**.
- Audit

  management permission: permission for management on the **Audit** page.
- Tenant

  management permission: permission for management on the **Tenant** page and permission to view alarms and events under **O&M** > **Alarm**.
- System

  management permission: permission for management on all pages except those under **Permission** on the **System** page and permission to view alarms and events under **O&M** > **Alarm**.

Step 4 Click **OK**.

----End

8.7.1.4 Security Policies

8.7.1.4.1 Configuring Password Policies

Scenario

To keep up with service security requirements, you can set password security rules, user login security rules, and user locking rules on FusionInsight Manager.

NOTICE

- Modify password policies based on service security requirements, because they involve user management security. Otherwise, security risks may be incurred.
- Change the user password after modifying the password policy, and then the new password policy can take effect.

Modifying a Password Policy (Versions Earlier Than MRS 3.1.2)

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **Security Policy** > **Password Policy**.

Step 3 Click **Modify** in the **Operation** column and modify the password policy as prompted.

For details about the parameters, see **Table 8-42**.

Table 8-42 Password policy parameters

Parameter	Description
Minimum Password Length	Indicates the minimum number of characters a password contains. The value ranges from **8** to **32**. The default value is **8**.
Character Types	Indicates how many character types in the following types a password can contain at least: uppercase letters, lowercase letters, digits, spaces, and special characters (~`!?,.:;-_'(){}[]/<>@#$%^&*+|\=). The value can be **4** or **5**. The default value is **4**, which means that a password can contain uppercase letters, lowercase letters, digits, and special characters. If you set the parameter to **5**, a password can contain all the five character types mentioned above.
Password Retries	Indicates the number of consecutive wrong password attempts allowed before the system locks the user. The value ranges from **3** to **30**. The default value is **5**.
User Lock Duration (Min)	Indicates the time period in which a user is locked when the user lockout conditions are met. The value ranges from **5** to **120**. The default value is **5**.

Parameter	Description
Password Validity Period (Day)	Indicates the validity period of a password. The value ranges from **0** to **90**. **0** indicates that the password is permanently valid. The default value is **90**.
Repetition Rule	Indicates the number of previous passwords that cannot be reused when you change the password. The value ranges from **1** to **5**. The default value is **1**. This policy applies to only human-machine accounts.
Password Expiration Notification (Days)	Indicates the number of days in advance users are notified that their passwords are about to expire. After the value is set, if the difference between the cluster time and the password expiration time is smaller than this value, the user receives password expiration notifications. When logging in to FusionInsight Manager, the user will be notified that the password is about to expire and a message is displayed asking the user to change the password. The value ranges from **0** to *X* (*X* must be set to the half of the password validity period and rounded down). Value **0** indicates that no notification is sent. The default value is **5**.
Interval for Deleting Authentication Failure Records (Min)	Indicates the interval of retaining incorrect password attempts. The value ranges from **0** to **1440**. **0** indicates that incorrect password attempts are permanently retained, and **1440** indicates that incorrect password attempts are retained for one day. The default value is **5**.

Step 4 Click **OK** to save the configurations. Change the user password after modifying the password policy, and then the new password policy can take effect.

----End

Adding a Password Policy (MRS 3.1.2 or Later)

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **Security Policy** > **Password Policy**.

Step 3 Click **Add Password Policy** and modify the password policy as prompted.

For details about the parameters, see **Table 8-43**.

Table 8-43 Password policy parameters

Parameter	Description
Password Policy Name	The value is a string of 3 to 32 characters, including case-insensitive letters, digits, underscores (_), and hyphens (-). It cannot start with a hyphen (-).

Parameter	Description
Minimum Password Length	Indicates the minimum number of characters a password contains. The value ranges from **8** to **32**.
Character Types	Indicates how many character types in the following types a password can contain at least: uppercase letters, lowercase letters, digits, spaces, and special characters (~`!?,.:;-_'(){}[]/<>@#$%^&*+|\=). The value can be **4** or **5**. The default value is **4**, which means that a password can contain uppercase letters, lowercase letters, digits, and special characters. If you set the parameter to **5**, a password can contain all the five character types mentioned above.
Password Retries	Indicates the number of consecutive wrong password attempts allowed before the system locks the user. The value ranges from **3** to **30**.
User Lock Duration (Min)	Indicates the time period in which a user is locked when the user lockout conditions are met. The value ranges from **5** to **120**.
Password Validity Period (Day)	Indicates the validity period of a password. The value ranges from **0** to **90**. **0** indicates that the password is permanently valid.
Repetition Rule	Indicates the number of previous passwords that cannot be reused when you change the password. The value ranges from **1** to **5**. The default value is **1**. This policy applies to only human-machine accounts.
Password Expiration Notification (Days)	Indicates the number of days in advance users are notified that their passwords are about to expire. After the value is set, if the difference between the cluster time and the password expiration time is smaller than this value, the user receives password expiration notifications. When logging in to FusionInsight Manager, the user will be notified that the password is about to expire and a message is displayed asking the user to change the password. The value ranges from **0** to *X* (*X* must be set to the half of the password validity period and rounded down). Value **0** indicates that no notification is sent.
Interval for Deleting Authentication Failure Records (Min)	Indicates the interval of retaining incorrect password attempts. The value ranges from **0** to **1440**. **0** indicates that incorrect password attempts are permanently retained, and **1440** indicates that incorrect password attempts are retained for one day.

Step 4 Click **OK** to save the configurations.

A new user uses the default password policy. After a new password policy is created, you can manually select the password policy when creating a user. You

can modify the password policy of an existing user. For details, see **Modifying User Information**.

----End

NOTE

A maximum of 32 password policies can be created.

Modifying a Password Policy (MRS 3.1.2 or Later)

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **Security Policy** > **Password Policy**.

Step 3 Click **Modify** in the row that contains the target password policy. On the **Modify Password Policy** page, modify the password policy as prompted.

For details about the parameters, see **Table 8-43**.

Step 4 Click **OK** to save the configurations.

----End

NOTE

- Users (except **admin**) cannot modify their own password policies.
- After the password policy bound to a user is modified, if the remaining password validity period is greater than the password validity period in the new password policy, the password validity period is set to the validity period in the new password policy. If the remaining password validity period is less than the password validity period in the new password policy, the password validity period remains unchanged.

Deleting a Password Policy (MRS 3.1.2 or Later)

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **Security Policy** > **Password Policy**.

Step 3 Click **Delete** in the row that contains the target password policy. In the dialog box that is displayed, click **OK**.

----End

NOTE

The default password policy and the password policy that has been bound to a user cannot be deleted.

8.7.1.4.2 Configuring the Independent Attribute

Scenario

User **admin** or administrators who are bound to the Manager_administrator role can configure the independent attribute on FusionInsight Manager so that common users (all service users in the cluster) can set or cancel their own independent attributes.

After the independent attribute option is toggled on, service users need to log in to the system and set the independent attribute.

Constraints

- Administrators cannot set or cancel the independent attribute of a user.
- Administrators cannot obtain the authentication credentials of independent users.

Prerequisites

You have obtained the required administrator username and password.

Procedure

Toggling On or Off the Independent Attribute

Step 1 Log in to FusionInsight Manager as user **admin** or a user bound to the Manager_administrator role.

Step 2 Choose **System** > **Permission** > **Security Policy** > **Independent Configurations**.

Step 3 Toggle on or off **Independent Attribute**, enter the password as prompted, and click **OK**.

Step 4 After the identity is authenticated, wait until the OMS configuration is modified and click **Finish**.

NOTE

After the independent attribute is disabled:

- A user who has the attribute can cancel it from the drop-down list of the username in the upper right corner of the page. The user cannot set the independent attribute again once it is cancelled. After the attribute is cancelled, existing independent tables will retain the attribute. However, the user cannot create independent tables again.
- Users without this attribute cannot set or cancel the attribute.

Configuring the Independent Attribute

Step 5 Log in to FusionInsight Manager as a service user.

NOTICE

Administrators cannot initialize the password of the user after the independent attribute is set. If the user password is forgotten, the password cannot be retrieved.

User **admin** cannot set the independent attribute.

Step 6 Move the cursor to the username in the upper right corner of the page.

Step 7 Select **Set Independent** or **Cancel Independent**.

NOTE

- If the independent attribute is toggled on and has been set for the service user, **Cancel Independent** is displayed.
- If the independent attribute is toggled on but has been cancelled for the service user, **Set Independent** is displayed.
- If the independent attribute is toggled off but has been set for the service user, **Cancel Independent** is displayed.
- If the independent attribute is toggled off and has been cancelled for the service user, no option related to the independent attribute is displayed.

Step 8 Enter the password as prompted and click **OK**.

Step 9 After the identity is authenticated, click **OK** in the dialog box.

----End

8.7.2 Configuring Interconnections

8.7.2.1 Configuring SNMP Northbound Parameters

Scenario

If users need to view alarms and monitoring data of a cluster on the O&M platform, you can use Simple Network Management Protocol (SNMP) on FusionInsight Manager to report related data to the network management system (NMS).

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Interconnection** > **SNMP**.

Step 3 Toggle on **SNMP Service**.

The SNMP service is disabled by default. indicates that the service is enabled.

Step 4 Set interconnection parameters according to **Table 8-44**.

Table 8-44 Interconnection parameters

Parameter	Description
Version	Specifies the version of SNMP, which can be: • **V2C**: This is an earlier version with low security. • **V3**: This is a later version with higher security than SNMPv2c SNMPv3 is recommended.
Local Port	Specifies the local port. The default value is **20000**. The value ranges from **1025** to **65535**.

Parameter	Description
Read Community Name	Specifies the read-only community name. This parameter is available only when **Version** is set to **V2C**.
Write Community Name	Specifies the write community name. This parameter is available only when **Version** is set to **V2C**.
Security Username	Specifies the SNMP security username. This parameter is available only when **Version** is set to **V3**.
Authentication Protocol	Specifies the authentication protocol. This parameter is available only when **Version** is set to **V3**. SHA is recommended.
Authentication Password	Specifies the authentication password. This parameter is available only when **Version** is set to **V3**.
Confirm Password	Used to confirm the authentication password. This parameter is available only when **Version** is set to **V3**.
Encryption Protocol	Specifies the encryption protocol. This parameter is available only when **Version** is set to **V3**. AES256 is recommended.
Encryption Password	Specifies the encryption password. This parameter is available only when **Version** is set to **V3**.
Confirm Password	Used to confirm the encryption password. This parameter is available only when **Version** is set to **V3**.

NOTE

- The value of **Security Username** cannot contain repeated strings with the unit length as a common factor of 64 (such as 1, 2, 4, and 8), for example, **abab** and **abcdabcd**.
- The **Authentication Password** and **Encryption Password** must contain 8 to 16 characters, including at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters. The two passwords must be different. The two passwords cannot be the same as the security username or the reverse of the security username.
- For security purposes, periodically change the authentication password and encryption password when the SNMP protocol is used.
- If SNMP v3 is used, a security user will be locked after five consecutive authentication failures within 5 minutes. The user will be automatically unlocked 5 minutes later.

Step 5 Click **Create Trap Target** in the **Trap Target** area. In the displayed dialog box, set the following parameters:

- **Target Symbol**: specifies the trap target ID, which is the ID of the NMS or host that receives traps. The value consists of 1 to 255 characters, including letters or digits.
- **Target IP Address Mode**: specifies the mode of the target IP address. The value can be **IPv4** or **IPv6**.
- **Target IP Address**: specifies the target IP address, which can communicate with the management plane IP address of the management node.
- **Target Port**: specifies the port receiving traps. The port number must be consistent with the peer end and ranges from 0 to 65535.
- **Trap Community Name**: This parameter is available only when **Version** is set to **V2C** and is used to report the community name.

Click **OK**.

The **Create Trap Target** dialog box is closed.

Step 6 Click **OK**.

----End

8.7.2.2 Configuring Syslog Northbound Parameters

Scenario

If users need to view alarms and events of a cluster on the unified alarm reporting platform, you can use the Syslog protocol on FusionInsight Manager to report related data to the alarm platform.

NOTICE

If the Syslog protocol is not encrypted, data may be stolen.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Interconnection** > **Syslog**.

Step 3 Toggle on **Syslog Service**.

The Syslog service is disabled by default. indicates that the service is enabled.

Step 4 Set northbound parameters according to **Table 8-45**.

Table 8-45 Syslog interconnection parameters

Parameter Area	Parameter	Description
Syslog Protocol	Server IP Address Mode	Specifies the IP address mode of the interconnected server. The value can be **IPV4** or **IPV6**.

Parameter Area	Parameter	Description
	Server IP Address	Specifies the IP address of the interconnected server.
	Server Port	Specifies the port number for interconnection.
	Protocol	Specifies the protocol type. The options are as follows: • **TCP** • **UDP**
	Severity Level	Specifies the severity of the reported message. The options are as follows: • **Emergency** • **Alert** • **Critical** • **Error** • **Warning** • **Notice** • **Informational** (default value) • **Debug** **NOTE** **Severity Level** and **Facility** determine the priority of the sent message. **Priority** = **Facility** × 8 + **Severity Level** For details about the values of **Severity Level** and **Facility**, see **Table 8-46**.
	Facility	Specifies the module where the log is generated. For details about the available values of this parameter, see **Table 8-46**. Default value **local use 0 (local0)** is recommended.
	Identifier	Specifies the product ID. The default value is **FusionInsight Manager**. The identifier can contain a maximum of 256 characters, including letters, digits, underscores (_), periods (.), hyphens (-), spaces, and the following special characters: | $ { }

Parameter Area	Parameter	Description
Report Message	Report Format	Specifies the message format of the alarm report. For details, see the help information on the page. The report format can contain a maximum of 1024 characters, including letters, digits, underscores (_), periods (.), hyphens (-), spaces, and the following special characters: | $ { } **NOTE** For details about each field in the report format, see **Table 8-47**.
	Alarm Type	Specifies the type of the alarm to be reported.
	Alarm Severities	Specifies the level of the alarm to be reported.
Uncleared Alarm Reporting	Periodic Uncleared Alarm Reporting	Specifies whether to report uncleared alarms in a specified period. You can toggle on or off the function. The function is toggled off by default.
	Report Interval (min)	Specifies the interval for periodically reporting uncleared alarms. This parameter is valid only when **Periodic Uncleared Alarm Reporting** is enabled. The default value is **15**, in minutes. The value ranges from **5** to **1440** (one day).
Heartbeat Settings	Heartbeat Reporting	Specifies whether to periodically report Syslog heartbeat messages. You can toggle on or off the function. The function is toggled off by default.
	Heartbeat Interval (minutes)	Specifies the interval for periodically reporting heartbeat messages. This parameter is valid only when **Heartbeat Reporting** is enabled. The default value is **15**, in minutes. The value ranges from **1** to **60**.
	Heartbeat Packet	Specifies the heartbeat message to be reported. This parameter is valid when **Heartbeat Reporting** is toggled on and cannot be left blank. The value can contain a maximum of 256 characters, including digits, letters, underscores (_), vertical bars (|), colons (:), spaces, commas (,), and periods (.).

NOTE

After the periodic heartbeat packet function is enabled, packets may be interrupted during automatic recovery of some cluster error tolerance (for example, active/standby OMS switchover). In this case, wait for automatic recovery.

Step 5 Click **OK**.

----End

Related Information

Table 8-46 Numeric codes of **Severity Level** and **Facility**

Severity Level	Facility	Numeric Code
Emergency	kernel messages	0
Alert	user-level messages	1
Critical	mail system	2
Error	system daemons	3
Warning	security/authorization messages (note 1)	4
Notice	messages generated internally by syslog	5
Informational	line printer subsystem	6
Debug	network news subsystem	7
-	UUCP subsystem	8
-	clock daemon (note 2)	9
-	security/authorization messages (note 1)	10
-	FTP daemon	11
-	NTP subsystem	12
-	log audit (note 1)	13
-	log alert (note 1)	14
-	clock daemon (note 2)	15
-	local use 0~7 (local0 ~ local7)	16 to 23

Table 8-47 Report format information fields

Information Field	Description
dn	Cluster name
id	Alarm ID
name	Alam name

Information Field	Description
serialNo	Alarm serial number **NOTE** The serial numbers of the fault alarms and the corresponding clear alarms are the same.
category	Alarm type. The options are as follows: • **0**: fault alarm • **1**: clear alarm • **2**: event
occurTime	Time when the alarm was generated
clearTime	Time when this alarm was cleared
isAutoClear	Whether an alarm is automatically cleared. The options are as follows: • **1**: yes • **0**: no
locationInfo	Location where the alarm was generated
clearType	Alarm clearance type. The options are as follows: • **-1**: not cleared • **0**: automatically cleared • **2**: manually cleared
level	Severity. The options are as follows: • **1**: critical alarm • **2**: major alarm • **3**: minor alarm • **4**: warning alarm
cause	Alarm cause
additionalInfo	Additional information
object	Alarm object

8.7.2.3 Configuring Monitoring Metric Dumping

Scenario

The monitoring data reporting function writes the monitoring data collected in the system into a text file and uploads the file to a specified server in FTP or SFTP mode.

Before using this function, you need to perform related configurations on FusionInsight Manager.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Interconnection** > **Upload Performance Data**.

Step 3 Toggle on **Upload Performance Data**.

The performance data upload service is disabled by default. indicates that the service is enabled.

Step 4 Set the upload parameters according to **Table 8-48**.

Table 8-48 Upload parameters

Parameter	Description
FTP IP Address Mode	Specifies the server IP address mode. This parameter is mandatory. The value can be **IPV4** or **IPV6**.
FTP IP Address	Specifies the IP address of the FTP server for storing monitoring files after the monitoring metric data is interconnected. This parameter is mandatory.
FTP Port	Specifies the port for connecting to the FTP server. This parameter is mandatory.
FTP Username	Specifies the username for logging in to the FTP server. This parameter is mandatory.
FTP Password	Specifies the password for logging in to the FTP server. This parameter is mandatory.
Save Path	Specifies the path for storing monitoring files on the FTP server. This parameter is mandatory.
Dump Interval (second)	Specifies the interval at which monitoring files are periodically stored on the FTP server, in seconds. This parameter is mandatory.
Dump Mode	Specifies the protocol used for sending monitoring files. This parameter is mandatory. The value can be **SFTP** or **FTP**. You are advised to use the SFTP mode based on SSH v2. Otherwise, security risks may be incurred.
SFTP Service Public Key	Specifies the public key of the FTP server. This parameter is optional. It is valid only when **Dump Mode** is set to **SFTP**.

Step 5 Click **OK**.

NOTE

If the dump mode is SFTP and the public key of the SFTP service is empty, the system displays a security risk warning. You need to evaluate the security risk and then save the configuration.

----End

Data Format

After the configuration is complete, the monitoring data reporting function periodically writes monitoring data in the cluster to text files and reports the files to the corresponding FTP/SFTP service based on the configured reporting period.

- Principles for generating monitoring files
  - The monitoring metrics are written to files generated every 30, 60, and 300 seconds based on the metric collection period.

    30s: real-time metrics that are collected every 30s by default

    60s: real-time metrics that are collected every 60s by default

    300s: all metrics that are not collected every 30s or 60s
  - File name format: *metric_{Interval}_{File creation time YYYYMMDDHHMMSS}.log*

    Example: **metric_60_20160908085915.log**

    metric_300_20160908085613.log
- Monitoring file content
  - Format of monitoring files:

    "Cluster ID|Cluster name|Displayed name|Service name|Metric ID|Collection time|Collection host@m@Sub-metric|Unit|Metric value", where fields are separated using vertical bars (|). For example:

    ```
    1|xx1|Host|Host|10000413|2019/06/18 10:05:00|189-66-254-146|KB/s|309.910
    1|xx1|Host|Host|10000413|2019/06/18 10:05:00|189-66-254-152|KB/s|72.870
    2|xx2|Host|Host|10000413|2019/06/18 10:05:00|189-66-254-163|KB/s|100.650
    ```

    Note: The actual files are not in that format.
  - Interval for uploading monitoring files:

    The interval for uploading monitoring files can be set using the **Dump Interval (second)** parameter on the page. Currently, the interval can range from **30** to **300**. After the configuration is complete, the system periodically uploads files to the corresponding FTP/SFTP server at the specified interval.
- Monitoring metric description file
  - Metric set file

    The metric set file **all-shown-metric-zh_CN** contains detailed information about all metrics. After obtaining the metric IDs from the files reported by the third-party system, you can query details about the metrics from the metric set file.

    Location of the metric set file:

    Active and standby OMS nodes: {*FusionInsight installation path*} **/om-server/om/etc/om/all-shown-metric-zh_CN**

    Content of the metric set file:

    ```
    Real-Time Metric ID,5-Minute Metric ID,Metric Name,Metric Collection Period (s),Collected by Default,Service Belonged To,Role Belonged To
    00101,10000101,JobHistoryServer non-heap memory usage,30,false,Mapreduce,JobHistoryServer
    00102,10000102,JobHistoryServer non-heap memory allocation volume,30,false,Mapreduce,JobHistoryServer
    00103,10000103,JobHistoryServer heap memory usage,30,false,Mapreduce,JobHistoryServer
    00104,10000104,JobHistoryServer heap memory allocation volume,30,false,Mapreduce,JobHistoryServer
    00105,10000105,Number of blocked threads,30,false,Mapreduce,JobHistoryServer
    ```



```
00106,10000106,Number of running threads,30,false,Mapreduce,JobHistoryServer
00107,10000107,GC time,30,false,Mapreduce,JobHistoryServer
00110,10000110,JobHistoryServer CPU usage,30,false,Mapreduce,JobHistoryServer
...
```

- Field description of critical metrics

  Real-Time Metric ID: indicates the ID of the metric whose collection period is 30s or 60s.

  5-Minute Metric ID: indicates the ID of a 5-minute (300s) metric.

  Metric Collection Period (s): indicates the collection period of real-time metrics. The value can be **30** or **60**.

  Service Belonged To: indicates the name of the service to which a metric belongs, for example, HDFS and HBase.

  Role Belonged To: indicates the name of the role to which a metric belongs, for example, JobServer and RegionServer.

- Description

  For metrics whose collection period is 30s/60s, you can find the corresponding metric description by referring to the first column, that is, **Real-Time Metric ID**.

  For metrics whose collection period is 300s, you can find the corresponding metric description by referring to the second column, that is, **5-Minute Metric ID**.

8.7.3 Importing a Certificate

Scenario

CA certificates are used to encrypt data during communication between FusionInsight Manager modules and between cluster component clients and servers to ensure security. CA certificates can be quickly imported to FusionInsight Manager for product security. Import CA certificates in following scenarios:

- When the cluster is installed for the first time, you need to replace the enterprise certificate.
- If the enterprise certificate has expired or security hardening is required, you need to replace it with a new certificate.

Impact on the System

- During certificate replacement, the cluster needs to be restarted. In this case, the system becomes inaccessible and cannot provide services.
- After the certificate is replaced, the certificates used by all components and FusionInsight Manager modules are automatically updated.
- After the certificate is replaced, you need to reinstall the certificate in the local environment where the certificate is not trusted.

Prerequisites

- You have generated the certificate file and key file or obtained them from the enterprise certificate administrator.
- You have obtained the files to be imported to the cluster, including the CA certificate file (***.crt**), key file (***.key**), and file that saves the key file password

(**password.property**). The certificate name and key name can contain uppercase letters, lowercase letters, and digits. After the preceding files are generated, compress them into a TAR package.

- You have obtained a password for accessing the key file, for example, **Userpwd@123**.

  To avoid potential security risks, the password must meet the following complexity requirements:

  – Contains at least 8 characters.

  – Contains at least four types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!?,.;-_'(){}[]/<>@#$%^&*+|\=).

- When applying for certificates from the certificate administrator, you have provided the password for accessing the key file and applied for the certificate files in CRT, CER, CERT, and PEM formats and the key files in KEY and PEM formats. The requested certificates must have the issuing function.

Procedure

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

Step 2 Click ··· or **Select File** on the right of **Upload Certificate**. In the file selection window, browse to select the obtained TAR package of the certificate files.

Step 3 Click **Upload**.

Manager uploads the compressed package and automatically imports the package.

Step 4 After the certificate is imported, the system displays a message asking you to synchronize the cluster configuration and restart the web service for the new certificate to take effect. Click **OK**.

Step 5 In the displayed dialog box, enter the password of the current login user and click **OK**. The cluster configuration is automatically synchronized and the web service is restarted.

NOTE

If the page is refreshed or the browser is closed during cluster configuration synchronization, perform the following operations to restart the web service:

1. Log in to the active OMS node as user **omm**.
2. Run the following command to restart HTTPD. **xxx** indicates the HTTPD version number. Replace it with the actual version number.

   sh ${BIGDATA_HOME}/om-server/Apache-httpd-*xxx*/setup/restarthttpd.sh
3. Run the following command to restart Tomcat:

   sh ${BIGDATA_HOME}/om-server/tomcat/bin/shutdown.sh;sh ${BIGDATA_HOME}/om-server/tomcat/bin/startup.sh
4. Run the following command to restart Knox:

   sh /opt/knox/bin/restart-knox.sh

Step 6 After the cluster is restarted, enter the URL for accessing FusionInsight Manager in the address box of the browser and check whether the FusionInsight Manager web page can be successfully displayed.

Step 7 Log in to FusionInsight Manager.

Step 8 Choose **Cluster** > **Dashboard**. In the upper right corner, click **More** > **Restart**. (For MRS 3.3.0 or later, choose **More** > **Restart** in the upper right corner of the home page.)

Step 9 In the displayed dialog box, enter the password of the current login user and click **OK**.

----End

8.7.4 OMS Management

8.7.4.1 Overview of the OMS Page

Overview

Log in to FusionInsight Manager and choose **System** > **OMS**. You can perform maintenance operations on the OMS page, including viewing basic information, viewing the service status of OMS service modules, and manually triggering health checks.

NOTE

OMS is the management node of the O&M system. Generally, there are two OMS nodes that work in active/standby mode.

Basic Information

OMS-associated information is displayed on FusionInsight Manager, as listed in Table 8-49.

Table 8-49 OMS information

Item	Description
Version	Indicates the OMS version, which is consistent with the FusionInsight Manager version.
IP Mode	Indicates the IP address mode of the current cluster network.
HA Mode	Indicates the OMS working mode, which is specified by the configuration file during FusionInsight Manager installation.
Current Active	Indicates the host name of the active OMS node, that is, the host name of the active management node. Click a host name to go to the host details page.
Current Standby	Indicates the host name of the standby OMS node, that is, the host name of the standby management node. Click a host name to go to the host details page.

Item	Description
Duration	Indicates the duration for starting the OMS process.

OMS Service Status

FusionInsight Manager displays the running status of all OMS service modules. If the status of each service module is displayed as ●, the OMS is running properly.

Health Check

You can click **Health Check** on the OMS page to check the OMS status. If some check items are faulty, you can view the check description for troubleshooting.

Figure 8-54 Health check

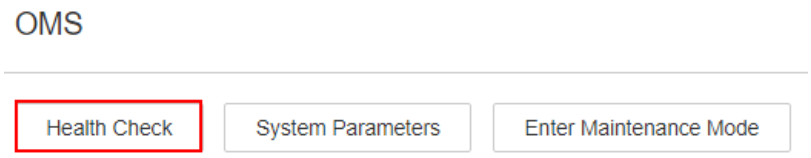

Entering or Exiting Maintenance Mode

Configure OMS to enter or exit the maintenance mode.

System Parameters

Connect to the DMPS cluster in large-scale cluster scenarios.

8.7.4.2 Modifying OMS Service Configuration Parameters

Scenario

Based on the security requirements of the user environment, you can modify the Kerberos and LDAP configurations in the OMS on FusionInsight Manager.

Impact on the System

After the OMS service configuration parameters are modified, the corresponding OMS module needs to be restarted. In this case, FusionInsight Manager cannot be used.

Procedure

Modifying the okerberos configuration

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

Step 2 Locate the row that contains okerberos and click **Modify Configuration**.

Step 3 Modify the parameters according to **Table 8-50**.

Table 8-50 okerberos parameters

Parameter	Description
KDC Timeout (ms)	Timeout duration for an application to connect to Kerberos, in milliseconds. The value must be an integer.
Max Retries	Maximum number of retries for an application to connect to Kerberos, in seconds. The value must be an integer.
LDAP Timeout (ms)	Timeout duration for Kerberos to connect to LDAP, in milliseconds.
LDAP Search Timeout (ms)	Timeout duration for Kerberos to query user information in LDAP, in milliseconds.
Kadmin Listening Port	Port number of the Kadmin service.
KDC Listening Port	Port number of the kinit service.
Kpasswd Listening Port	Port number of the Kpasswd service.
Reset LDAP Account Password	Machine-machine users (**cn=krbadmin,ou=Users,dc=hadoop,dc=com** and **cn=krbkdc,ou=Users,dc=hadoop,dc=com**) used by Kerberos to access LDAP. If this parameter is selected, the passwords will be replaced by random passwords. **NOTE** This parameter is available only in MRS 3.1.2 or later.

Step 4 Click **OK**.

In the displayed dialog box, enter the password of the current login user and click **OK**. In the displayed confirmation dialog box, click **OK**.

Modifying the oldap configuration

Step 5 Locate the row that contains the oldap and click **Modify Configuration**.

Step 6 Modify the parameters according to **Table 8-51**.

Table 8-51 OLDAP parameters

Parameter	Description
LDAP Listening Port	Port number of the LDAP service.
Reset LDAP Account Password	Machine-machine users (**cn=root,dc=hadoop,dc=com** and **cn=pg_search_dn,ou=Users,dc=hadoop,dc=com**) used by LDAP for data management, synchronization, and status check. If this parameter is selected, the passwords will be replaced by random passwords. **NOTE** This parameter is available only in MRS 3.1.2 or later.

Step 7 Click **OK**.

In the displayed dialog box, enter the password of the current login user and click **OK**. In the displayed confirmation dialog box, click **OK**.

NOTE

To reset the password of the LDAP account, you need to restart ACS. The procedure is as follows:

1. Log in to the active management node as user **omm** using PuTTY, and run the following command to update the domain configuration:

   sh ${BIGDATA_HOME}/om-server/om/sbin/restart-RealmConfig.sh

   The command is run successfully if the following information is displayed:

   ```
   Modify realm successfully. Use the new password to log into FusionInsight again.
   ```

2. Run the **sh $CONTROLLER_HOME/sbin/acs_cmd.sh stop** command to stop ACS.
3. Run the **sh $CONTROLLER_HOME/sbin/acs_cmd.sh start** command to start ACS.

Restarting the cluster

Step 8 Log in to FusionInsight Manager and restart the cluster by referring to **Performing a Rolling Restart of a Cluster**.

----End

8.7.5 Viewing Component Packages

Scenario

A complete MRS cluster consists of multiple component packages. Before installing some services on FusionInsight Manager, check whether the component packages of those services have been installed.

Procedure

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

Step 2 On the **Installed Component** page, view all components.

NOTE

In the **Platform Type** column, you can view the registered OS and platform type of the component.

Step 3 Click ⌄ on the left of a component name to view the services and version numbers contained in the component.

----End

8.8 Cluster Management

8.8.1 Cluster Mutual Trust Management

8.8.1.1 Overview of Mutual Trust Between Clusters

Function Description

By default, users of a big data cluster in security mode can only access resources in the cluster but cannot perform identity authentication or access resources in other clusters in security mode.

Feature Description

- **Domain**

  The secure usage scope of users in each system is called a domain. Each FusionInsight Manager must have a unique domain name. Cross-Manager access allows users to use resources across domains.

- **User Encryption**

  Mutual trust can be configured across FusionInsight Managers. The current Kerberos server supports only the aes256-cts-hmac-sha1-96:normal and aes128-cts-hmac-sha1-96:normal encryption types for encrypting cross-domain users, and the encryption types cannot be changed.

- **User Authentication**

  After cross-Manager mutual trust is configured, if a user with the same name exists in two systems and the user in the peer system has the permission to access a resource in that system, this user can also access the remote resource.

- **Direct Mutual Trust**

  The system saves the mutual trust ticket of the peer system in two clusters with mutual trust configured and uses the mutual trust ticket to access the peer system.

8.8.1.2 Changing Manager's Domain Name

Scenario

The secure usage scope of users in each system is called a domain. Each system must have a unique domain name. The domain name of FusionInsight Manager is

generated during installation. The system administrator can change the domain name on FusionInsight Manager.

NOTICE

- Changing the system domain name is a high-risk operation. Before performing operations in this section, ensure that the OMS data has been backed up by referring to **Backing Up Manager Data**.

Impact on the System

- During the configuration, all of the clusters need to be restarted and are unavailable during restart.
- After the domain name is changed, the passwords of the Kerberos administrator and OMS Kerberos administrator will be initialized. You need to use the default passwords and then change the passwords. If a component user whose password is generated randomly by the system is used for identity authentication, see **Exporting an Authentication Credential File** to download the keytab file again.
- After the domain name is changed, passwords of the **admin** user, component user, and human-machine user added by the system administrator before the domain name change will be reset to the same one. Change these passwords. The reset password consists of two parts: one part is generated by the system and the other is set by the user. The system generating part is **Admin@123**, which is the default password. For details about the user-defined part, see descriptions of **Password Suffix** in **Table 8-53**. For example, if the system generates **Admin@123** and the user sets **Test#$%@123**, the new password after reset is **Admin@123Test#$%@123**.
- The new password must meet the password policies. To obtain the new human-machine user password, log in to the active OMS as user **omm** and run the following script:

  sh ${BIGDATA_HOME}/om-server/om/sbin/get_reset_pwd.sh *Password suffix user_name*

  - *Password suffix* is a parameter set by the user. If it is not specified, the default value **Admin@123** is used.
  - *user_name* is optional. The default value is **admin**.
  - There can be security risks if a command contains the authentication password. You are advised to disable the command recording function (history) before running the command.

  Example:

  sh ${BIGDATA_HOME}/om-server/om/sbin/get_reset_pwd.sh Test#$%@123

  ```
  To get the reset password after changing cluster domain name.
    pwd_min_len     : 8
    pwd_char_types  : 4
  The password reset after changing cluster domain name is: "Admin@123Test#$%@123"
  ```

  In this example, **pwd_min_len** and **pwd_char_types** indicate the minimum password length and number of password character types respectively defined in the password policies. **Admin@123Test#$%@123** indicates the human-machine user password after the system domain name is changed.

- After the system domain name is changed, the reset password consists of two parts: one part is generated by the system and the other is set by the user. The reset password must meet the password policies. If the password is not long enough, one or multiple at signs (@) are added between **Admin@123** and the user-defined part. If there are five character types, a space is added after **Admin@123**.

  When the user-defined part is **Test@123** and the default user password policy is used, the new password is **Admin@123Test@123**. The password contains 17 characters of four types. To meet the current password policy, the new password is processed according to **Table 8-52**.

  Table 8-52 Password processing

Minimum Password Length	Number of Character Types	Processing Against the Password Policy	New Password
8 to 17 characters	4	The user password policy is met.	Admin@123Test@123
18 characters	4	Add an at sign (@).	Admin@123@Test@123
19 characters	4	Add two at signs (@).	Admin@123@@Test@123
8 to 18 characters	5	Add a space.	Admin@123 Test@123
19 characters	5	Add a space and an at sign (@).	Admin@123 @Test@123
20 characters	5	Add a space and two at signs (@).	Admin@123 @@Test@123

- After the system domain name is changed, download the **keytab** file for the machine-machine user added by the system administrator before the domain name is changed.
- After the system domain name is changed, download and install the client again.
- After the system domain name is changed, if there is any running HetuEngine compute instance, restart the instance.

Prerequisites

- The system administrator has clarified service requirements and planned domain names for the systems.

  A domain name can contain only uppercase letters, numbers, periods (.), and underscores (_), and must start with a letter or number, for example, **DOMAINA.HW** and **DOMAINB.HW**.
- The running status of all components in the Manager clusters is **Normal**.

- The **acl.compare.shortName** parameter of the ZooKeeper service of all clusters in Manager is set to default value **true**. Otherwise, change the value to **true** and restart the ZooKeeper service.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **Domain and Mutual Trust**.

Figure 8-55 Domain and mutual trust

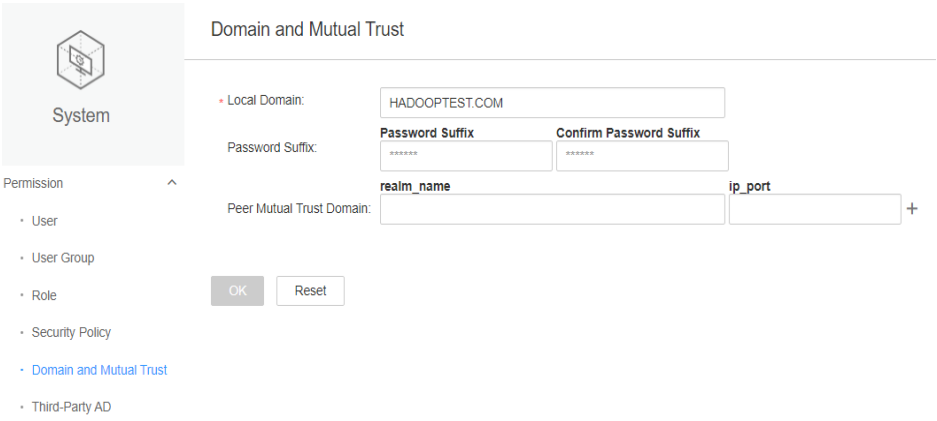

Step 3 Modify required parameters.

Table 8-53 Related parameters

Parameter	Description
Local Domain	Planned domain name of the system.
Password Suffix	Part of the password set by the user after the password of the human-machine user is reset. This parameter is mandatory. The default value is **Admin@123**. **NOTE** This parameter takes effect only after **Local Domain** is modified. The following conditions must be met: • The password ranges from 8 to 16 characters. • The password must contain at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (`~!@#$%^&*()-_=+|[{}];:',<.>/? and spaces).

Step 4 Determine whether to set up mutual trust between managers across cluster.

- If you need to do so, perform operations by referring to **Configuring Cross-Manager Mutual Trust Between Clusters** and skip subsequent steps in this section.
- If you do not need such configuration, go to **Step 5**.

Step 5 Click **OK**. Proceed with the subsequent steps only after the modification is complete.

Step 6 Log in to the active management node as user **omm**.

Step 7 Run the following command to update the domain configuration:

sh ${BIGDATA_HOME}/om-server/om/sbin/restart-RealmConfig.sh

The command is executed successfully if the following information is displayed:

```
Modify realm successfully. Use the new password to log in to FusionInsight again.
```

NOTE

After the restart, some hosts and services cannot be accessed and an alarm is generated. This problem can be automatically resolved in about 1 minute after **restart-RealmConfig.sh** is run.

Step 8 Log in to FusionInsight Manager using the new password of user **admin** (for example, **Admin@123Admin@123**). On the dashboard, click or **More** and select **Restart**.

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

In the displayed dialog box, click **OK**. Wait for a while until a message indicating that the operation is successful is displayed. Click **Finish**.

Step 9 Log out of FusionInsight Manager and then log in again. If the login is successful, the configuration is successful.

Step 10 Log in to the active management node as user **omm** and run the following command to update the configurations of the job submission client:

sh /opt/executor/bin/refresh-client-config.sh

Step 11 If a HetuEngine compute instance is running, restart the compute instance.

1. Log in to FusionInsight Manager as the user who is used to access the HetuEngine web UI.
2. Choose **Cluster** > **Services** > **HetuEngine** to go to the HetuEngine service page.
3. In the **Basic Information** area on the **Dashboard** page, click the link next to **HSConsole WebUI**. The HSConsole page is displayed.
4. For a running compute instance, click **Stop** in the **Operation** column. After the compute instance is in the **Stopped** state, click **Start** to restart the compute instance.

----End

8.8.1.3 Configuring Cross-Manager Mutual Trust Between Clusters

Scenario

When two security-mode clusters managed by different FusionInsight Managers need to access each other's resources, the system administrator can configure cross-Manager mutual trust for them.

The secure usage scope of users in each system is called a domain. Each FusionInsight Manager must have a unique domain name. Cross-Manager access allows users to use resources across domains.

NOTE

A maximum of 500 mutually trusted clusters can be configured.

Impact on the System

- After cross-Manager cluster mutual trust is configured, users of an external system can be used in the local system. The system administrator needs to periodically check the user permissions in Manager based on enterprise service and security requirements.
- If you set up mutual trust between clusters, affected services need to be restarted and services will be interrupted.
- After cross-Manager cluster mutual trust is configured, internal Kerberos users **krbtgt/***Local cluster domain name*@*External cluster domain name* and **krbtgt/***External cluster domain name*@*Local cluster domain name* are added to the two mutually trusted clusters. The internal users cannot be deleted. The system administrator needs to change the passwords periodically based on enterprise service and security requirements. The passwords of these four users in the two systems must be the same. For details, see **Changing the Password for a Component Running User**. When the passwords are changed, the connectivity between cross-cluster service applications may be affected.
- If the system domain name is changed and there is any running HetuEngine compute instance, restart the compute instance.
- After cross-Manager cluster mutual trust is configured, the clients of each cluster need to be downloaded and installed again.
- After cross-Manager cluster mutual trust is configured, you need to check whether the system works properly and how to access resources of the peer system as a user of the local system. For details, see **Assigning User Permissions After Cross-Cluster Mutual Trust Is Configured**.

Prerequisites

- The system administrator has clarified service requirements and planned domain names for the systems. A domain name can contain only uppercase letters, numbers, periods (.), and underscores (_), and must start with a letter or number, for example, **DOMAINA.HW** and **DOMAINB.HW**.
- The domain names of the two Managers are different. When an ECS or BMS cluster is created on MRS, a unique system domain name is randomly generated. Generally, you do not need to change the system domain name.
- The two clusters do not have the same host name or the same IP address.
- The system time of the two clusters is consistent, and the NTP services in the two systems use the same clock source.
- The running status of all components in the Manager clusters is **Normal**.
- The **acl.compare.shortName** parameter of the ZooKeeper service of all clusters in Manager is set to default value **true**. Otherwise, change the value to **true** and restart the ZooKeeper service.

- The two clusters are in the same VPC. If they are not, create a VPC peering connection between them. For details, see **VPC Peering Connection**.

Procedure

Step 1 Log in to a FusionInsight Manager of one cluster.

Step 2 Choose **System** > **Permission** > **Domain and Mutual Trust**.

Step 3 Modify **Peer Mutual Trust Domain**.

Table 8-54 Related parameters

Parameter	Description
realm_name	Enter the domain name of the peer system.
ip_port	Enter the KDC address of the peer system. Value format: *IP address of the node accommodating the Kerberos service in the peer system:Port number* • In dual-plane networking, enter the service plane IP address. • If an IPv6 address is used, the IP address must be enclosed in square brackets ([]). • Use commas (,) to separate the KDC addresses if the active and standby Kerberos services are deployed or multiple clusters in the peer system need to establish mutual trust with the local system. • You can obtain the port number from the **kdc_ports** parameter of the KrbServer service. The default value is **21732**. To obtain the IP address of the node where the service is deployed, click the **Instance** tab on the KrbServer page and view **Service IP Address** of the KerberosServer role. For example, if the Kerberos service is deployed on nodes at **10.0.0.1** and **10.0.0.2** that have established mutual trust with the local system, the parameter value is **10.0.0.1:21732,10.0.0.2:21732**.

NOTE

If you need to configure mutual trust for multiple Managers, click + to add a new item and set parameters. Click — to delete unnecessary configurations.

Step 4 Click **OK**.

Step 5 Log in to the active management node as user **omm**, and run the following command to update the domain configuration:

sh ${BIGDATA_HOME}/om-server/om/sbin/restart-RealmConfig.sh

The command is executed successfully if the following information is displayed:


```
Modify realm successfully. Use the new password to log in to FusionInsight again.
```

After the restart, some hosts and services cannot be accessed and an alarm is generated. This problem can be automatically resolved in about 1 minute after **restart-RealmConfig.sh** is run.

Step 6 Log in to FusionInsight Manager and restart the cluster or configure expired instances:

Check whether the system domain name of Manager is changed.

- If the system domain name is changed, click or **More** on the home page, click **Start**, enter the password, select the checkbox for confirming the impact, and click **OK**. Wait until the cluster is started successfully.
- If the system domain name is not changed, click or **More** on the home page and select **Restart Configuration-Expired Instances**. Enter the password, select the checkbox for confirming the impact, and click **OK**. Wait until the service is restarted.

Step 7 Log out of FusionInsight Manager and then log in again. If the login is successful, the configuration is successful.

Step 8 If a HetuEngine compute instance is running, restart the compute instance.

1. Log in to FusionInsight Manager as the user who is used to access the HetuEngine web UI.
2. Choose **Cluster** > **Services** > **HetuEngine** to go to the HetuEngine service page.
3. In the **Basic Information** area on the **Dashboard** page, click the link next to **HSConsole WebUI**. The HSConsole page is displayed.
4. For a running compute instance, click **Stop** in the **Operation** column. After the compute instance is in the **Stopped** state, click **Start** to restart the compute instance.

Step 9 Log in to FusionInsight Manager of another cluster and repeat the preceding steps.

----End

8.8.1.4 Assigning User Permissions After Cross-Cluster Mutual Trust Is Configured

Scenario

After cross-Manager cluster mutual trust is configured, assign user access permissions on FusionInsight Managers so that these users can perform service operations in the mutually trusted Managers.

Prerequisites

The mutual trust between the two Managers has been configured.

Procedure

Step 1 Log in to the local FusionInsight Manager.

Step 2 Choose **System** > **Permission** > **User** to check whether the target user exists.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Click ⌄ on the left of the target user, and check whether the permissions assigned to the user group of the user and the roles meet service requirements. If not, create a role and bind the role to the user by referring to **Configuring Permissions**, or modify the user group or role permissions of the user.

Step 4 Create a user required by the service operations and associate the required user group or role. For details, see **Creating a User**.

Step 5 Log in to the other FusionInsight Manager and repeat **Step 2** to **Step 4** to create a user with the same name and set permissions.

----End

8.8.2 Configuring Scheduled Backup of Alarm and Audit Information

Scenario

You can modify the configuration file to periodically back up FusionInsight Manager alarm information, FusionInsight Manager audit information, and audit information of all services to the specified storage location.

The backup can be performed using FTP or SFTP. FTP does not encrypt data, which may cause security risks. Therefore, SFTP is recommended.

Procedure

Step 1 Log in to the active management node as user **omm**.

NOTE

Perform this operation only on the active management node. Scheduled backup is not supported on the standby management node.

Step 2 Run the following command to switch the directory:

cd ${BIGDATA_HOME}/om-server/om/sbin

Step 3 Run the following command to configure scheduled backup of FusionInsight Manager's alarm and audit information or service audit information:

./setNorthBound.sh -t *Information type* **-i** *Remote server IP address* **-p** *SFTP or FTP port used by the server*-**u** *Username* **-d** *Save path* **-c** *Interval (minutes)* **-m** *Number of records in each file* **-s** *Whether to enable backup* **-e** *Protocol*

Example:

./setNorthBound.sh -t alarm -i 10.0.0.10 -p 22 -u sftpuser -d /tmp/ -c 10 -m 100 -s true -e sftp

This script modifies the alarm backup configuration file **alarm_collect_upload.properties**. The file save path is **${BIGDATA_HOME}/om-server/tomcat/webapps/web/WEB-INF/classes/config**.

./setNorthBound.sh -t audit -i 10.0.0.10 -p 22 -u sftpuser -d /tmp/ -c 10 -m 100 -s true -e sftp

This script modifies the audit backup configuration file **audit_collect_upload.properties**. The file save path is **${BIGDATA_HOME}/om-server/tomcat/webapps/web/WEB-INF/classes/config**.

./setNorthBound.sh -t service_audit -i 10.0.0.10 -p 22 -u sftpuser -d /tmp/ -c 10 -m 100 -s true -e sftp

This script modifies the service audit backup configuration file **service_audit_collect_upload.properties**. The file save path is **${BIGDATA_HOME}/om-server/tomcat/webapps/web/WEB-INF/classes/config**.

Step 4 Enter the password as prompted. The password is encrypted and saved in the configuration file.

```
Please input sftp/ftp server password:
```

Step 5 Check the configuration result. If the following information is displayed, the configuration is successful. The configuration file will be automatically synchronized to the standby management node.

```
execute command syncfile successfully.
Config Succeed.
```

----End

8.8.3 Modifying the FusionInsight Manager Routing Table

Scenario

When FusionInsight Manager is installed, two pieces of routing information are automatically created on the active management node. You can run the **ip rule list** command to view the routing information, as shown in the following example:

```
0:from all lookup local
32764:from all to 10.10.100.100 lookup ntp_rt   #NTP routing information created by FusionInsight Manager (this information is unavailable if no external NTP clock source is configured).
32765:from 192.168.0.117 lookup om_rt   #OM routing information created by the FusionInsight Manager.
32766:from all lookup main
32767:from all lookup default
```

NOTE

If no external NTP server has been configured, only OM routing information **om_rt** will be created.

If the routing information created by FusionInsight Manager conflicts with the routing information configured in the enterprise network planning, the cluster administrator can use **autoroute.sh** to disable or enable the routing information created by FusionInsight Manager.

Impact on the System

After the routing information created by FusionInsight Manager is disabled and before the new routing information is set, FusionInsight Manager cannot be accessed but the clusters are running properly.

Prerequisites

FusionInsight Manager has been installed.

You have obtained routing information about the WS floating IP address.

Disable the Routing Information Created by the System

Step 1 Log in to the active management node as user **omm**. Run the following commands to disable the routing information created by the system:

cd ${BIGDATA_HOME}/om-server/om/sbin

./autoroute.sh disable

```
Deactivating Route.
Route operation (disable) successful.
```

Step 2 Run the following command to view the execution result:

ip rule list

```
0:from all lookup local
32766:from all lookup main
32767:from all lookup default
```

Step 3 Run the following command and enter the password of user **root** to switch to user **root**:

su - root

Step 4 Run the following commands to manually create the routing information about the WS floating IP address:

ip route add *Network segment of the WS floating IP address/Subnet mask of the WS floating IP address* **scope link src** *WS floating IP address* **dev** *NIC of the WS floating IP address* **table om_rt**

ip route add default via *Gateway of the WS floating IP address* **dev** *NIC of the WS floating IP address* **table om_rt**

ip rule add from *WS floating IP address* **table om_rt**

Example:

ip route add 192.168.0.0/255.255.255.0 scope link src 192.168.0.117 dev eth0:ws table om_rt

ip route add default via 192.168.0.254 dev eth0:ws table om_rt

ip rule add from 192.168.0.117 table om_rt

NOTE

If IPv6 addresses are used, run the **ip -6 route add** command.

Step 5 Run the following commands to manually create the NTP service routing information. Skip this step when no external NTP clock source is configured.

ip route add default via *IP gateway of the NTP service* **dev** *NIC of the local IP address* **table ntp_rt**

ip rule add to *ntpIP* **table ntp_rt**

NIC of the local IP address indicates the NIC that can communicate with the network segment where the NTP server is located.

Example:

ip route add default via 10.10.100.254 dev eth0 table ntp_rt

ip rule add to 10.10.100.100 table ntp_rt

Step 6 View the execution result.

In the following example, if the command output contains **om_rt** and **ntp_rt**, the operation is successful.

ip rule list

```
0:from all lookup local
32764:from all to 10.10.100.100 lookup ntp_rt  #This information is not displayed if no external NTP clock source is configured.
32765:from 192.168.0.117 lookup om_rt
32766:from all lookup main
32767:from all lookup default
```

----End

Enable the Routing Information Created by the System

Step 1 Log in to the active management node as user **omm**.

Step 2 Run the following commands to enable the routing information created by the system:

cd ${BIGDATA_HOME}/om-server/om/sbin

./autoroute.sh enable

```
Activating Route.
Route operation (enable) successful.
```

Step 3 View the execution result.

In the following example, if the command output contains **om_rt** and **ntp_rt**, the operation is successful.

ip rule list

```
0:from all lookup local
32764:from all to 10.10.100.100 lookup ntp_rt  #This information is not displayed if no external NTP clock source is configured.
32765:from 192.168.0.117 lookup om_rt
32766:from all lookup main
32767:from all lookup default
```

----End

8.8.4 Replacing the NTP Server for the Cluster

Scenario

After FusionInsight Manager is installed, if no NTP server is configured or the configured NTP server is no longer used, you can specify a new NTP server for the cluster or replace the NTP server with a new one to enable the cluster to synchronize time with the new NTP clock source.

Impact on the System

- Replacing the NTP server is a high-risk operation and may result in time change in the cluster.
- If the time difference between the NTP server and the cluster is greater than 150s before the NTP server replacement, you need to stop the cluster first to prevent data loss. Services are unavailable when the cluster is stopped.

Prerequisites

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

Procedure

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

- If yes, clear the alarm. After the alarm is cleared, go to **Step 2**.
- If no, go to **Step 2**.

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

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

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

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

- If yes, go to **Step 5**.
- If no, contact the network administrator to rectify the network fault. After the fault is rectified, go to **Step 5**.

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

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

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

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

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

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

In the preceding information, **+2.118107 sec** indicates the time offset. A positive value indicates that the NTP server time is earlier than the current cluster time. A negative value indicates the opposite.

- NOTE

  You can run the **ntpq -v** or **ntpq --version** command to query the NTP version. The command output may vary with the actual service environment.

  - Output of the **ntpq -v** command:
    ```
    10.1.1.112: ~# ntpq -v
    ntpq - standard NTP query program - Ver. 4.2.4p8
    ```
  - Output of the **ntpq --version** command:
    ```
    10.1.1.112: ~# ntpq --version
    ntpq 4.2.8p10@1.3728-o Mon Jun  6 08:01:59 UTC 2016 (1)
    ```

Step 7 Check whether the absolute value of the time difference exceeds **150**.

- If yes, go to **Step 8**.
- If no, perform **Step 10** as user **omm**.

Step 8 Check whether the cluster can be stopped.

- If yes, stop upper-layer services and the cluster, and go to **Step 9**.
- If no, no further action is required.

Step 9 Check whether the time of the NTP server is slower than the time of the cluster.

- If yes, wait a period of the time difference obtained in **Step 6** after message **Operation successful** is displayed on the UI, perform **Step 11** as user **omm**.
- If no, after message **Operation successful** is displayed on the UI, perform **Step 11** as user **omm**.

Step 10 Run the following command on the active management node to replace the NTP server:

sh ${BIGDATA_HOME}/om-server/om/bin/tools/modifyntp.sh --ntp_server_ip ntp.myhuaweicloud.com

NOTE

The IP address of the NTP server cannot be set to the IP address of a node in the cluster. Otherwise, the service network between the node and the active/standby OMS node may be disconnected.

Step 11 Run the following command on the active management node to forcibly synchronize time from the NTP server at **ntp.myhuaweicloud.com** immediately and replace the NTP server:

sh ${BIGDATA_HOME}/om-server/om/bin/tools/modifyntp.sh --ntp_server_ip ntp.myhuaweicloud.com --force_sync_time

NOTE

- If the cluster is stopped, start the cluster after the NTP server is replaced.
- After the command for forcible time synchronization is executed, it takes about five minutes for time synchronization on cluster nodes.

----End

8.8.5 Switching to the Maintenance Mode

Scenario

FusionInsight Manager allows you to set clusters, services, hosts, or OMSs to the maintenance mode. Objects in maintenance mode do not report alarms. This prevents the system from generating a large number of unnecessary alarms during maintenance changes, such as upgrade, because these alarms may influence O&M personnel's judgment on the cluster status.

- Cluster maintenance mode

  If a cluster is not brought online or has been brought offline due to O&M operations (for example, non-rolling upgrade), you can set the entire cluster to the maintenance mode.

- Service maintenance mode

  When performing maintenance operations on a specific service (for example, performing service-affecting commissioning operations like batch restart of service instances, directly powering on or off nodes of the service, or repairing the service), you can set only this service to the maintenance mode.

- Host maintenance mode

  When performing maintenance operations on a host (such as powering on or off, isolating, or reinstalling the host, upgrading its OS, or replacing the host), you can set only this host to the maintenance mode.

- OMS maintenance mode

  When restarting, replacing, or repairing an OMS node, you can set the OMS node to the maintenance mode.

Impact on the System

After the maintenance mode is set, alarms caused by non-maintenance operations are suppressed and cannot be reported. Alarms can be reported only when faults persist after the system exits the maintenance mode. Therefore, exercise caution when setting the maintenance mode.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Set the maintenance mode.

Determine the object to set the maintenance mode based on the service scenario. For details, see **Table 8-55**.

Table 8-55 Setting to the maintenance mode

Scenario	Operation
Configure a cluster to enter the maintenance mode.	1. On the management page, choose ··· or **More** > **Enter Maintenance**. 2. In the displayed dialog box, click **OK**. After the cluster enters the maintenance state, the status of the cluster becomes . After maintenance is complete, click **Exit Maintenance Mode**. The cluster then exits the maintenance mode.
Configure a service to enter the maintenance mode.	1. On FusionInsight Manager, choose **Cluster** > **Services** > *Service name*. 2. On the service details page, click **More** and select **Enter Maintenance Mode**. 3. In the displayed dialog box, click **OK**. After a service enters the maintenance mode, the status of the service becomes in the service list. After maintenance is complete, click **Exit Maintenance Mode**. The service then exits the maintenance mode. **NOTE** When configuring a service to enter the maintenance mode, you are advised to set the upper-layer services that depend on this service to the maintenance mode as well.
Configure a host to the maintenance mode.	1. On FusionInsight Manager, choose **Hosts**. 2. On the **Hosts** page, select the target host, click **More**, and select **Enter Maintenance Mode**. 3. In the displayed dialog box, click **OK**. After the host enters the maintenance mode, the status of the host becomes in the host list. After maintenance is complete, click **Exit Maintenance Mode**. The host then exits the maintenance mode.
Configure the OMS to enter the maintenance mode.	1. On FusionInsight Manager, choose **System** > **OMS** > **Enter Maintenance Mode**. 2. In the displayed dialog box, click **OK**. After the OMS enters the maintenance state, the OMS status becomes . After maintenance is complete, click **Exit Maintenance Mode**. The OMS then exits the maintenance mode.

Step 3 Check the cluster maintenance view.

On FusionInsight Manager, click ··· or **More** next to the cluster name and select **Maintenance Mode View**. In the displayed window, you can view the services and hosts in maintenance mode in the cluster.

After maintenance is complete, you can select services and hosts in batches in the maintenance mode view and click **Exit Maintenance Mode** to make them exit the maintenance mode.

Figure 8-56 Exiting the maintenance mode in batches

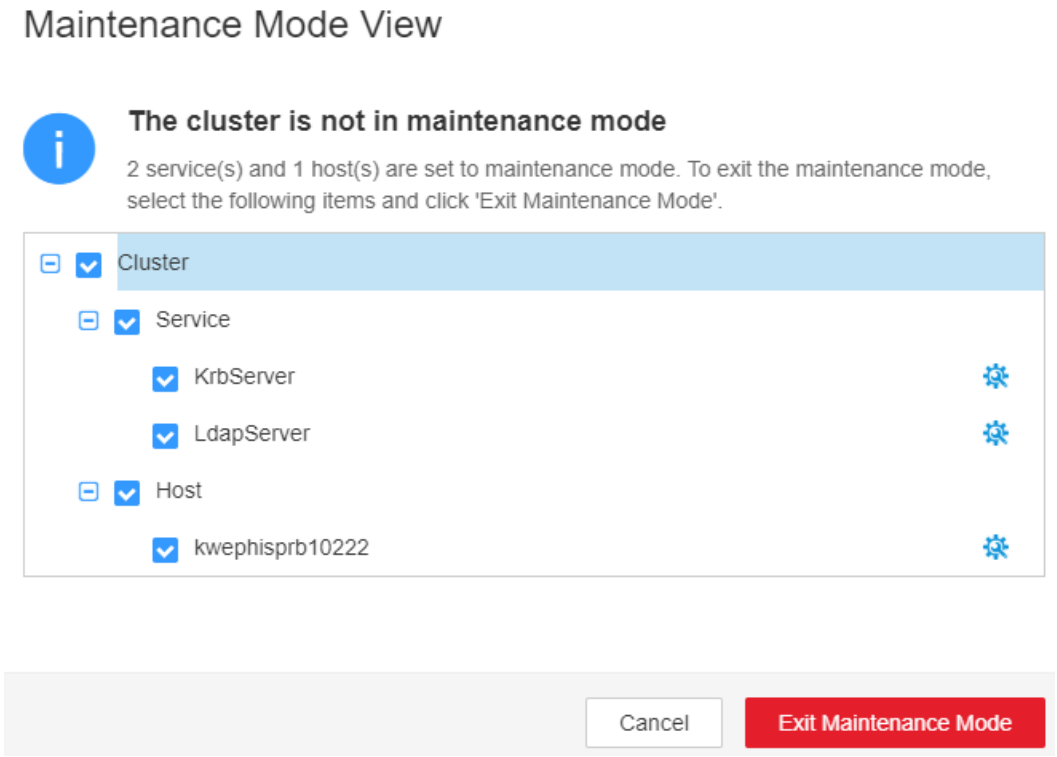

----End

8.8.6 Routine Maintenance of Manager

To ensure long-term and stable running of the system, administrators or maintenance engineers need to periodically check items listed in **Table 8-56** and rectify the detected faults based on the check results. It is recommended that administrators or engineers record the result in each task scenario and sign off based on the enterprise management regulations.

Table 8-56 Routine maintenance check items

Routine Maintenance Frequency	Task Scenario	Check Item
Daily	Check the cluster service status.	• Check whether the running status and configuration status of each service are normal and whether the status icons are green. • Check whether the running status and configuration status of the role instances in each service are normal and whether the status icons are green. • Check whether the active/standby status of role instances in each service can be properly displayed. • Check whether the dashboard of the services and role instances can be displayed properly.
	Check the cluster host status.	• Check whether the running status of each host is normal and whether the status icon is green. • Check the current disk usage, memory usage, and CPU usage of each host. Check whether the current memory usage and CPU usage are increasing.
	Check the cluster alarm information.	Check whether alarms were generated for unhandled exceptions on the previous day, including alarms that were automatically cleared.
	Check the cluster audit information.	Check whether critical and major operations are performed on the previous day and whether the operations are valid.
	Check the cluster backup status.	Check whether OMS, LDAP, DBService, and NameNode have been automatically backed up on the previous day.
	View the health check result.	Perform a health check on FusionInsight Manager and download the health check report to check whether the current cluster is abnormal. You are advised to enable the automatic health check, export the latest cluster health check result, and repair unhealthy items based on the result.

Routine Maintenance Frequency	Task Scenario	Check Item
	Check the network communication.	Check the cluster network status and check whether the network communication between nodes is delayed.
	Check the storage status.	Check whether the total data storage volume of the cluster increases abruptly. ● Check whether the disk usage is close to the threshold. If yes, locate the causes. For example, check whether the junk data or cold data left by services needs to be cleared. ● Check whether disk partitions need to be expanded based on the service growth trend.
	Check logs.	● Check whether there are failed or unresponsive MapReduce and Spark tasks. Check the **/tmp/logs/${username}/logs/${application id}** log file in HDFS and rectify faults. ● Check Yarn task logs, view the logs of failed and unresponsive tasks, and delete duplicate data. ● Check the worker logs of Storm. ● Back up logs to the storage server.
Weekly	Manage users.	Check whether the user password is about to expire and notify the user of changing the password. To change the password of a machine-machine user, you need to download the keytab file again.
	Analyze alarms.	Export and analyze alarms generated in a specified period.
	Scan disks.	Check the disk health status. You are advised to use a dedicated disk check tool.
	Collect statistics on storage.	Check in batches whether the disk data of cluster nodes is evenly stored, filter out the disks whose data increases significantly or is insufficient, and check whether the disks are normal.
	Record changes.	Arrange and record the operations on cluster configuration parameters and files to provide reference for fault analysis and handling.

Routine Maintenance Frequency	Task Scenario	Check Item
Monthly	Analyze logs.	• Collect and analyze hardware logs of cluster node servers, such as BMC system logs. • Collect and analyze the OS logs of the cluster node servers. • Collect and analyze cluster logs.
	Diagnose the network.	Analyze the network health status of the cluster.
	Manage hardware.	Check the equipment room environment and clean the devices.

8.9 Log Management

8.9.1 About Logs

Log Description

MRS cluster logs are stored in the **/var/log/Bigdata** directory. The following table lists the log types.

Table 8-57 Log types

Log Type	Description
Installation logs	Installation logs record information about FusionInsight Manager, cluster, and service installation to help users locate installation errors.
Run logs	Run logs record the running track information, debugging information, status changes, potential problems, and error information generated during the running of services.
Audit logs	Audit logs record information about users' activities and operation instructions, which can be used to locate fault causes in security events and determine who are responsible for these faults.

The following table lists the MRS log directories.

Table 8-58 Log directories

Directory	Log
/var/log/Bigdata/audit	Component audit log.
/var/log/Bigdata/controller	Log collecting script log. Controller process log. Controller monitoring log.
/var/log/Bigdata/dbservice	DBService log.
/var/log/Bigdata/flume	Flume log.
/var/log/Bigdata/hbase	HBase log.
/var/log/Bigdata/hdfs	HDFS log.
/var/log/Bigdata/hive	Hive log.
/var/log/Bigdata/hetuengine	HetuEngine log.
/var/log/Bigdata/httpd	HTTPd log.
/var/log/Bigdata/hue	Hue log.
/var/log/Bigdata/kerberos	Kerberos log.
/var/log/Bigdata/ldapclient	LDAP client log.
/var/log/Bigdata/ldapserver	LDAP server log.
/var/log/Bigdata/loader	Loader log.
/var/log/Bigdata/logman	Logman script log management log.
/var/log/Bigdata/mapreduce	MapReduce log.
/var/log/Bigdata/nodeagent	NodeAgent log.
/var/log/Bigdata/okerberos	OMS Kerberos log.
/var/log/Bigdata/oldapserver	OMS LDAP log.
/var/log/Bigdata/metric_agent	Run log file of MetricAgent.
/var/log/Bigdata/omm	**oms**: complex event processing log, alarm service log, HA log, authentication and authorization management log, and monitoring service run log of the OMM server. **oma**: installation log and run log of the OMM agent. **core**: dump log generated when the OMM agent and the HA process are suspended.
/var/log/Bigdata/spark2x	Spark2x log.

Directory	Log
/var/log/Bigdata/sudo	Log generated when the **sudo** command is executed by user **omm**.
/var/log/Bigdata/timestamp	Time synchronization management log.
/var/log/Bigdata/tomcat	Tomcat log.
/var/log/Bigdata/watchdog	Watchdog log.
/var/log/Bigdata/yarn	Yarn log.
/var/log/Bigdata/zookeeper	ZooKeeper log.
/var/log/Bigdata/oozie	Oozie log.
/var/log/Bigdata/kafka	Kafka log.
/var/log/Bigdata/storm	Storm log.
/var/log/Bigdata/iotdb	IoTDB log.
/var/log/Bigdata/cdl	CDL log.
/var/log/Bigdata/upgrade	OMS upgrade log.
/var/log/Bigdata/update-service	Upgrade service log.

Installation Logs

Table 8-59 Installation logs

Installation Log	Description
Configuration log	Records information about the configuration process before the installation.
FusionInsight Manager installation log	Records information about the two-node FusionInsight Manager installation.
Cluster installation log	Records information about the cluster installation.

Run Logs

Table 8-60 describes the running information recorded in run logs.

Table 8-60 Running information

Run Log	Description
Installation preparation log	Records information about preparations for the installation, such as the detection, configuration, and feedback operation information.
Process startup log	Records information about the commands executed during the process startup.
Process startup exception log	Records information about exceptions during process startup, such as dependent service errors and insufficient resources.
Process run log	Records information about the process running track information and debugging information, such as function entries and exits as well as cross-module interface messages.
Process running exception log	Records errors that cause process running errors, for example, the empty input objects or encoding or decoding failure.
Process running environment log	Records information about the process running environment, such as resource status and environment variables.
Script log	Records information about the script execution process.
Resource reclamation log	Records information about the resource reclaiming process.
Uninstallation clearing logs	Records information about operations performed during service uninstallation, such as directory and execution time deletion.

Audit Logs

Audit information recorded in audit logs includes FusionInsight Manager audit information and component audit information.

Table 8-61 Audit information of FusionInsight Manager

Operation Type	Operation
User management	Creating a user. Modifying a user. Deleting a user. Creating a user group. Modifying a user group. Deleting a group. Adding a role. Changing the user's roles. Deleting a role. Changing a password policy. Changing a password. Resetting a password. Logging in. Logging out. Unlocking the screen. Downloading the authentication credential. Unauthorized operation. Unlocking a user account. Locking a user account Locking the screen. Exporting a user. Exporting a user group. Exporting a role.

Operation Type	Operation
Cluster management	Starting a cluster. Stopping a cluster. Restarting a cluster. Performing a rolling restart of a cluster. Restarting all expired instances. Saving configurations. Synchronizing cluster configurations. Customizing cluster monitoring metrics. Configuring monitoring dumping. Saving monitoring thresholds. Downloading a client configuration file. Configuring the northbound Syslog interface. Configuring the northbound SNMP interface. Clearing alarms using SNMP. Adding a trap target using SNMP. Deleting a trap target using SNMP. Checking alarms using SNMP. Synchronizing alarms using SNMP. Creating a threshold template. Deleting a threshold template. Applying a threshold template. Saving cluster monitoring configurations. Exporting configurations. Importing cluster configurations. Exporting an installation template. Modifying a threshold template. Canceling the application of a threshold template. Masking an alarm. Sending an alarm. Changing the OMS database password. Resetting the component database password. Restarting OMM and Controller. Starting the health check of a cluster. Importing a certificate file. Configuring SSO information. Deleting historical health check reports. Modifying cluster properties.

Operation Type	Operation
	Running maintenance commands in synchronous mode. Running maintenance commands in asynchronous mode. Customizing report monitoring metrics. Exporting report monitoring data. Runing a command in asynchronous mode using SNMP. Restarting the Web service. Customizing monitoring metrics for static resource pools. Exporting monitoring data of a static resource pool. Customizing dashboard monitoring metrics. Stopping a task. Restoring configurations. Modifying domain and mutual trust configurations. Modifying system parameters. Making a cluster enter the maintenance mode. Making a cluster exit the maintenance mode. Making OMS enter the maintenance mode. Making OMS exit the maintenance mode. Making services in a cluster exit the maintenance mode in batches. Modifying OMS configurations. Enabling threshold alarms. Synchronizing all cluster configurations.

Operation Type	Operation
Service management	Starting a service. Stopping a service. Synchronizing service configurations. Refreshing a service queue. Customizing service monitoring metrics. Restarting a service. Performing a rolling service restart. Exporting service monitoring data. Importing service configuration data. Starting the health check of a service. Configuring a service. Uploading a configuration file. Downloading a configuration file. Synchronizing instance configurations. Commissioning an instance. Decommissioning an instance. Starting an instance. Stopping an instance. Customizing instance monitoring metrics. Restarting an instance. Performing a rolling restart of an instance. Exporting instance monitoring data. Importing instance configuration data. Creating an instance group. Modifying an instance group. Deleting an instance group. Moving an instance to another instance group. Making a service enter the maintenance mode. Making a service exit the maintenance mode. Changing the name of a service. Modifying service association. Downloading monitoring data. Masking alarms. Unmasking alarms. Exporting report data of a service. Adding custom parameters for a report. Modifying custom parameters of a report. Deleting custom parameters of a report.

Operation Type	Operation
	Switching over control nodes. Adding a mount table. Modifying a mount table.
Host management	Setting a node rack. Starting all roles. Stopping all roles. Isolating a host. Canceling isolation of a host. Customizing host monitoring metrics. Exporting host monitoring data. Making a host enter the maintenance mode. Making a host exit the maintenance mode. Exporting basic host information. Exporting host distribution report data. Exporting host trend report data. Exporting host cluster report data. Exporting report data of a service. Customizing host cluster monitoring metrics. Customizing host cluster trend monitoring metrics.
Alarm management	Exporting alarms. Clearing alarms. Exporting events. Clearing alarms in batches.
Log collection	Collecting log files. Downloading log files. Collecting service stack information. Collecting instance stack information. Preparing service stack information. Preparing instance stack information. Clearing service stack information. Clearing instance stack information.
Audit log management	Modifying audit dumping configurations. Exporting audit logs.

Operation Type	Operation
Data backup and restoration	Creating a backup task. Executing a backup task. Executing backup tasks in batches. Stopping a backup task. Deleting a backup task. Modifying a backup task. Locking a backup task. Unlocking a backup task. Creating a restoration task. Executing a restoration task. Stopping a restoration task. Retrying a restoration task. Deleting a restoration task.

Operation Type	Operation
Multi-tenant management	Saving static configurations. Adding a tenant. Deleting a tenant. Associating a service with a tenant. Deleting a service from a tenant. Configuring resources. Creating a resource. Deleting a resource. Adding a resource pool. Modifying a resource pool. Deleting a resource pool. Restoring tenant data. Modifying global configurations of a tenant. Modifying queue configurations of a capacity scheduler. Modifying queue configurations of a super scheduler. Modifying resource distribution of a capacity scheduler. Clearing resource distribution of a capacity scheduler. Modifying resource distribution of a super scheduler. Clearing resource distribution of a super scheduler. Adding a resource catalog. Modifying a resource catalog. Deleting a resource catalog. Customizing tenant monitoring metrics.

Operation Type	Operation
Health check	Starting the health check of a cluster. Starting the health check of a service. Starting the health check of a host. Starting the health check of OMS. Starting the system health check. Updating the health check configurations. Exporting health check reports. Exporting health check results of a cluster. Exporting health check results of a service. Exporting health check results of a host. Deleting historical health check reports. Exporting historical health check reports. Downloading a health check report.

Table 8-62 Component audit information

Audit Log	Operation Type	Operation
CDL audit log	Service operations	Creating a link. Deleting a link. Creating a job. Starting a Job. Deleting a job.
IoTDB audit log	Maintenance management	Granting permissions. Revoking permissions. Recording authentication and login information.
	Service operations	Deleting a time series, partition, function, or index. Modifying a time series.
ClickHouse audit log	Maintenance management	Granting permissions. Revoking permissions. Recording authentication and login information.
	Service operations	Creating databases or tables. Inserting, deleting, querying, and migrating data.
DBService audit log	Maintenance management	Performing backup restoration operations.

Audit Log	Operation Type	Operation
HBase audit log	Data definition language (DDL) statements	Creating a table. Deleting a table. Modifying a table. Adding a column family. Modifying a column family. Deleting a column family. Enabling a table. Disabling a table. Modifying user information. Changing a password. Logging in.
	Data manipulation language (DML) statements	Putting data (to the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables). Deleting data (from the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables). Checking and putting data (to the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables). Checking and deleting data (from the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables).
	Permission control	Assigning permissions to a user. Canceling permission assigning.
HDFS audit log	Permission management	Managing access permissions on files or folders. Managing the owner information of files or folders.
	File operations	Creating a folder. Creating a file. Opening a file. Appending file content. Changing a file name. Deleting a file or folder. Setting time property of a file. Setting the number of file copies. Merging files. Checking the file system. Linking to a file.

Audit Log	Operation Type	Operation
Hive audit log	Metadata operations	Defining metadata, such as creating databases and tables. Deleting metadata, such as deleting databases and tables. Modifying metadata, such as adding columns and renaming tables. Importing and exporting metadata.
	Data maintenance	Loading data to a table. Inserting data into a table.
	Permission management	Creating or deleting a role. Granting/Reclaiming roles. Granting/Reclaiming permissions.
Hue audit log	Service startup	Starting Hue.
	User operations	Logging in. Logging out.
	Task operations	Creating a task. Modifying a task. Deleting a task. Submitting a task. Saving a task. Updating the status of a task.
KrbServer audit log	Maintenance management	Changing the password of a Kerberos account Adding a Kerberos account Deleting a Kerberos account Authenticating users.
LdapServer audit log	Maintenance management	Adding an OS user. Adding a user group. Adding a user to a user group. Deleting a user. Deleting a group.
Loader audit log	Security management	Logging in.
	Metadata management	Querying connector information. Querying a framework. Querying step information.

Audit Log	Operation Type	Operation
	Data source connection management	Querying a data source connection. Adding a data source connection. Updating a data source connection. Deleting a data source connection. Activating a data source connection. Disabling a data source connection.
	Job management	Querying a job. Creating a job. Updating a job. Deleting a job. Activating a job. Disabling a job. Querying all execution records of a job. Querying the latest execution record of a job. Submitting a job. Stopping a job.
MapReduce audit log	Application running	Starting a container request. Stopping a container request. After a container request is complete, the status of the request becomes successful. After a container request is complete, the status of the request becomes failed. After a container request is complete, the status of the request becomes suspended. Submitting a task. Ending a task.
Oozie audit log	Task management	Submitting a task. Starting a task. Killing a task. Suspending a task. Resuming a task. Running a task again.

Audit Log	Operation Type	Operation
Spark2x audit log	Metadata operations	Defining metadata, such as creating databases and tables. Deleting metadata, such as deleting databases and tables. Modifying metadata, such as adding columns and renaming tables. Importing and exporting metadata.
	Data maintenance	Loading data to a table. Inserting data into a table.
Storm audit log	Nimbus operations	Submitting a topology. Stopping a topology. Reallocating a topology. Deactivating a topology. Activating a topology.
	UI operations	Stopping a topology. Reallocating a topology. Deactivating a topology. Activating a topology.
Yarn audit log	Job submission	Submitting a job to a queue.
ZooKeeper audit log	Permission management	Setting access permissions to Znode.
	Znode operations	Creating Znodes. Deleting Znodes. Configuring Znode data.
HetuEngine audit log	Job management	Adding an external data source. Deleting an external data source. Modifying an external data source. Creating a compute instance. Starting a compute instance. Stopping a compute instance. Deleting a compute instance. Querying a compute instance. Modifying compute instance configurations.

FusionInsight Manager audit logs are stored in the database. You can view and export the audit logs on the **Audit** page.

The following table lists the directories to store component audit logs. Audit log files of some components are stored in **/var/log/Bigdata/audit**, such as HDFS, HBase, MapReduce, Hive, Hue, Yarn, Storm, and ZooKeeper. The component audit logs are automatically compressed and backed up to **/var/log/Bigdata/audit/bk** at 03: 00 every day. A maximum of latest 90 compressed backup files are retained, and the backup time cannot be changed. For details about how to configure the number of reserved audit log files, see **Configuring the Number of Local Audit Log Backups**.

Audit log files of other components are stored in the component log directory.

Table 8-63 Directories for storing component audit logs

Component	Audit Log Directory
DBService	/var/log/Bigdata/audit/dbservice/dbservice_audit.log
HBase	/var/log/Bigdata/audit/hbase/hm/hbase-audit-hmaster.log /var/log/Bigdata/audit/hbase/hm/hbase-ranger-audit-hmaster.log /var/log/Bigdata/audit/hbase/rs/hbase-audit-regionserver.log /var/log/Bigdata/audit/hbase/rs/hbase-ranger-audit-regionserver.log /var/log/Bigdata/audit/hbase/rt/hbase-audit-restserver.log /var/log/Bigdata/audit/hbase/ts/hbase-audit-thriftserver.log
HDFS	/var/log/Bigdata/audit/hdfs/nn/hdfs-audit-namenode.log /var/log/Bigdata/audit/hdfs/nn/ranger-plugin-audit.log /var/log/Bigdata/audit/hdfs/dn/hdfs-audit-datanode.log /var/log/Bigdata/audit/hdfs/jn/hdfs-audit-journalnode.log /var/log/Bigdata/audit/hdfs/zkfc/hdfs-audit-zkfc.log /var/log/Bigdata/audit/hdfs/httpfs/hdfs-audit-httpfs.log /var/log/Bigdata/audit/hdfs/router/hdfs-audit-router.log
HetuEngine	/var/log/Bigdata/audit/hetuengine/hsbroker/hsbroker-audit.log.0 /var/log/Bigdata/audit/hetuengine/hsconsole/hsconsole-audit.log.0 /var/log/Bigdata/audit/hetuengine/hsfabric/hsfabric-audit.log.0 hdfs://hacluster/hetuserverhistory/*Tenant name*/coordinator/application_*ID*/container_*ID*/*yyyyMMdd*/hetuserver-engine-audit.log hdfs://hacluster/hetuserverhistory/*Tenant name*/coordinator or worker/application_*ID*/container_*ID*/*yyyyMMdd*/server.log
Hive	/var/log/Bigdata/audit/hive/hiveserver/hive-audit.log /var/log/Bigdata/audit/hive/hiveserver/hive-rangeraudit.log /var/log/Bigdata/audit/hive/metastore/metastore-audit.log /var/log/Bigdata/audit/hive/webhcat/webhcat-audit.log

Component	Audit Log Directory
Hue	/var/log/Bigdata/audit/hue/hue-audits.log
Kafka	/var/log/Bigdata/audit/kafka/audit.log
Loader	/var/log/Bigdata/loader/audit/default.audit
CDL	/var/log/Bigdata/audit/cdl/service/cdl-audit.log
MapReduce	/var/log/Bigdata/audit/mapreduce/jobhistory/mapred-audit-jobhistory.log
Oozie	/var/log/Bigdata/audit/oozie/oozie-audit.log
Spark2x	/var/log/Bigdata/audit/spark2x/jdbcserver/jdbcserver-audit.log /var/log/Bigdata/audit/spark2x/jdbcserver/ranger-audit.log /var/log/Bigdata/audit/spark2x/jobhistory/jobhistory-audit.log
Storm	/var/log/Bigdata/audit/storm/logviewer/audit.log /var/log/Bigdata/audit/storm/nimbus/audit.log /var/log/Bigdata/audit/storm/supervisor/audit.log /var/log/Bigdata/audit/storm/ui/audit.log
Yarn	/var/log/Bigdata/audit/yarn/rm/yarn-audit-resourcemanager.log /var/log/Bigdata/audit/yarn/rm/ranger-plugin-audit.log /var/log/Bigdata/audit/yarn/nm/yarn-audit-nodemanager.log
ZooKeeper	/var/log/Bigdata/audit/zookeeper/quorumpeer/zk-audit-quorumpeer.log
IoTDB	/var/log/Bigdata/audit/iotdb/iotdbserver/log_audit.log

8.9.2 Manager Log List

Log Description

Log path: The default storage path of Manager log files is **/var/log/Bigdata/***Manager component*.

- ControllerService: **/var/log/Bigdata/controller/** (OMS installation and run logs)
- HTTPd: **/var/log/Bigdata/httpd** (HTTPd installation and run logs)
- Logman: **/var/log/Bigdata/logman** (log packaging tool logs)
- NodeAgent: **/var/log/Bigdata/nodeagent** (NodeAgent installation and run logs)
- okerberos: **/var/log/Bigdata/okerberos** (okerberos installation and run logs)
- oldapserver: **/var/log/Bigdata/oldapserver** (oldapserver installation and run logs)

- MetricAgent: **/var/log/Bigdata/metric_agent** (MetricAgent run logs)
- OMM: **/var/log/Bigdata/omm** (OMM installation and run logs)
- Timestamp: **/var/log/Bigdata/timestamp** (NodeAgent startup time logs)
- Tomcat: **/var/log/Bigdata/tomcat** (Web process logs)
- Watchdog: **/var/log/Bigdata/watchdog** (watchdog logs)
- Upgrade: **/var/log/Bigdata/upgrade** (OMS upgrade logs)
- UpdateService: **/var/log/Bigdata/update-service** (upgrade service logs)
- Sudo: **/var/log/Bigdata/sudo** (sudo script execution logs)
- OS: **/var/log/***message file* (OS system logs)
- OS performance: **/var/log/osperf** (OS performance statistics logs)
- OS statistics: **/var/log/osinfo/statistics** (OS parameter configuration logs)

Log archive rule:

The automatic compression and archiving function is enabled for Manager logs. By default, when the size of a log file exceeds 10 MB, the log file is automatically compressed. The naming rule of a compressed log file is as follows: <*Original log name*>-<*yyyy-mm-dd_hh-mm-ss*>.[*ID*].**log.zip** A maximum of 20 latest compressed files are retained.

Table 8-64 Manager logs

Log Type	Log File Name	Description
Controller run logs	controller.log	Log file that records component installation, upgrade, configuration, monitoring, alarm reporting, and routine O&M operations
	controller_client.log	Run log file of the Representational State Transfer (REST) APIs
	acs.log	ACS run log file
	acs_spnego.log	spnego user logs in ACS
	aos.log	AOS run log file
	plugin.log	AOS plug-in logs
	backupplugin.log	Run log file that records the backup and restoration operations
	controller_config.log	Configuration run log file
	controller_nodesetup.log	Controller loading task log file

Log Type	Log File Name	Description
	controller_root.log	System log file of the Controller process
	controller_trace.log	Log file that records the remote procedure call (RPC) communication between Controller and NodeAgent
	controller_monitor.log	Monitoring log file
	controller_fsm.log	State machine log file
	controller_alarm.log	Controller alarm log file
	controller_backup.log	Controller backup and recovery log file
	install.log, restore_package.log, installPack.log, distributeAdapterFiles.log, and install_os_optimization.log	OMS installation log file
	oms_ctl.log	OMS startup and stop log file
	preInstall_client.log	Preprocessing log file before client installation
	installntp.log	NTP installation log file
	modify_manager_param.log	Manager parameter modification log file
	backup.log	OMS backup script run log file
	supressionAlarm.log	Alarm script run log file
	om.log	OM certificate generation log file
	backupplugin_ctl.log	Startup log file of the backup and restoration plug-in process
	getLogs.log	Run log of the log collection script
	backupAuditLogs.log	Run log of the audit log backup script

Log Type	Log File Name	Description
	certStatus.log	Log file that records regular certificate checks
	distribute.log	Certificate distribution log
	ficertgenetrate.log	Certificate replacement log file, covering level-2 certificates, CAS certificates, and HTTPd certificates
	genPwFile.log	Log file that records the generation of certificate password files
	modifyproxyconf.log	Log file that records the modification of the HTTPd proxy configuration
	importTar.log	Log file that records the process for importing certificates into the trust store.
HTTPd	install.log	HTTPd installation log file
	access_log, error_log	HTTPd run log file
Logman	logman.log	Log packaging tool log file
NodeAgent	install.log and install_os_optimization.log	NodeAgent installation log file
	installntp.log	NTP installation log file
	start_ntp.log	NTP startup log file
	ntpChecker.log	NTP check log file
	ntpMonitor.log	NTP monitoring log file
	heartbeat_trace.log	Log file that records heartbeats between NodeAgent and Controller
	alarm.log	Alarm log file
	monitor.log	Monitoring log file

Log Type	Log File Name	Description
	nodeagent_ctl.log and start-agent.log	NodeAgent startup log file
	agent.log	NodeAgent run log file
	cert.log	Certificate log file
	agentplugin.log	Log file that records the Agent plug-in running status
	omaplugin.log	OMA plug-in run log file
	diskhealth.log	Disk health check log file
	supressionAlarm.log	Alarm script run log file
	updateHostFile.log	Host list update log file
	collectLog.log	Run log file of the node log collection script
	host_metric_collect.log	Run log file of host metric collection
	checkfileconfig.log	Run log file of file permission check
	entropycheck.log	Entropy check run log file
	timer.log	Log file of scheduled node scheduling
	pluginmonitor.log	Component monitoring plug-in log file
	agent_alarm_py.log	Log file that records alarms upon insufficient NodeAgent file permission
oKerberos	addRealm.log and modifyKerberosRealm.log	Realm handover log file
	checkservice_detail.log	Okerberos health check log file
	genKeytab.log	keytab generation log file

Log Type	Log File Name	Description
	KerberosAdmin_genConfigDetail.log	Run log file of **kadmin.conf** generated during start of the kadmin process
	KerberosServer_genConfigDetail.log	Run log file of **krb5kdc.conf** generated during start of the krb5kdc process
	oms-kadmind.log	Run log file of the kadmin process
	oms_kerberos_install.log and postinstall_detail.log	Okerberos installation log file
	oms-krb5kdc.log	Run log file of the krbkdc process
	start_detail.log	Okerberos startup log file
	realmDataConfigProcess.log	Log file that records the rollback upon a realm handover failure
	stop_detail.log	Okerberos stop log file
oldapserver	ldapserver_backup.log	Oldapserver backup log file
	ldapserver_chk_service.log	Oldapserver health check log file
	ldapserver_install.log	Oldapserver installation log file
	ldapserver_start.log	Oldapserver startup log file
	ldapserver_status.log	Log file that records the status of the Oldapserver process
	ldapserver_stop.log	Oldapserver stop log file
	ldapserver_wrap.log	Oldapserver service management log file
	ldapserver_uninstall.log	Oldapserver uninstallation log file
	restart_service.log	Oldapserver restart log file

Log Type	Log File Name	Description
	ldapserver_unlockUser.log	Log file that records information about unlocking LDAP users and managing accounts
metric_agent	gc.log	MetricAgent JVM GC log file
	metric_agent.log	Run log file of MetricAgent
	metric_agent_qps.log	Log file that records MetricAgent Internal queue length and QPS information
	metric_agent_root.log	All run log files of MetricAgent
	start.log	Log file that records information about the MetricAgent startup and stop
OMM	omsconfig.log	OMS configuration log file
	check_oms_heartbeat.log	OMS heartbeat log file
	monitor.log	OMS monitoring log file
	ha_monitor.log	HA_Monitor operation log file
	ha.log	HA operation log file
	fms.log	Alarm log file
	fms_ha.log	HA alarm monitoring log file
	fms_script.log	Alarm control log file
	config.log	Alarm configuration log file
	iam.log	IAM log file
	iam_script.log	IAM control log file
	iam_ha.log	IAM HA monitoring log file
	config.log	IAM configuration log file

Log Type	Log File Name	Description
	operatelog.log	IAM operation log file
	heartbeatcheck_ha.log	OMS heartbeat HA monitoring log file
	install_oms.log	OMS installation log file
	pms_ha.log	HA monitoring log file
	pms_script.log	Monitoring control log file
	config.log	Monitoring configuration log file
	plugin.log	Monitoring plug-in run log file
	pms.log	Monitoring log file
	ha.log	HA run log file
	cep_ha.log	CEP HA monitoring log file
	cep_script.log	CEP control log file
	cep.log	CEP log file
	config.log	CEP configuration log file
	omm_gaussdba.log	GaussDB HA monitoring log file
	gaussdb-<SERIAL>.log	GaussDB run log file
	gs_ctl-<DATE>.log	Archive log file of GaussDB control logs
	gs_ctl-current.log	GaussDB control log file
	gs_guc-current.log	GaussDB operation log file
	encrypt.log	OMM encryption log file
	omm_agent_ctl.log	OMA control log file
	oma_monitor.log	OMA monitoring log file
	install_oma.log	OMA installation log file
	config_oma.log	OMA configuration log file

Log Type	Log File Name	Description
	omm_agent.log	OMA run log file
	acs.log	ACS resource log file
	aos.log	AOS resource log file
	controller.log	Controller resource log file
	floatip.log	Floating IP address resource log file
	ha_ntp.log	NTP resource log file
	httpd.log	HTTPd resource log file
	okerberos.log	Okerberos resource log file
	oldap.log	OLdap resource log file
	tomcat.log	Tomcat resource log file
	send_alarm.log	Run log file of the HA alarm sending script of the management node
	feed_watchdog.log	feed_watchdog resource log
Timestamp	restart_stamp	NodeAgent startup time log file
Tomcat	cas.log and localhost_access_cas_log.log	CAS run log file
	catalina.log, catalina.out, host-manager.log, localhost.log, and manager.log	Tomcat run log file
	localhost_access_web_log.log	Log file that records the access to REST APIs of FusionInsight Manager
	web.log	Run log file of the Web process
	northbound_ftp_sftp.log and snmp.log	Northbound log file
	perfStats.log	Performance statistics log file

Log Type	Log File Name	Description
Watchdog	watchdog.log and feed_watchdog.log	watchdog.log run log file
update-service	omm_upd_server.log	UPDServer run log file
	omm_upd_agent.log	UPDAgent run log file
	update-manager.log	UPDManager run log file
	install.log	Installation log file of the upgrade service
	uninstall.log	Uninstallation log file of the upgrade service
	catalina.<*Time*>.log, catalina.out, host-manager.<*Time*>.log, localhost.<*Time*>.log, manager.<*Time*>.log, manager_access_log.<*Time*>.txt, web_service_access_log.<*Time*>.txt, catalina.log, gc-update-service.log.0.current, update-manager.controller, update-web-service.controller, update-web-service.log, commit_rm_distributed.log, commit_rm_upload_package.log, common_omagent_operator.log, forbid_monitor.log, initialize_package_atoms.log, initialize_unzip_pack.log, omm-upd.log, register_patch_pack.log, resume_monitor.logrollback_clear_patch.log, unregister_patch_pack.log, update-rcommupd.log, update-rcupdatemanager.log, and update-service.log	Run log file of the upgrade service
Upgrade	upgrade.log_<*Time*>	OMS upgrade log file

Log Type	Log File Name	Description
	rollback.log_<*Time*>	OMS rollback log file
sudo	sudo.log	Sudo script execution log file

Log Levels

Table 8-65 describes the log levels provided by Manager. The log levels are FATAL, ERROR, WARN, INFO, and DEBUG in descending order. Logs whose levels are higher than or equal to the set level are printed by the program. The number of printed logs decreases as the set log level increases.

Table 8-65 Log levels

Level	Description
FATAL	Logs of this level record fatal error information about the current event processing that may result in a system crash.
ERROR	Logs of this level record error information about the current event processing, which indicates that system running is abnormal.
WARN	Logs of this level record abnormal information about the current event processing. These abnormalities will not result in system faults.
INFO	Logs of this level record normal running status information about the system and events.
DEBUG	Logs of this level record system information and debugging information.

Log Formats

The following table lists the Manager log formats.

Table 8-66 Log formats

Log Type	Component	Format	Example
Controller, HTTPd, Logman, NodeAgent, oKerberos, oldapserver, OMM, Tomcat, and upgrade	Controller, HTTPd, Logman, NodeAgent, oKerberos, oldapserver, OMM, Tomcat, and upgrade	<yyyy-MM-dd HH:mm:ss, SSS>|<Log Level>|<*Name of the thread for which the log is generated*>|<*Log message*>|<*Location where the log event occurs*>	2020-06-30 00:37:09,067 INFO [pool-1-thread-1] Completed Discovering Node. com.xxx.hadoop.om.controller.tasks.nodesetup.DiscoverNodeTask.execute(DiscoverNodeTask.java:299)

8.9.3 Configuring the Log Level and Log File Size

Scenario

You can change the log levels of FusionInsight Manager. For a specific service, you can change the log level and the log file size to prevent the failure in saving logs due to insufficient disk space.

Impact on the System

The services need to be restarted for the new configuration to take effect. During the restart, the services are unavailable.

Changing the FusionInsight Manager Log Level

1. Log in to the active management node as user **omm**.
2. Run the following command to switch to the required directory:

   cd ${BIGDATA_HOME}/om-server/om/sbin
3. Run the following command to change the log level:

   ./setLogLevel.sh*Log level parameters*

   The priorities of log levels are FATAL, ERROR, WARN, INFO, and DEBUG in descending order. Logs whose levels are higher than or equal to the set level are printed. The number of printed logs decreases as the configured log level increases.

   - **DEFAULT**: After this parameter is set, the default log level is used.
   - **FATAL**: critical error log level. After this parameter is set, only logs of the **FATAL** level are printed.
   - **ERROR**: error log level. After this parameter is set, logs of the **ERROR** and **FATAL** levels are printed.
   - **WARN**: warning log level. After this parameter is set, logs of the **WARN**, **ERROR**, and **FATAL** levels are printed.

- **INFO** (default): informational log level. After this parameter is set, logs of the **INFO**, **WARN**, **ERROR**, and **FATAL** levels are printed.
- **DEBUG**: debugging log level. After this parameter is set, logs of the **DEBUG**, **INFO**, **WARN**, **ERROR**, and **FATAL** levels are printed.
- **TRACE**: tracing log level. After this parameter is set, logs of the **TRACE**, **DEBUG**, **INFO**, **WARN**, **ERROR**, and **FATAL** levels are printed.

NOTE

The log levels of components are different from those defined in open-source code.

4. Download and view logs to verify that the log level settings have taken effect. For details, see **Log**.

Changing the Service Log Level and Log File Size

NOTE

KrbServer, LdapServer, and DBService do not support the changing of service log levels and log file sizes.

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > **Services**.

Step 3 Click a service in the service list. On the displayed page, click the **Configuration** page.

Step 4 On the displayed page, click the **All Configuration** tab. Expand the role instance displayed on the left of the page. Click **Log** of the role to be modified.

Step 5 Search for each parameter and obtain the parameter description. On the parameter configuration page, select the required log level or change the log file size. The unit of the log file size is MB.

NOTICE

- The system automatically deletes logs based on the configured log size. To save more information, set the log file size to a larger value. To ensure the integrity of log files, you are advised to manually back up the log files to another directory based on the actual service volume before the log files are cleared according to clearance rules.
- Some services do not support change of the log level on the UI.

Step 6 Click **Save**. In the **Save Configuration** dialog box, click **OK**.

Step 7 Download and view logs to verify that the log level settings have taken effect.

----End

8.9.4 Configuring the Number of Local Audit Log Backups

Scenario

Audit logs of cluster components are classified by name and stored in the **/var/log/Bigdata/audit** directory on each cluster node. The OMS automatically backs up the audit log directories at 03:00 every day.

The audit log directory on each node is compressed and named in the *<Node IP address>***.tar.gz** format. All compressed files are compressed and named in the *<yyyy-MM-dd_HH-mm-ss>***.tar.gz** format and saved in the **/var/log/Bigdata/audit/bk/** directory on the active management node. In addition, the standby management node saves a copy of the file.

By default, a maximum of 90 OMS backup files can be retained. This section describes how to configure the maximum number.

Procedure

Step 1 Log in to the active management node as user **omm**.

NOTE

Perform this operation only on the active management node. This operation is not supported on the standby management nodes; otherwise, the cluster cannot work properly.

Step 2 Run the following command to switch to the required directory:

cd ${BIGDATA_HOME}/om-server/om/sbin

Step 3 Run the following command to change the maximum number of audit log backup files to be retained:

./modifyLogConfig.sh -m *Maximum number of backup files that can be retained*

The default value is **90**. The value ranges from **0** to **365**. A larger value means to consume more disk space.

If the following information is displayed, the operation is successful:

```
Modify log config successfully
```

----End

8.9.5 Viewing Role Instance Logs

Scenario

FusionInsight Manager allows users to view logs of each role instance.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > **Services**, and click a service name. Then click the **Instances** tab of the service and click the name of the target instance to access the instance status page.

Step 3 In the **Log** area, click the name of a log file to preview its content online.

NOTE

- On the **Hosts** page, click a host name. In the instance list of the host, you can view the log files of all role instances on the host.
- By default, a maximum of 100 lines of logs can be displayed. You can click **Load More** to view more logs. Click **Download** to download the log file to the local PC. For details about how to download service logs in batches, see **Log Download**.

Figure 8-57 Viewing instance logs

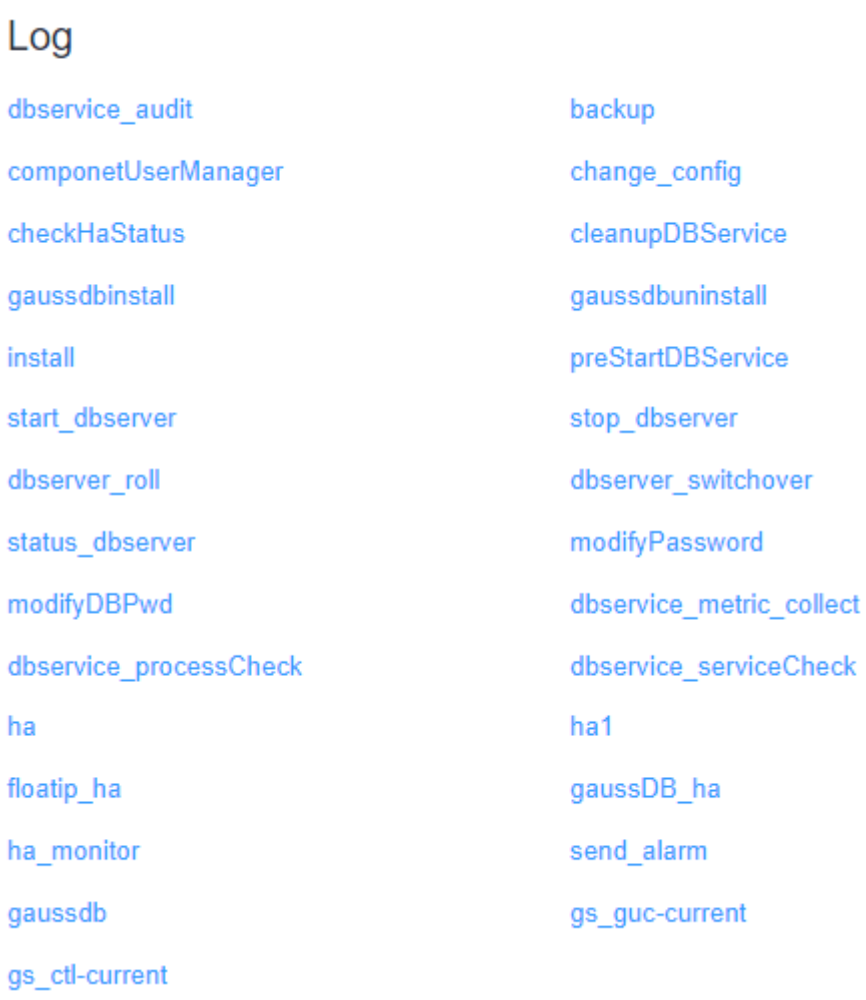

----End

8.10 Backup and Recovery Management

8.10.1 Introduction

Overview

FusionInsight Manager provides the backup and restoration of system data and user data by component. The system can back up Manager data, component metadata, and service data.

Data can be backed up to local disks (LocalDir), local HDFS (LocalHDFS), remote HDFS (RemoteHDFS), NAS (NFS/CIFS), Object Storage Service (OBS), and SFTP server (SFTP). For details, see **Backing Up Data**.

For a component that supports multiple services, multiple instances of a service can be backed up and restored. The backup and restoration operations are consistent with those of a service instance.

NOTE

Only MRS 3.1.0 or later supports data backup to OBS.

Backup and restoration tasks are performed in the following scenarios:

- Routine backup is performed to ensure the data security of the system and components.
- If the system is faulty, the data backup can be used to recover the system.
- If the active cluster is completely faulty, a mirrored cluster identical to the active cluster needs to be created. You can use the backup data to restore the active cluster.

Table 8-67 Manager configuration data to be backed up

Backup Type	Backup Content	Backup Directory Type
OMS	Database data (excluding alarm data) and configuration data in the cluster management system by default	• LocalDir • LocalHDFS • RemoteHDFS • NFS • CIFS • SFTP • OBS

Table 8-68 Component metadata or other data to be backed up

Backup Type	Backup Content	Backup Directory Type
DBService	Metadata of the components (including Loader, Hive, Spark, Oozie, CDL, and Hue) managed by DBService. For a cluster with multiple services installed, back up the metadata of multiple Hive and Spark service instances.	• LocalDir • LocalHDFS • RemoteHDFS • NFS • CIFS • SFTP • OBS
Flink (Applicable to MRS 3.2.0 and later versions)	Flink metadata.	• LocalDir • LocalHDFS • RemoteHDFS
Kafka	Kafka metadata.	• LocalDir • LocalHDFS • RemoteHDFS • NFS • CIFS • OBS
NameNode	HDFS metadata. After multiple NameServices are added, backup and restoration are supported for all of them and the operations are consistent with those of the default hacluster instance.	• LocalDir • RemoteHDFS • NFS • CIFS • SFTP • OBS
Yarn	Information about the Yarn service resource pool.	
HBase	**tableinfo** files and data files of HBase system tables.	
IoTDB	IoTDB metadata.	• LocalDir • NFS • RemoteHDFS • CIFS • SFTP
ClickHouse	ClickHouse metadata.	• LocalDir • RemoteHDFS

Table 8-69 Service data of specific components to be backed up

Backup Type	Backup Content	Backup Directory Type
HBase	Table-level user data. For a cluster with multiple services installed, backup and restoration are supported for multiple HBase service instances and the backup and restoration operations are consistent with those of a single HBase service instance.	• RemoteHDFS • NFS • CIFS • SFTP
HDFS	Directories or files of user services. **NOTE** Encrypted directories cannot be backed up or restored.	
Hive	Table-level user data. For a cluster with multiple services installed, backup and restoration are supported for multiple Hive service instances and the backup and restoration operations are consistent with those of a single Hive service instance.	
IoTDB	IoTDB service data.	• RemoteHDFS
ClickHouse	Table-level user data.	• RemoteHDFS

Note that some components do not provide data backup or restoration:

- Kafka supports replicas and allows multiple replicas to be specified when a topic is created.
- CDL data is stored in DBService and Kafka. A system administrator can create DBService and Kafka backup tasks to back up data.
- MapReduce and Yarn data is stored in HDFS. Therefore, they rely on the backup and restoration provided by HDFS.
- Backup and restoration of service data in ZooKeeper are performed by their own upper-layer components.

Principles

Task

Before backup or restoration, you need to create a backup or restoration task and set task parameters, such as the task name, backup data source, and type of the directory for storing backup files. Then you can execute the tasks to back up or restore data. When Manager is used to restore the data of HDFS, HBase, Hive, and NameNode, the cluster cannot be accessed.

Each backup task can back up data of different data sources and generate an independent backup file for each data source. All the backup files generated in a

backup task form a backup file set, which can be used in restoration tasks. Backup data can be stored on Linux local disks, local cluster HDFS, and standby cluster HDFS.

Backup tasks support full backup and incremental backup policies. Cloud data backup tasks do not support incremental backup. If the backup directory type is NFS or CIFS, incremental backup is not recommended. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

NOTE

Task execution rules:

- If a task is being executed, the task cannot be executed repeatedly and other tasks cannot be started at the same time.
- The interval at which a periodic task is automatically executed must be greater than 120s. Otherwise, the task is postponed and will be executed in the next period. Manual tasks can be executed at any interval.
- When a periodic task is to be automatically executed, the current time cannot be 120s later than the task start time. Otherwise, the task is postponed and executed in the next period.
- When a periodic task is locked, it cannot be automatically executed and needs to be manually unlocked.
- Before an OMS, DBService, Kafka, or NameNode backup task starts, ensure that the LocalBackup partition on the active management node has not less than 20 GB of available space. Otherwise, the backup task cannot be started.

When planning backup and restoration tasks, select the data to be backed up or restored strictly based on the service logic, data store structure, and database or table association. By default, the system creates periodic backup tasks **default-oms** and **default-*cluster ID*** at an interval of one hour. OMS metadata and cluster metadata, such as DBService and NameNode, can be fully backed up to local disks.

Snapshot

The system uses the snapshot technology to quickly back up data. Snapshots include HBase and HDFS snapshots.

- HBase snapshots

  An HBase snapshot is a backup file of HBase tables at a specified time point. This backup file does not replicate service data or affect the RegionServer. The HBase snapshot replicates table metadata, including table descriptor, region info, and HFile reference information. The metadata can be used to restore data before the snapshot creation time.

- HDFS snapshots

  An HDFS snapshot is a read-only backup of HDFS at a specified time point. The snapshot is used in data backup, misoperation protection, and disaster recovery scenarios.

  The snapshot function can be enabled for any HDFS directory to create the related snapshot file. Before creating a snapshot for a directory, the system automatically enables the snapshot function for the directory. Creating a snapshot does not affect any HDFS operation. A maximum of 65,536 snapshots can be created for each HDFS directory.

When a snapshot is being created for an HDFS directory, the directory cannot be deleted or modified before the snapshot is created. Snapshots cannot be created for the upper-layer directories or subdirectories of the directory.

DistCp

Distributed copy (DistCp) is a tool used to replicate a large amount of data in HDFS in a cluster or between the HDFSs of different clusters. In a backup or restoration task of HBase, HDFS, or Hive, if you back up the data to HDFS of the standby cluster, the system invokes DistCp to perform the operation. Install the MRS software of the same version for the active and standby clusters and install the cluster.

DistCp uses MapReduce to implement data distribution, troubleshooting, restoration, and report. DistCp specifies different Map jobs for various source files and directories in the specified list. Each Map job copies the data in the partition that corresponds to the specified file in the list.

If you use DistCp to replicate data between HDFSs of two clusters, configure the cross-cluster mutual trust (mutual trust does not need to be configured for clusters managed by the same FusionInsight Manager) and cross-cluster replication for both clusters. When backing up the cluster data to HDFS in another cluster, you need to install the Yarn component. Otherwise, the backup fails.

Local rapid restoration

After using DistCp to back up the HBase, HDFS, and Hive data of the local cluster to the HDFS of the standby cluster, the HDFS of the local cluster retains the backup data snapshots. You can create local rapid restoration tasks to restore data by using the snapshot files in the HDFS of the local cluster.

NAS

Network Attached Storage (NAS) is a dedicated data storage server which includes the storage components and embedded system software. It provides the cross-platform file sharing function. By using NFS (supporting NFSv3 and NFSv4) and CIFS (supporting SMBv2 and SMBv3), you can connect the service plane of MRS to the NAS server to back up data to the NAS or restore data from the NAS.

NOTE

- Before data is backed up to the NAS, the system automatically mounts the NAS shared address to a local partition of the backup task execution node. After the backup is complete, the system unmounts the NAS shared partition from the backup task execution node.
- To prevent backup and restoration failures, do not access the shared address where the NAS server has been mounted to, for example, **/srv/BigData/LocalBackup/nas**, during data backup and restoration.
- When service data is backed up to the NAS, DistCp is used.

Specifications

Table 8-70 Specifications of the backup and restoration feature

Item	Specification
Maximum number of backup or restoration tasks	100
Number of concurrent tasks in a cluster	1
Maximum number of waiting tasks	199
Maximum size (GB) of backup files on a Linux local disk	600

NOTE

If service data is stored in the ZooKeeper upper-layer components, ensure that the number of znodes in a single backup or restoration task is not too large. Otherwise, the task will fail, and the ZooKeeper service performance will be affected. To check the number of znodes in a single backup or restoration task, perform the following operations:

- Ensure that the number of znodes in a single backup or restoration task is smaller than the upper limit of OS file handles. Specifically:
  1. To check the upper limit at the system level, run the **cat /proc/sys/fs/file-max** command.
  2. To check the upper limit at the user level, run the **ulimit -n** command.
- If the number of znodes in the parent directory exceeds the upper limit, back up and restore data in its sub-directories in batches. To check the number of znodes using ZooKeeper client scripts, perform the following operations:
  1. On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services** > **ZooKeeper** > **Instance**, and view the management IP address of each ZooKeeper role.
  2. Log in to the node where the client is located and run the following command:

     zkCli.sh -server *ip*:*port*, where, *ip* can be any management IP address, and the default port number is 2181.
  3. If the following information is displayed, login to the ZooKeeper server is successful:
     ```
     WatchedEvent state:SyncConnected type:None path:null
     [zk: ip:port(CONNECIED) 0]
     ```
  4. Run the **getusage** command to check the number of znodes in the directory to be backed up.

     For example, **getusage /hbase/region**. In the command output, **Node count=xxxxxx** indicates the number of znodes stored in the **region** directory.

Table 8-71 Specifications of the default task

Item	OMS	HBase	Kafka	DBService	NameNode
Backup period	1 hour				

Item	OMS	HBase	Kafka	DBService	NameNode
Maximum number of backups	168 (7-day historical data)				24 (one-day historical data)
Maximum size of a backup file	10 MB	10 MB	512 MB	100 MB	20 GB
Maximum size of disk space used	1.64 GB	1.64 GB	84 GB	16.41 GB	480 GB
Storage path of backup data	*Data storage path*/**LocalBackup/** of the active and standby management nodes				

NOTE

- The backup data of the default backup task must be periodically transferred and saved outside the cluster based on the enterprise O&M requirements.
- Administrators can create DistCp backup tasks to save OMS, DBService, and NameNode data to external clusters.
- The execution time of a cluster data backup task can be calculated using the following formula: Task execution time = Volume of data to be backed up/Network bandwidth between the cluster and the backup device. In practice, you are advised to multiply the calculated time by 1.5 to get the reference value of the task execution time.
- Executing a data backup task affects the maximum I/O performance of the cluster. Therefore, you are advised to execute a backup task during off-peak hours.

8.10.2 Backing Up Data

8.10.2.1 Backing Up Manager Data

Scenario

To ensure data security of FusionInsight Manager routinely or before and after a critical operation (such as capacity expansion and reduction) on FusionInsight Manager, you need to back up FusionInsight Manager data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Manager data. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.

- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Figure 8-58 Creating a backup task.

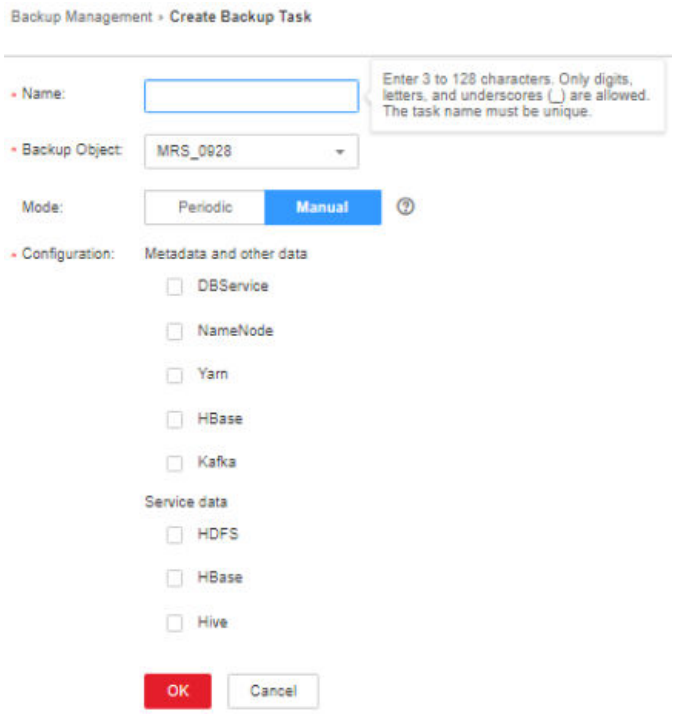

Step 3 Set **Name** to the name of the backup task.

Step 4 Set **Backup Object** to **OMS**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-72 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, select **OMS**.

Step 7 Set **Path Type** of **OMS** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

  The default storage directory is *Data storage path*/**LocalBackup/**, for example, **/srv/BigData/LocalBackup**.

  If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.
- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select this option, set the following parameters:

  – **Target Path**: indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.

  – **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

  – **Cluster for Backup**: Enter the cluster name mapping to the backup directory.

  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select this option, set the following parameters:

  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
  - **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Source Cluster**: Select the cluster of the Yarn queue used by the backup data.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.

  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **OBS**: indicates that backup files are stored in OBS.

  If you select this option, set the following parameters:
  - **Target Path**: indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

    NOTE

    Only MRS 3.1.0 or later supports data backup to OBS.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files.

The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

8.10.2.2 Backing Up CDL Data

Scenario

To ensure CDL service data security routinely or before a major operation on CDL (such as upgrade or migration), you need to back up CDL data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

CDL data is stored in DBService and Kafka. You can create DBService and Kafka backup tasks on FusionInsight Manager to back up CDL data. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-73 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 Set **Configuration** to **DBService** and **Kafka**.

Step 7 Set **Path Type** of **DBService** to a backup directory type. For details about how to set the parameters, see **Step 7**.

Step 8 Set **Path Type** of **Kafka** to a backup directory type. For details about how to set the parameters, see **Step 7**.

Step 9 Click **OK**.

Step 10 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files.

The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

8.10.2.3 Backing Up ClickHouse Metadata

Scenario

To ensure ClickHouse metadata security or before a major operation (such as upgrade or migration), you need to back up ClickHouse metadata. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up ClickHouse metadata. Both automatic and manual backup tasks are supported.

NOTICE

This function is supported only by MRS 3.1.0 or later.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If the active and standby clusters are deployed in security mode and they are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active and standby clusters are deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- In the active/standby cluster, if data is remotely backed up to HDFS, ensure that the value of **HADOOP_RPC_PROTECTION** of ClickHouse is the same as that of **hadoop.rpc.protection** of HDFS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task. **Periodic** indicates that the backup task is periodically executed. **Manual** indicates that the backup task is manually executed.

To create a periodic backup task, set the following parameters:

- **Started**: indicates the time when the task is started for the first time.
- **Period**: indicates the task execution interval. The options include **Hours** and **Days**.
- **Backup Policy**: Only **Full backup every time** is supported.

Step 6 In **Configuration**, select **ClickHouse** under **Metadata and other data**.

Step 7 Set **Path Type** of **ClickHouse** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

  The default storage directory is *Data storage path*/**LocalBackup/**, for example, **/srv/BigData/LocalBackup**.

  If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  This value option is available only after you configure the environment by referring to **How Do I Configure the Environment When I Create a ClickHouse Backup Task on FusionInsight Manager and Set the Path Type to RemoteHDFS?**.

  You also need to configure the following parameters:

  - **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
  - **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Data source_Task execution time***.tar.gz**.

----End

8.10.2.4 Backing Up ClickHouse Service Data

Scenario

To ensure ClickHouse service data security routinely or before a major operation on ClickHouse (such as upgrade or migration), you need to back up ClickHouse service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up ClickHouse service data. Both automatic and manual backup tasks are supported.

NOTICE

This function is supported only by MRS 3.1.0 or later.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active and standby clusters are deployed in security mode and they are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active and standby clusters are deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- You have planned the backup type, period, object, and directory based on service requirements.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
- In the active/standby cluster, if data is remotely backed up to HDFS, ensure that the value of **HADOOP_RPC_PROTECTION** of ClickHouse is the same as that of **hadoop.rpc.protection** of HDFS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-74 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported.

Step 6 In **Configuration**, select **ClickHouse** under **Service Data**.

Step 7 Set **Path Type** of **ClickHouse** to a backup directory type.

Currently, only the **RemoteHDFS** type is available.

RemoteHDFS: indicates that backup files are stored in HDFS of the standby cluster.

This value option is available for MRS 3.1.0 or 3.1.2 clusters only after you configure the environment by referring to **How Do I Configure the Environment When I Create a ClickHouse Backup Task on FusionInsight Manager and Set the Path Type to RemoteHDFS?**.

You also need to configure the following parameters for MRS 3.1.3 or later clusters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Destination Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
- **Destination Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Destination NameNode RPC Port**: indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the destination cluster.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a

snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.

You also need to configure the following parameters for MRS 3.1.0 or 3.1.2 clusters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

Step 8 Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

Step 9 Set **Backup Content** to one or multiple ClickHouse tables to be backed up.

You can select backup data using either of the following methods:

- Adding a backup data file

  Click the name of a database in the navigation tree to show all the tables in the database, and select specified tables.

  MRS 3.2.0 or later:

  a. Click **Add**.

  b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.

  c. Click **OK**.

- MRS 3.2.0 or later: Regular expression filtering

  a. Click **Query Regular Expression**.

  b. Enter the logical cluster and database to which the ClickHouse table belongs in the first text box as prompted. The logical cluster and database must match the existing logical cluster and database, for example, **/default_cluster/database**.

  c. Enter a regular expression in the second box. Standard regular expressions are supported. For example, to filter all tables that contain the keyword **test** in the database, enter **test.***.

  d. Click **Refresh** to view the displayed tables in **Directory Name**.

e. Click **Synchronize** to save the result.

NOTE

- When entering regular expressions, click + or − to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.

- Versions earlier than MRS 3.2.0: Regular expression filtering
  a. Click **Query Regular Expression**.
  b. Enter the database where the ClickHouse tables are located in the first text box as prompted. The database must be the same as the existing database, for example, **default**.
  c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all tables in the database, enter **([\s\S]*?)**. To get tables named in the format of letters and digits, for example, **tb1**, enter **tb\d***.
  d. Click **Refresh** to view the displayed tables in **Directory Name**.
  e. Click **Synchronize** to save the result.

NOTE

- When entering regular expressions, click + or − to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.

Step 10 Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The directory or table to be backed up does not exist.
- The name of the NameService is incorrect.

Step 11 Click **OK**.

Step 12 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Data source_Task creation time*, and the subdirectory is used to save latest data source backup files.

----End

8.10.2.5 Backing Up DBService Data

Scenario

To ensure DBService service data security routinely or before a major operation on DBService (such as upgrade or migration), you need to back up DBService data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up DBService data. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-75 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.

Parameter	Description
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, select **DBService**.

NOTE

If there are multiple DBService services, all DBService services are backed up by default. You can click **Assign Service** to specify the services to be backed up.

Step 7 Set **Path Type** of **DBService** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

  The default storage directory is *Data storage path*/**LocalBackup/**, for example, **/srv/BigData/LocalBackup**.

  If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.
- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select this option, set the following parameters:

  – **Target Path**: indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.

  – **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

  – **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select this option, set the following parameters:

  – **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**,

haclusterX1, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)

  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **OBS**: indicates that backup files are stored in OBS.

  If you select this option, set the following parameters:
  - **Target Path**: indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

    NOTE

    Only MRS 3.1.0 or later supports data backup to OBS.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files.

The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

8.10.2.6 Backing Up Flink Metadata

Scenario

To ensure Flink metadata security or before a major operation on Flink (such as upgrade or migration), you need to back up Flink metadata. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Flink metadata. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- The HDFS and Yarn services have been installed if data needs to be backed up to HDFS.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task. **Periodic** indicates that the backup task is periodically executed. **Manual** indicates that the backup task is manually executed.

To create a periodic backup task, set the following parameters:

- **Started**: indicates the time when the task is started for the first time.
- **Period**: indicates the task execution interval. The options include **Hours** and **Days**.
- **Backup Policy**: Only **Full backup every time** is supported.

Step 6 In **Configuration**, select **Flink** under **Metadata and other data**.

Step 7 Set **Path Type** of **Flink** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

The default storage directory is *Data storage path*/**LocalBackup/**, for example, **/srv/BigData/LocalBackup**.

If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select this option, set the following parameters:

  - **Target Path**: indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select this option, set the following parameters:

  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
  - **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data

source backup files. The format of the backup file name is *Data source*_*Task execution time*.tar.gz.

----End

8.10.2.7 Backing Up HBase Metadata

Scenario

To ensure HBase metadata security (including tableinfo files and HFiles) or before a major operation on HBase system tables (such as upgrade or migration), you need to back up HBase metadata to prevent HBase service unavailability caused by HBase system table directory or file damages. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up HBase metadata. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- The **fs.defaultFS** parameter settings of HBase are the same as those of Yarn and HDFS.
- If HBase data is stored in the local HDFS, HBase metadata can be backed up to OBS. If HBase data is stored in OBS, data backup is not supported.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-76 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, select **HBase** under **Metadata and other data**.

NOTE

If there are multiple HBase services, all HBase services are backed up by default. You can click **Assign Service** to specify the services to be backed up.

Step 7 Set **Path Type** of **HBase** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files.

  The default storage directory is *Data storage path*/**LocalBackup/**, for example, **/srv/BigData/LocalBackup**.

  If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory,

and the user group and owner group of the shared path must be **nobody:nobody**.)

- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

- **OBS**: indicates that backup files are stored in OBS.

  If you select this option, set the following parameters:

  - **Target Path**: indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

    NOTE

    Only MRS 3.1.0 or later supports data backup to OBS.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

8.10.2.8 Backing Up HBase Service Data

Scenario

To ensure HBase service data security routinely or before a major operation on HBase (such as upgrade or migration), you need to back up HBase service data. The backup data can be used to recover the system if an exception occurs or the

operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up HBase service data. Both automatic and manual backup tasks are supported.

The following situations may occur during the HBase service data backup:

- When a user creates an HBase table, **KEEP_DELETED_CELLS** is set to **false** by default. When the user backs up this HBase table, deleted data will be backed up and junk data may exist after data restoration. This parameter can be set to **true** manually when an HBase table is created based on service requirements.
- When a user manually specifies the timestamp when writing data into an HBase table and the specified time is earlier than the last backup time of the HBase table, new data may not be backed up in incremental backup tasks.
- The HBase backup function cannot back up the access control lists (ACLs) for reading, writing, executing, creating, and managing HBase global or namespaces. After HBase data is restored, you need to reset the role permissions on FusionInsight Manager.
- If the backup data of the standby cluster is lost in an existing HBase backup task, the next incremental backup will fail, and you need to create an HBase backup task again. However, the next full backup task will be normal.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
- On the HDFS client, you have executed the **hdfs lsSnapshottableDir** command as user **hdfs** to check the list of directories for which HDFS snapshots have been created in the current cluster and ensured that the HDFS parent directory or subdirectory where data files to be backed up are stored does not have HDFS snapshots. Otherwise, the backup task cannot be created.
- If you want to back up data to NAS, you have deployed the NAS server in advance.

- The **fs.defaultFS** parameter settings of HBase are the same as those of Yarn and HDFS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-77 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, choose **HBase** > **HBase** under **Service data**.

Step 7 Set **Path Type** of **HBase** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select this option, set the following parameters:

  – **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in

remote cluster of the cluster, or the NameService name of a configured remote cluster.

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.

  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

Step 8 Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

Step 9 Set **Backup Content** to one or multiple HBase tables to be backed up.

You can select backup data using either of the following methods:

- Adding a backup data file

  Click the name of a database in the navigation tree to show all the tables in the database, and select specified tables.

  MRS 3.2.0 or later:

  a. Click **Add**.
  b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.
  c. Click **OK**.

- Selecting using regular expressions

  a. Click **Query Regular Expression**.
  b. Enter the namespace where the HBase tables are located in the first text box as prompted. The namespace must be the same as the existing namespace, for example, **default**.
  c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all tables in the namespace, enter **([\s\S]*?)**. To get tables whose names consist of letters and digits, for example, **tb1**, enter **tb\d***.
  d. Click **Refresh** to view the displayed tables in **Directory Name**.
  e. Click **Synchronize** to save the result.

  NOTE

  - When entering regular expressions, click + or — to add or delete an expression.
  - If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.

Step 10 Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The queue name is incorrect.
- The parent directory or subdirectory of the HDFS directory where HBase table data files to be backed up are stored has HDFS snapshots.
- The directory or table to be backed up does not exist.

Step 11 Click **OK**.

Step 12 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *xxx/Backup task name_Data source_Task creation time*, and the subdirectory is used to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

8.10.2.9 Backing Up NameNode Data

Scenario

To ensure NameNode service data security routinely or before a major operation on NameNode (such as upgrade or migration), you need to back up NameNode data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up NameNode data. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-78 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	Only **Full backup every time** is supported. **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, select **NameNode**.

Step 7 Set **Path Type** of **NameNode** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path*/**LocalBackup/**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select this option, set the following parameters:
  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the service plane IP address of the NameNode in the standby cluster.
  - **Target Path**: indicates the path for storing backup files.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.

  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **OBS**: indicates that backup files are stored in OBS.

  If you select this option, set the following parameters:
  - **Target Path**: indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

    NOTE

    Only MRS 3.1.0 or later supports data backup to OBS.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files.

The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

8.10.2.10 Backing Up HDFS Service Data

Scenario

To ensure HDFS service data security routinely or before a major operation on HDFS (such as upgrade or migration), you need to back up HDFS service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up HDFS service data. Both automatic and manual backup tasks are supported.

NOTE

Encrypted directories cannot be backed up or restored.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
- On the HDFS client, you have executed the **hdfs lsSnapshottableDir** command as user **hdfs** to check the list of directories for which HDFS snapshots have been created in the current cluster and ensured that the HDFS parent directory or subdirectory where data files to be backed up are stored does not have HDFS snapshots. Otherwise, the backup task cannot be created.
- If you want to back up data to NAS, you have deployed the NAS server in advance.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-79 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.

Parameter	Description
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, select **HDFS**.

Step 7 Set **Path Type** of **HDFS** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select this option, set the following parameters:

  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
  - **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.

  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

Step 8 Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

Step 9 Set **Backup Content** to one or multiple HDFS directories to be backed up based on service requirements.

You can select backup data using either of the following methods:

- Adding a backup data file

  Click the name of a database in the navigation tree to show all the tables in the database, and select specified tables.

  MRS 3.2.0 or later:

  a. Click **Add**.

  b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.

  c. Click **OK**.
- Selecting using regular expressions

a. Click **Query Regular Expression**.
b. Enter the parent directory full path of the directory in the first text box as prompted. The directory must be the same as the existing directory, for example, **/tmp**.
c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all files or subdirectories in the parent directory, enter **([\s\S]*?)**. To get files whose names consist of letters and digits, for example, **file*1***, enter **file\d***.
d. Click **Refresh** to view the displayed directories in **Directory Name**.
e. Click **Synchronize** to save the result.

NOTE

- When entering regular expressions, click + or — to add or delete an expression.
- If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.
- The backup directory cannot contain files that have been written for a long time. Otherwise, the backup task will fail. Therefore, you are not advised to perform operations on the top-level directory, such as **/user**, **/tmp**, and **/mr-history**.

Step 10 Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The queue name is incorrect.
- The parent directory or subdirectory of the HDFS directory where data files to be backed up are stored has HDFS snapshots.
- The directory or table to be backed up does not exist.
- The name of the NameService is incorrect.

Step 11 Click **OK**.

Step 12 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Data source_Task creation time*, and the subdirectory is used to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

8.10.2.11 Backing Up Hive Service Data

Scenario

To ensure Hive service data security routinely or before a major operation on Hive (such as upgrade or migration), you need to back up Hive service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Hive service data. Both automatic and manual backup tasks are supported.

- Hive backup and restoration cannot identify the service and structure relationships of objects such as Hive tables, indexes, and views. When executing backup and restoration tasks, you need to manage unified restoration points based on service scenarios to ensure proper service running.
- Hive backup and restoration do not support Hive on RDB data tables. You need to back up and restore original data tables in external databases independently.
- If the backup data of the standby cluster is lost in an existing Hive backup task that contains Hive on HBase tables, the next incremental backup will fail, and you need to create a Hive backup task again. However, the next full backup task will be normal.
- After the backup function of FusionInsight Manager is used to back up the HDFS directories at the Hive table level, the Hive tables cannot be deleted and recreated.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.
- On the HDFS client, you have executed the **hdfs lsSnapshottableDir** command as user **hdfs** to check the list of directories for which HDFS snapshots have been created in the current cluster and ensured that the HDFS parent directory or subdirectory where data files to be backed up are stored does not have HDFS snapshots. Otherwise, the backup task cannot be created.
- If you want to back up data to NAS, you have deployed the NAS server in advance.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-80 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, choose **Hive** > **Hive**.

Step 7 Set **Path Type** of **Hive** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select this option, set the following parameters:
  - **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.

  - **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol. If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol. If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.

  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

Step 8 Set **Maximum Number of Recovery Points** to the number of snapshots that can be retained in the cluster.

Step 9 Set **Backup Content** to one or multiple Hive tables to be backed up.

You can select backup data using either of the following methods:

- Adding a backup data file

  Click the name of a database in the navigation tree to show all the tables in the database, and select specified tables.

  MRS 3.2.0 or later:

  a. Click **Add**.
  b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.
  c. Click **OK**.

- Selecting using regular expressions

  a. Click **Query Regular Expression**.
  b. Enter the database where the Hive tables are located in the first text box as prompted. The database must be the same as the existing database, for example, **default**.
  c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all tables in the database, enter **([\s\S]*?)**. To get tables whose names consist of letters and digits, for example, **tb1**, enter **tb\d***.
  d. Click **Refresh** to view the displayed tables in **Directory Name**.
  e. Click **Synchronize** to save the result.

  NOTE

  - When entering regular expressions, click + or − to add or delete an expression.
  - If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.

Step 10 Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The queue name is incorrect.
- The parent directory or subdirectory of the HDFS directory where data files to be backed up are stored has HDFS snapshots.
- The directory or table to be backed up does not exist.
- The name of the NameService is incorrect.

Step 11 Click **OK**.

Step 12 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Data source_Task creation time*, and the subdirectory is used to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

8.10.2.12 Backing Up IoTDB Metadata

Scenario

To ensure IoTDB metadata security and prevent the IoTDB service from being unavailable due to IoTDB metadata file damages, you need to back up IoTDB metadata. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up IoTDB metadata. Both automatic and manual backup tasks are supported.

Prerequisites

- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- Backup policies, including the backup task type, period, backup object, backup directory, and Yarn queue required by the backup task are planned based on service requirements.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-81 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	The task execution interval. Value options include **Hours** and **Days**.

Parameter	Description
Backup Policy	Indicates a periodic backup policy. • **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** Incremental backup is not supported when component metadata is backed up. Only **Full backup every time** is supported.

Step 6 In **Configuration**, select **IoTDB** under **Metadata and other data**.

Step 7 Set **Path Type** of **IoTDB** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path*/**LocalBackup/**.

  If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select this option, set the following parameters:
  - **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Destination Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
  - **Destination Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.

  - **Destination NameNode RPC Port**: indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the standby cluster.
  - **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Server Shared Path**: indicates the backup path on the SFTP server.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files.

The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

8.10.2.13 Backing Up IoTDB Service Data

Scenario

To ensure IoTDB service data security routinely or before a major operation on IoTDB (such as upgrade or migration), you need to back up IoTDB service data. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up IoTDB service data. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster. Currently, IoTDB data can be backed up only to HDFS.
- For the IoTDB cluster in normal mode, service data cannot be backed up to HDFS in a cluster in security mode.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The HDFS in the standby cluster has sufficient space. You are advised to save backup files in a custom directory.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-82 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	The task execution interval. Value options include **Hours** and **Days**.
Backup Policy	Indicates a periodic backup policy. • **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times**

Step 6 In **Configuration**, choose **IoTDB** > **IoTDB** under **Service data**.

Step 7 Set **Path Type** of **IoTDB** to a backup directory type.

The following backup directory types are supported:

RemoteHDFS: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Destination Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
- **Destination Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Destination NameNode RPC Port**: indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the standby cluster.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

Step 8 Set **Backup Content** to one or multiple service data records to be backed up.

You can select backup data using either of the following methods:

- Adding a backup data file

  Click the name of a database in the navigation tree to show all the tables in the database, and select specified tables.

  MRS 3.2.0 or later:

  a. Click **Add**.

  b. Select the table to be backed up under **File Directory**, and click **Add** to add the table to **Backup Content**.

  c. Click **OK**.

- Selecting using regular expressions

  a. Click **Query Regular Expression**.

  b. Enter the parent directory full path of the directory in the first text box as prompted. The directory must be the same as the existing directory, for example, **/root**.

  c. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get all files or subdirectories in the parent directory, enter **([\s\S]*?)**. To get files whose names consist of letters and digits, for example, **file*1***, enter **file\d***.

  d. Enter a regular expression in the second text box. Standard regular expressions are supported. For example, to get objects containing **test**, enter **.*test.***. To get objects starting with **test**, enter **test.***. To get objects ending with **test**, enter **.*test**.

  e. Click **Refresh** to view the displayed directories in **Directory Name**.

  f. Click **Synchronize** to save the result.

  NOTE

  - When entering regular expressions, click + or — to add or delete an expression.
  - If the selected table or directory is incorrect, click **Clear Selected Node** to deselect it.
  - The backup directory cannot contain files that have been written for a long time. Otherwise, the backup task will fail. Therefore, you are not advised to perform operations on the top-level directory, such as **/user**, **/tmp**, and **/mr-history**.

Step 9 Click **Verify** to check whether the backup task is configured correctly.

The possible causes of the verification failure are as follows:

- The target NameNode IP address is incorrect.
- The data to be backed up does not exist.

Step 10 Click **OK**.

Step 11 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Data source_Task creation time*, and the subdirectory is used

to save latest data source backup files. All the backup file sets are stored in the related snapshot directories.

----End

8.10.2.14 Backing Up Kafka Metadata

Scenario

To ensure Kafka metadata security or before a major operation on ZooKeeper (such as upgrade or migration), you need to back up Kafka metadata. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

You can create a backup task on FusionInsight Manager to back up Kafka metadata. Both automatic and manual backup tasks are supported.

Prerequisites

- If data needs to be backed up to the remote HDFS, you have prepared a standby cluster for data backup. The authentication mode of the standby cluster is the same as that of the active cluster. For other backup modes, you do not need to prepare the standby cluster.
- If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.
- If you want to back up data to NAS, you have deployed the NAS server in advance.
- If you want to back up data to OBS, you have connected the current cluster to OBS and have the permission to access OBS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 Click **Create**.

Step 3 Set **Name** to the name of the backup task.

Step 4 Select the cluster to be operated from **Backup Object**.

Step 5 Set **Mode** to the type of the backup task.

Periodic indicates that the backup task is executed by the system periodically. **Manual** indicates that the backup task is executed manually.

Table 8-83 Periodic backup parameters

Parameter	Description
Started	Indicates the time when the task is started for the first time.
Period	Indicates the task execution interval. The options include **Hours** and **Days**.
Backup Policy	• **Full backup at the first time and incremental backup subsequently** • **Full backup every time** • **Full backup once every n times** **NOTE** • Incremental backup is not supported when Manager data and component metadata are backed up. Only **Full backup every time** is supported. • If **Path Type** is set to **NFS** or **CIFS**, incremental backup cannot be used. When incremental backup is used for NFS or CIFS backup, the latest full backup data is updated each time the incremental backup is performed. Therefore, no new recovery point is generated.

Step 6 In **Configuration**, select **Kafka**.

Step 7 Set **Path Type** of **Kafka** to a backup directory type.

The following backup directory types are supported:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path*/**LocalBackup/**.

  If you select this option, you need to set the maximum number of replicas to specify the number of backup file sets that can be retained in the backup directory.
- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select this option, set the following parameters:

  – **Target Path**: indicates the HDFS directory for storing the backup files. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.

  – **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

  – **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, set the following parameters:

- **Destination NameService Name**: indicates the NameService name of the standby cluster. You can set it to the NameService name (**haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**) of the built-in remote cluster of the cluster, or the NameService name of a configured remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Target NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Target Path**: indicates the HDFS directory for storing standby cluster backup data. The storage path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory, such as **/hbase** or **/user/hbase/backup**.
- **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory, and the user group and owner group of the shared path must be **nobody:nobody**.)
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.

  If you select this option, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Server Shared Path**: indicates the configured shared directory of the NAS server. (The shared path of the server cannot be set to the root directory,

and the user group and owner group of the shared path must be **nobody:nobody**.)

  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

- **OBS**: indicates that backup files are stored in OBS.

  If you select this option, set the following parameters:

  - **Target Path**: indicates the OBS directory for storing backup data.
  - **Maximum Number of Backup Copies**: indicates the number of backup file sets that can be retained in the backup directory.

  NOTE

  Only MRS 3.1.0 or later supports data backup to OBS.

Step 8 Click **OK**.

Step 9 In the **Operation** column of the created task in the backup task list, click **More** and select **Back Up Now** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

8.10.3 Recovering Data

8.10.3.1 Restoring Manager Data

Scenario

Manager data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in FusionInsight Manager, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable.

System administrators can create a restoration task in FusionInsight Manager to recover Manager data. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that of data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the Manager data that is generated after the data backup and before the data restoration will be lost.

Impact on the System

- In the restoration process, the Controller needs to be restarted and FusionInsight Manager cannot be logged in or operated during the restart.
- In the restoration process, all clusters need to be restarted and cannot be accessed during the restart.
- After data restoration, the data, such as system configuration, user information, alarm information, and audit information, that is generated after the data backup and before the data restoration will be lost. This may result in data query failure or cluster access failure.
- After the Manager data is recovered, the system forces the LdapServer of each cluster to synchronize data from the OLadp.

Prerequisites

- To restore data from a remote HDFS, you need to prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, system mutual trust needs to be configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The status of the OMS resources and the LdapServer instances of each cluster is normal. If the status is abnormal, data restoration cannot be performed.
- The status of the cluster hosts and services is normal. If the status is abnormal, data restoration cannot be performed.
- The cluster host topologies during data restoration and data backup are the same. If the topologies are different, data restoration cannot be performed and you need to back up data again.
- The services added to the cluster during data restoration and data backup are the same. If the topologies are different, data restoration cannot be performed and you need to back up data again.
- The upper-layer applications that depend on the cluster are stopped.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view the historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restore** > **Restoring Management**. On the displayed page, click **Create**.

Figure 8-59 Creating a restoration task

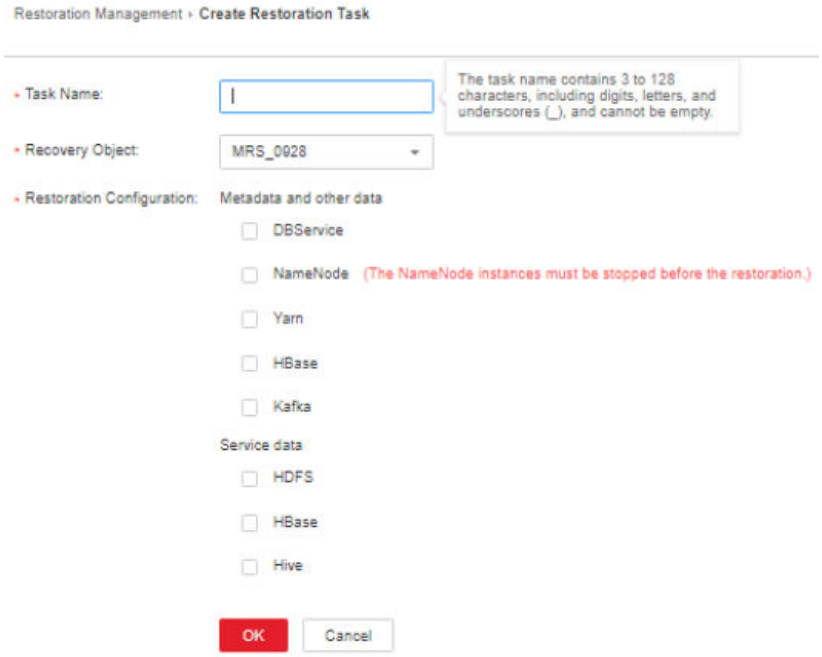

Step 4 Set **Task Name** to the name of the restoration task.

Step 5 Set **Recovery Object** to **OMS**.

Step 6 Select **OMS**.

Step 7 Set **Path Type** of **OMS** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.

  If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version_Data source_Task execution time*.**tar.gz**.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select **LocalHDFS**, set the following parameters:

  - **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name_Task creation time/Version_Data source_Task execution time*.**tar.gz**.
  - **Cluster for Restoration**: Enter the name of the cluster used during restoration task execution.
  - **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select **RemoteHDFS**, set the following parameters:

  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **Source Cluster**: Select the cluster of the Yarn queue used by the recovery data.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol. If you select **NFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the complete path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select **SFTP**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

- **OBS**: indicates that backup files are stored in OBS.

  If you select **OBS**, set the following parameters:

  - **Source Path**: indicates the full OBS path of a backup file, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

    NOTE

    Only MRS 3.1.0 or later supports saving backup files in OBS.

Step 8 Click **OK**.

Step 9 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

Step 10 Log in to the active and standby management nodes as user **omm**.

Step 11 Run the following command to restart OMS:

sh ${BIGDATA_HOME}/om-server/om/sbin/restart-oms.sh

The command is run successfully if the following information is displayed:

```
start HA successfully.
```

Run **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** to check whether **HAAllResOK** of the management node is **Normal** and whether FusionInsight Manager can be logged in again. If yes, OMS is restarted successfully.

Step 12 On FusionInsight Manager, click **Cluster**, click the name of the target cluster, and choose **Services** > **KrbServer**. On the displayed page, choose **More** > **Synchronize**

Configuration, click **OK**, and wait for the KrbServer configuration to be synchronized and the service to be restarted.

Step 13 Choose **Cluster**, click the name of the desired cluster, and choose **More** > **Synchronize Configurations**, click **OK**, and wait until the cluster configuration is synchronized successfully.

Step 14 On FusionInsight Manager, click **Cluster**, click the name of the target cluster, and choose **More** > **Restart**. On the displayed page, enter the password of the current login user, click **OK**, and wait for the cluster to be restarted.

----End

8.10.3.2 Restoring CDL Data

Scenario

CDL data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs major operations (such as upgrade and data adjustment) on CDL, an exception occurs or the expected result is not achieved. All modules are faulty and become unavailable. Data is migrated to a new cluster.

CDL metadata is stored in DBService and Kafka. A system administrator can create DBService and Kafka restoration tasks on FusionInsight Manager to restore CDL data. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the DBService and Kafka data that is generated after the data backup and before the data restoration will be lost.
- By default, MRS clusters use DBService to store metadata of Hive, Hue, Loader, Spark, CDL, and Oozie. Restoring DBService data will restore the metadata of all these components.

Impact on the System

- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the configurations of the components that depend on DBService may expire and these components need to be restarted.
- After the metadata is restored, the offset information stored on ZooKeeper by Kafka consumers is rolled back, resulting in repeated consumption.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby

clusters are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The status of the active and standby DBService instances is normal. If the status is abnormal, data restoration cannot be performed.
- The Kafka service is disabled first, and then enabled upon data restoration.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the row where the specified backup task is located, choose **More** > **View History** in the **Operation** column to display the historical execution records of the backup task.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object**: indicates the backup data source.
- **Backup Path**: indicates the full path where backup files are stored.

  Select the correct path, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In the **Restoration Configuration** area, select **DBService** and **Kafka**.

Step 8 Set **Path Type** of **DBService** to a backup directory type. For details about how to configure the parameters, see **Step 8**.

Step 9 Set **Path Type** of **Kafka** to a backup directory type. For details about how to configure the parameters, see **Step 8**.

Step 10 Click **OK**.

Step 11 In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.

- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

8.10.3.3 Restoring ClickHouse Metadata

Scenario

ClickHouse metadata needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After a user performs major operations (such as upgrade and migration) on ClickHouse, an exception occurs or the expected result is not achieved. The ClickHouse component is faulty and becomes unavailable. Data is migrated to a new cluster.

Users can create a ClickHouse restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

NOTICE

- This function is supported only by MRS 3.1.0 or later.
- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore ClickHouse metadata when the service is running properly, you are advised to manually back up the latest ClickHouse metadata before restoration. Otherwise, the ClickHouse metadata that is generated after the data backup and before the data restoration will be lost.
- ClickHouse metadata restoration and service data restoration cannot be performed at the same time. Otherwise, service data restoration fails. You are advised to restore service data after metadata restoration is complete.

Impact on the System

- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the ClickHouse upper-layer applications need to be started.

Prerequisites

- You have checked the path for storing ClickHouse metadata backup files.
- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active and standby clusters are deployed in security mode and they are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active and standby clusters are deployed in normal mode, no mutual trust is required.
- In the active/standby cluster, when restoring data from the remote HDFS to the local host, ensure that the value of **HADOOP_RPC_PROTECTION** of ClickHouse is the same as that of **hadoop.rpc.protection** of HDFS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of the specified task in the task list, choose **More** > **View History**.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object**: indicates the backup data source.
- **Backup Path**: indicates the full path where backup files are stored.

  Select the correct path, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, select **ClickHouse** under **Metadata and other data**.

Step 8 Set **Path Type** of **ClickHouse** to a restoration directory type.

The configurations vary based on backup directory types:

- **LocalDir**: indicates that data is restored from the local disk of the active management node.

  If you select this value, you also need to configure the following parameters:

  - **Source Path**: backup file to be restored, for example, *Backup task name_Data source_Task execution time*.**tar.gz**.

- **RemoteHDFS**: indicates that data is restored from the HDFS directory of the standby cluster.

  If you select this value option for MRS 3.2.0 or later clusters, you also need to configure the following parameters:

  - **Source NameService Name**: indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
  - **Source Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.

- **Source NameNode RPC Port**: indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the destination cluster.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Data source_Task execution time*.**tar.gz**.
- **Logical Cluster** of MRS 3.2.0 or later: Enter the ClickHouse logical cluster whose data has been backed up.

If you select this value option for MRS 3.1.0 or 3.1.2 clusters, you also need to configure the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Data source_Task execution time*.**tar.gz**.

Step 9 Click **OK**.

Step 10 In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column. In the displayed dialog box, click **OK** to start the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

Step 11 Choose **Cluster** > **Services** and start the ClickHouse service.

----End

8.10.3.4 Restoring ClickHouse Service Data

Scenario

ClickHouse data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After a user performs major operations (such as upgrade and migration) on ClickHouse, an exception occurs or the expected result is not achieved. All modules are faulty and become unavailable. Data is migrated to a new cluster.

Users can create a ClickHouse restoration task on FusionInsight Manager to restore data. Only manual restoration tasks are supported.

The ClickHouse backup and restoration functions cannot identify the service and structure relationships of objects such as ClickHouse tables, indexes, and views. When executing backup and restoration tasks, you need to manage unified restoration points based on service scenarios to ensure proper service running.

NOTICE

- This function is supported only by MRS 3.1.0 or later.
- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the ClickHouse data that is generated after the data backup and before the data restoration will be lost.
- ClickHouse metadata restoration and service data restoration cannot be performed at the same time. Otherwise, service data restoration fails. You are advised to restore service data after metadata restoration is complete.

Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the ClickHouse upper-layer applications need to be started.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active and standby clusters are deployed in security mode and they are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active and standby clusters are deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The database for storing restored data tables, the HDFS save path of data tables, and the list of users who can access restored data are planned.
- The ClickHouse backup file save path is correct.
- The ClickHouse upper-layer applications are stopped.
- In the active/standby cluster, when restoring data from the remote HDFS to the local host, ensure that the value of **HADOOP_RPC_PROTECTION** of ClickHouse is the same as that of **hadoop.rpc.protection** of HDFS.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the row where the specified backup task is located, choose **More** > **View History** in the **Operation** column to display the historical execution records of the backup task.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object**: indicates the backup data source.
- **Backup Path**: indicates the full path where backup files are stored.

  Select the correct path, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, select **ClickHouse** under **Service data**.

Step 8 Set **Path Type** of **ClickHouse** to a backup directory type.

Currently, the backup directory supports only the **RemoteHDFS** type.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select this value option, you also need to configure the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the IP address of the NameNode service plane in the standby cluster. It can be of an active or standby node.
  - Source Path: indicates the full path of HDFS directory for storing backup data of the standby cluster. For details, see the **Backup Path** obtained in step **Step 2**, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

Step 9 Click **OK**.

Step 10 In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.

- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

8.10.3.5 Restoring DBService data

Scenario

DBService data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in DBService, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover DBService data. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the DBService data that is generated after the data backup and before the data recovery will be lost.
- By default, MRS clusters use DBService to store metadata of Hive, Hue, Loader, Spark, CDL, and Oozie. Restoring DBService data will restore the metadata of all these components.

Impact on the System

- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the configurations of the components that depend on DBService may expire and these components need to be restarted.

Prerequisites

- To restore data from a remote HDFS, you need to prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.

- The status of the active and standby DBService instances is normal. If the status is abnormal, data restoration cannot be performed.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view the historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In the **Restoration Configuration** area, select **DBService**.

NOTE

If multiple DBServices are installed, select the DBServices to be restored.

Step 8 Set **Path Type** of **DBService** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.

  If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version_Data source_Task execution time*.**tar.gz**.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select **LocalHDFS**, set the following parameters:

  - **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name_Task creation time/Version_Data source_Task execution time*.**tar.gz**.
  - **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select **RemoteHDFS**, set the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select **NFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol.

  If you select **CIFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select **SFTP**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **OBS**: indicates that backup files are stored in OBS.

  If you select **OBS**, set the following parameters:

  - **Source Path**: indicates the full OBS path of a backup file, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

    NOTE

    Only MRS 3.1.0 or later supports saving backup files in OBS.

Step 9 Click **OK**.

Step 10 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

8.10.3.6 Restoring Flink Metadata

Scenario

Flink metadata needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs major operations (such as upgrade and data adjustment) on Flink, an exception occurs or the expected result is not achieved. The Flink component is faulty and becomes unavailable. Data is migrated to a new cluster.

System administrators can create a Flink restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore Flink metadata when the service is running properly, you are advised to manually back up the latest Flink metadata before restoration. Otherwise, the Flink metadata that is generated after the data backup and before the data restoration will be lost.
- Flink metadata restoration and service data restoration cannot be performed at the same time. Otherwise, service data restoration fails. You are advised to restore service data after metadata restoration is complete.

Impact on the System

- Before restoring the metadata, you need to stop the Flink service. During this period, all upper-layer applications are affected and cannot work properly.
- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the Flink upper-layer applications of Solr need to be started.

Prerequisites

- You have checked the path for storing Flink metadata backup files.
- The Flink service has been stopped before its metadata is restored.
- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of the specified task in the task list, choose **More** > **View History**.

In the displayed window, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct path, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, select **Flink** under **Metadata and other data**.

Step 8 Set **Path Type** of **Flink** to a restoration directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that data is restored from the local disk of the active management node.

  If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Backup task name_Data source_Task execution time***.tar.gz**.

- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select **LocalHDFS**, set the following parameters:

  - **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed.

- **RemoteHDFS**: indicates that data is restored from the HDFS directory of the standby cluster.

  If you select **RemoteHDFS**, set the following parameters:

  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Data source_Task execution time***.tar.gz**.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the source cluster.

Step 9 Click **OK**.

Step 10 In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column. In the displayed dialog box, click **OK** to start the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

Step 11 Choose **Cluster** > **Services** and start the Flink service.

----End

8.10.3.7 Restoring HBase Metadata

Scenario

To ensure HBase metadata security (including tableinfo files and HFiles) or before a major operation on HBase system tables (such as upgrade or migration), you need to back up HBase metadata to prevent HBase service unavailability caused by HBase system table directory or file damages. The backup data can be used to recover the system if an exception occurs or the operation has not achieved the expected result, minimizing the adverse impacts on services.

System administrators can create a recovery task in FusionInsight Manager to recover HBase metadata. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the HBase data that is generated after the data backup and before the data recovery will be lost.
- It is recommended that a data restoration task restore the metadata of only one component to prevent the data restoration of other components from being affected by stopping a service or instance. If data of multiple components is restored at the same time, data restoration may fail.

  HBase metadata cannot be restored at the same time as NameNode metadata. As a result, data restoration fails.

Impact on the System

- Before restoring the metadata, you need to stop the HBase service, during which the HBase upper-layer applications are unavailable.
- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the upper-layer applications of HBase need to be started.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- You have checked the path for storing HBase metadata backup files.
- The HBase service has been stopped before its metadata is restored.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view the historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, select **HBase** under **Metadata and other data**.

NOTE

If multiple HBase services are installed, select the HBase services to be restored.

Step 8 Set **Path Type** of **HBase** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.

  If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version_Data source_Task execution time*.**tar.gz**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select **RemoteHDFS**, set the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol.

  If you select **NFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

- **CIFS**: indicates that backup files are stored in NAS using the CIFS protocol.

  If you select **CIFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select **SFTP**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **OBS**: indicates that backup files are stored in OBS.

  If you select **OBS**, set the following parameters:

  - **Source Path**: indicates the full OBS path of a backup file, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

    NOTE

    Only MRS 3.1.0 or later supports saving backup files in OBS.

Step 9 Click **OK**.

Step 10 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

8.10.3.8 Restoring HBase Service Data

Scenario

HBase data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in HBase, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover HBase data. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the HBase data that is generated after the data backup and before the data recovery will be lost.

Impact on the System

- During the data recovery process, the system disables the HBase table to be recovered and the table cannot be accessed in this moment. The data recovery process takes several minutes, during which the HBase upper-layer applications are unavailable.
- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the HBase upper-layer applications need to be started.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The directory for saving the backup file has been checked.
- The HBase upper-layer applications have been stopped.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.

- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, select **HBase** under **Service Data**.

Step 8 Set **Path Type** of **HBase** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select **RemoteHDFS**, set the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/xxx/Backup task name_Data source_Task creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More** > **View History** in the **Operation** column, and click **View** in the **Backup Path** column.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol. If you select **NFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.

  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/xxx/Backup task name_Data source_Task creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More** > **View History** in the **Operation** column, and click **View** in the **Backup Path** column.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/xxx/Backup task name_Data source_Task creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More** > **View History** in the **Operation** column, and click **View** in the **Backup Path** column.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select **SFTP**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.

- **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
- **Username**: indicates the username for connecting to the server using the SFTP protocol.
- **Password**: indicates the password for connecting to the server using the SFTP protocol.
- **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/xxx/Backup task name_Data source_Task creation time*. To obtain the backup path, locate the row that contains the target backup task in the backup management list, click **More** > **View History** in the **Operation** column, and click **View** in the **Backup Path** column.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
- **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
- **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
- **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

Step 9 Set **Backup Data** column in **Data Configuration** to one or multiple backup data sources to be recovered. In the **Target Namespace** column, specify the target naming space after backup data recovery.

You are advised to set **Target Namespace** to a location that is different from the backup naming space.

Step 10 Set **Force recovery** to **true**, which indicates to forcibly recover all backup data when a data table with the same name already exists. If the data table contains new data added after backup, the new data will be lost after the data recovery. If you set the parameter to **false**, the restoration task is not executed if a data table with the same name exists.

Step 11 Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified naming space does not exist, the verification fails.
- If the forcible overwrite conditions are not met, the verification fails.

Step 12 Click **OK** to save the settings.

Step 13 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

Step 14 Check whether HBase data is restored in an environment where HBase is newly installed or reinstalled.

- If yes, the administrator needs to set new permission for roles on FusionInsight Manager based on the original service plan.
- If no, no further operation is required.

----End

8.10.3.9 Restoring NameNode Data

Scenario

NameNode data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in NameNode, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover NameNode data. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the NameNode data that is generated after the data backup and before the data recovery will be lost.
- It is recommended that a data restoration task restore the metadata of only one component to prevent the data restoration of other components from being affected by stopping a service or instance. If data of multiple components is restored at the same time, data restoration may fail.

  HBase metadata cannot be restored at the same time as NameNode metadata. As a result, data restoration fails.

Impact on the System

- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the NameNode needs to be restarted and is unavailable during the restart.
- After data is restored, metadata and service data may not be matched, the HDFS enters the security mode, and the HDFS service fails to be started. .

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- On FusionInsight Manager, all the NameNode role instances whose data is to be recovered are stopped. Other HDFS role instances must keep running. After data is recovered, the NameNode role instances need to be restarted. The NameNode role instances cannot be accessed during the restart.
- The NameNode backup files are stored ***Data path*/LocalBackup/** on the active management node.

Procedure

Step 1 On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Services** > **HDFS**. On the displayed page, click **Instance** and click **NameNode** to check whether the NameNode instances of the data to be restored are stopped. If the NameNode instances are not stopped, stop them.

Step 2 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 3 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 4 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 5 Click **Create**.

Step 6 Set **Task Name** to the name of the restoration task.

Step 7 Select the cluster to be operated from **Recovery Object**.

Step 8 In the **Restoration Configuration** area, select **NameNode**.

Step 9 Set **Path Type** of **NameNode** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.

  If you select **LocalDir**, set the following parameters:

  - **Source Path**: indicates the full path of the backup file on the local disk, for example, *Backup path/Backup task name_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select **RemoteHDFS**, set the following parameters:

  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **NFS**: indicates that backup files are stored in the NAS using the NFS protocol. If you select **NFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

- **CIFS**: indicates that backup files are stored in the NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.

  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select **SFTP**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
- **OBS**: indicates that backup files are stored in OBS.

  If you select **OBS**, set the following parameters:
  - **Source Path**: indicates the full OBS path of a backup file, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
  - **NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

    NOTE

    Only MRS 3.1.0 or later supports saving backup files in OBS.

Step 10 Click **OK**.

Step 11 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

Step 12 On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Services** > **HDFS**. On the displayed page, click **Configurations** and click **All Configurations**.

On the displayed page, enter the password of the administrator who has logged in for authentication and click **OK**. After the system displays "Operation succeeded", click **Finish**. The service is started successfully.

----End

8.10.3.10 Restoring HDFS Service Data

Scenario

HDFS data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in the HDFS, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover HDFS data. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the HDFS data that is generated after the data backup and before the data recovery will be lost.
- The HDFS restoration operation cannot be performed for the directories used by running Yarn tasks, for example, **/tmp/logs**, **/tmp/archived**, and **/tmp/hadoop-yarn/staging**. Otherwise, data restoration using Distcp tasks fails due to file loss.

Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the HDFS upper-layer applications need to be started.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.

- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The HDFS backup file save path is correct.
- The HDFS upper-layer applications are stopped.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, select **HDFS** under **Service Data**.

Step 8 Set **Path Type** of **HDFS** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select **RemoteHDFS**, set the following parameters:

  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time*.

  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **NFS**: indicates that backup files are stored in NAS using the NFS protocol. If you select **NFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **CIFS**: indicates that backup files are stored in NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.

  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select **SFTP**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

Step 9 In the **Backup Data** column of the **Data Configuration** page, select one or more pieces of backup data that needs to be restored based on service requirements. In the **Target Path** column, specify the target location after backup data restoration.

You are advised to set **Target Path** to a new path that is different from the backup path.

Step 10 Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified directory to be restored does not exist, the verification fails.

Step 11 Click **OK**.

Step 12 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

8.10.3.11 Restoring Hive Service Data

Scenario

Hive data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in the Hive, an exception occurs or the operation has not achieved the expected result. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover Hive data. Only manual restoration tasks are supported.

Hive backup and restoration cannot identify the service and structure relationships of objects such as Hive tables, indexes, and views. When executing backup and restoration tasks, you need to manage unified restoration points based on service scenarios to ensure proper service running.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To recover data when the service is running properly, you are advised to manually back up the latest management data before recovering data. Otherwise, the Hive data that is generated after the data backup and before the data recovery will be lost.

Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is recovered, the Hive upper-layer applications need to be started.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust

has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.

- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The database for storing restored data tables, the HDFS save path of data tables, and the list of users who can access restored data are planned.
- The Hive backup file save path is correct.
- The Hive upper-layer applications are stopped.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In the **Restoration Configuration** area, select **Hive**.

Step 8 Set **Path Type** of **Hive** to a backup directory type.

The following backup directory types are supported:

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster. If you select **RemoteHDFS**, set the following parameters:
  - **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**, **haclusterX3**, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.

  - **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select a Hive backup file set that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **NFS**: indicates that backup files are stored in NAS using the NFS protocol. If you select **NFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select a Hive backup file set that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.
- **CIFS**: indicates that backup files are stored in NAS using the CIFS protocol. If you select **CIFS**, set the following parameters:
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.

  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select a Hive backup file set that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **100**.

- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select **SFTP**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path**: indicates the full path of the backup file on the backup server, for example, *Backup path/Backup task name_Data source_Task creation time*.
  - **Queue Name**: indicates the name of the Yarn queue used for backup task execution.
  - **Recovery Point List**: Click **Refresh** and select an HDFS directory that has been backed up in the standby cluster.
  - **Target NameService Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.
  - **Maximum Number of Maps**: indicates the maximum number of maps in a MapReduce task. The default value is **20**.
  - **Maximum Bandwidth of a Map (MB/s)**: indicates the maximum bandwidth of a map. The default value is **1**.

Step 9 Set **Backup Data** in the **Data Configuration** to one or multiple backup data sources to be recovered based on service requirements. In the **Target Database** and **Target Path** columns, specify the target database and file save path after backup data recovery.

Configuration restrictions:

- Data can be restored to the original database, but data tables must be stored in a new path that is different from the backup path.
- To restore Hive index tables, select the Hive data tables that correspond to the Hive index tables to be restored.
- If a new restoration directory is selected to avoid affecting the current data, HDFS permission must be manually granted so that users who have permission of backup tables can access this directory.
- Data can be restored to other databases. In this case, HDFS permission must be manually granted so that users who have permission of backup tables can access the HDFS directory that corresponds to the database.

Step 10 Set **Force recovery** to **true**, which indicates to forcibly recover all backup data when a data table with the same name already exists. If the data table contains new data added after backup, the new data will be lost after the data recovery. If you set the parameter to **false**, the restoration task is not executed if a data table with the same name exists.

Step 11 Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified directory to be restored does not exist, the verification fails.
- If the forcible overwrite conditions are not met, the verification fails.

Step 12 Click **OK**.

Step 13 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

8.10.3.12 Restoring IoTDB Metadata

Scenario

To ensure IoTDB metadata security and prevent the IoTDB service from being unavailable due to IoTDB file damage, IoTDB metadata needs to be backed up. In this way, the system can restore data timely when an exception is reported or an operation does not achieve the expected result, minimizing the impact on services.

System administrators can create an IoTDB restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the IoTDB data that is generated after the data backup and before the data restoration will be lost.
- You are advised to restore the metadata of only one component in a restoration task to prevent the stop of a service or instance from affecting the data restoration of other components. If data of multiple components is restored at the same time, data restoration may fail.

Impact on the System

After the metadata is restored, the data generated after the data backup and before the data restoration is lost.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the row where the specified backup task is located, choose **More** > **View History** in the **Operation** column to display the historical execution records of the backup task.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object**: indicates the backup data source.
- **Backup Path**: indicates the full path where backup files are stored.

  Select the correct path, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, select **IoTDB** under **Metadata and other data**.

Step 8 Select a backup directory type for **Path Type**.

The configurations vary based on backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.

If you select this option, you also need to set **Source Path**, which indicates the backup file to be restored, for example, *Version_Data source_Task execution time*.**tar.gz**.

- **NFS**: indicates that backup files are stored in NAS using the NFS protocol.

  If you select this option, configure the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the complete path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time*.**tar.gz**.

- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select this option, configure the following parameters:

  - **Source NameService Name**: indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Source Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
  - **Source Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.
  - **Source NameNode RPC Port**: indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the standby cluster.
  - **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time*.**tar.gz**.

- **CIFS**: indicates that backup files are stored in NAS using the CIFS protocol.

  If you select this option, configure the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Source Path**: indicates the complete path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time*.**tar.gz**.

- **SFTP**: indicates that backup files are stored in the server using the SFTP protocol.

  If you select this option, configure the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the server where the backup data is stored.
  - **Port**: indicates the port number used to connect to the backup server over the SFTP protocol. The default value is **22**.
  - **Username**: indicates the username for connecting to the server using the SFTP protocol.
  - **Password**: indicates the password for connecting to the server using the SFTP protocol.
  - **Source Path**: indicates the complete path of the backup file on the backup server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time*.**tar.gz**.

Step 9 Click **OK**.

Step 10 In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

Step 11 Choose **Cluster** > **Services** and start the IoTDB service.

----End

8.10.3.13 Restoring IoTDB Service Data

Scenario

IoTDB service data needs to be restored in the following scenarios: Data is modified or deleted unexpectedly and needs to be restored. After an administrator performs major operations (such as upgrade and data adjustment) on IoTDB, an exception occurs or the expected result is not achieved. All modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create an IoTDB restoration task on FusionInsight Manager. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the IoTDB data that is generated after the data backup and before the data restoration will be lost.

Impact on the System

- During data restoration, user authentication stops and users cannot create new connections.
- After the data is restored, the data generated after the data backup and before the data restoration is lost.
- After the data is restored, the IoTDB upper-layer applications need to be started.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust must be configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.
- The IoTDB backup file save path is correct.
- The IoTDB upper-layer applications are stopped.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the row where the specified backup task is located, choose **More** > **View History** in the **Operation** column to display the historical execution records of the backup task.

In the window that is displayed, select a success record and click **View** in the **Backup Path** column to view its backup path information and find the following information:

- **Backup Object**: indicates the backup data source.
- **Backup Path**: indicates the full path where backup files are stored.

  Select the correct path, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In **Restoration Configuration**, choose **IoTDB** > **IoTDB** under **Service Data**.

Step 8 Set **Path Type** of **IoTDB** to a backup directory type.

The following backup directory types are supported:

RemoteHDFS: indicates that the backup files are stored in the HDFS directory of the standby cluster.

If you select this option, configure the following parameters:

- **Source NameService Name**: indicates the NameService name of the backup data cluster, for example, **hacluster**. You can obtain it from the **NameService Management** page of HDFS of the standby cluster.
- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source Active NameNode IP Address**: indicates the service plane IP address of the active NameNode in the standby cluster.
- **Source Standby NameNode IP Address**: indicates the service plane IP address of the standby NameNode in the standby cluster.
- **Source NameNode RPC Port**: indicates the value of **dfs.namenode.rpc.port** in the HDFS basic configuration of the standby cluster.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time*.
- **Recovery Point List**: Click **Refresh** and select an IoTDB directory that has been backed up in the standby cluster.

Step 9 In the **Backup Data** column of the **Data Configuration** page, select one or more pieces of backup data that needs to be restored based on service requirements. In the **Target Path** column, specify the target location after backup data restoration.

You are advised to set **Target Path** to a new path that is different from the backup path.

Step 10 Set **Force recovery** to **true**, which indicates that all backup data is forcibly restored when a data table with the same name already exists. If the data table contains new data added after backup, the new data will be lost after the data restoration. If you set the parameter to **false**, the restoration task is not executed if a data table with the same name exists.

Step 11 Click **Verify** to check whether the restoration task is configured correctly.

- If the queue name is incorrect, the verification fails.
- If the specified directory to be restored does not exist, the verification fails.

Step 12 Click **OK**.

Step 13 In the restoration task list, locate the row where the created task is located, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

----End

8.10.3.14 Restoring Kafka Metadata

Scenario

Kafka data needs to be recovered in the following scenarios: data is modified or deleted unexpectedly and needs to be restored. After an administrator performs critical data adjustment in ZooKeeper, an exception occurs or the operation has not achieved the expected result. All Kafka modules are faulty and become unavailable. Data is migrated to a new cluster.

System administrators can create a recovery task in FusionInsight Manager to recover Kafka data. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore Kafka metadata when the service is running properly, you are advised to manually back up the latest Kafka metadata before restoration. Otherwise, the Kafka metadata that is generated after the data backup and before the data restoration will be lost.

Impact on the System

- After the metadata is restored, the data generated after the data backup and before the data restoration is lost.
- After the metadata is restored, the offset information stored on ZooKeeper by Kafka consumers is rolled back, resulting in repeated consumption.

Prerequisites

- If you need to restore data from a remote HDFS, prepare a standby cluster. If the active cluster is deployed in security mode and the active and standby clusters are not managed by the same FusionInsight Manager, mutual trust has been configured. For details, see **Configuring Cross-Manager Mutual Trust Between Clusters**. If the active cluster is deployed in normal mode, no mutual trust is required.
- Cross-cluster replication has been configured for the active and standby clusters. For details, see **Enabling Cross-Cluster Replication**.
- Time is consistent between the active and standby clusters and the NTP services on the active and standby clusters use the same time source.

- The Kafka service is disabled first, and then enabled upon data restoration.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view historical backup task execution records.

In the displayed window, locate a specified success record and click **View** in the **Backup Path** column to view the backup path information of the task and find the following information:

- **Backup Object** specifies the data source of the backup data.
- **Backup Path** specifies the full path where the backup files are saved.

  Select the correct item, and manually copy the full path of backup files in **Backup Path**.

Step 3 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 4 Click **Create**.

Step 5 Set **Task Name** to the name of the restoration task.

Step 6 Select the cluster to be operated from **Recovery Object**.

Step 7 In the **Restoration Configuration** area, select **Kafka**.

Step 8 Set **Path Type** of **Kafka** to a backup directory type.

The settings vary according to backup directory types:

- **LocalDir**: indicates that the backup files are stored on the local disk of the active management node.

  If you select **LocalDir**, you also need to set **Source Path** to select the backup file to be restored, for example, *Version_Data source_Task execution time*.**tar.gz**.
- **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster.

  If you select **LocalHDFS**, set the following parameters:

  – **Source Path**: indicates the full path of the backup file in the HDFS, for example, *Backup path/Backup task name_Task creation time/Version_Data source_Task execution time*.**tar.gz**.

  – **Source NameService Name**: indicates the NameService name that corresponds to the backup directory when a restoration task is executed. The default value is **hacluster**.
- **RemoteHDFS**: indicates that the backup files are stored in the HDFS directory of the standby cluster.

  If you select **RemoteHDFS**, set the following parameters:

  – **Source NameService Name**: indicates the NameService name of the backup data cluster. You can enter the built-in NameService name of the remote cluster, for example, **haclusterX**, **haclusterX1**, **haclusterX2**,

haclusterX3, or **haclusterX4**. You can also enter a configured NameService name of the remote cluster.

- **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
- **Source NameNode IP Address**: indicates the NameNode service plane IP address of the standby cluster, supporting the active node or standby node.
- **Source Path**: indicates the full path of HDFS directory for storing backup data of the standby cluster, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **Queue Name**: indicates the name of the Yarn queue used for backup task execution. The name must be the same as the name of the queue that is running properly in the cluster.

- **NFS**: indicates that backup files are stored in NAS using the NFS protocol.

  If you select **NFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **CIFS**: indicates that backup files are stored in NAS using the CIFS protocol.

  If you select **CIFS**, set the following parameters:

  - **IP Mode**: indicates the mode of the target IP address. The system automatically selects the IP address mode based on the cluster network type, for example, **IPv4** or **IPv6**.
  - **Server IP Address**: indicates the IP address of the NAS server.
  - **Port**: indicates the port number used to connect to the NAS server over the CIFS protocol. The default value is **445**.
  - **Username**: indicates the username set when the CIFS protocol is configured.
  - **Password**: indicates the password set when the CIFS protocol is configured.
  - **Source Path**: indicates the full path of the backup file on the NAS server, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.
- **OBS**: indicates that backup files are stored in OBS.

  If you select **OBS**, set the following parameters:

  - **Source Path**: indicates the full OBS path of a backup file, for example, *Backup path/Backup task name_Data source_Task creation time/Version_Data source_Task execution time***.tar.gz**.

    NOTE

    Only MRS 3.1.0 or later supports saving backup files in OBS.

Step 9 Click **OK**.

Step 10 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and click **Retry** to execute the task again.

NOTICE

– If the Kafka service is deleted after the backup is complete, reinstall the Kafka service, restore its metadata, and restart the Kafka service. It is found that the Broker service cannot be started. In this case, the **/var/log/Bigdata/kafka/broker/server.log** file contains an error. An error example is as follows:

```
ERROR Fatal error during KafkaServer startup. Prepare to shutdown (kafka.server.KafkaServer)kafka.common.InconsistentClusterIdException: The Cluster ID kVSgfurUQFGGpHMTBqBPiw doesn't match stored clusterId Some(0Qftv9yBTAmf2iDPSlIk7g) in meta.properties. The broker is trying to join the wrong cluster. Configured zookeeper.connect may be wrong. at kafka.server.KafkaServer.startup(KafkaServer.scala:220) at kafka.server.KafkaServerStartable.startup(KafkaServerStartable.scala:44) at kafka.Kafka$.main(Kafka.scala:84) at kafka.Kafka.main(Kafka.scala)
```

Check the value of **log.dirs** in the Kafka Broker configuration file **${BIGDATA_HOME}/Fusionsight_Current/*Broker/etc/server.properties**. The value is the Kafka data directory. Go to the Kafka data directory and change the value **0Qftv9yBTAmf2iDPSlIk7g** of **cluster.id** in **meta.properties** to **kVSgfurUQFGGpHMTBqBPiw** (the latest value in the error log).

– The preceding modification must be performed on each node where Broker is located. After the modification, restart the Kafka service.

----End

8.10.4 Enabling Cross-Cluster Replication

Scenario

DistCp is used to replicate the data stored in HDFS from a cluster to another cluster. DistCp depends on the cross-cluster replication function, which is disabled by default. You need to enable it for both clusters.

This section describes how to modify parameters on FusionInsight Manager to enable the cross-cluster replication function. After this function is enabled, you can create a backup task for backing up data to the remote HDFS (RemoteHDFS).

Impact on the System

Yarn needs to be restarted to enable the cross-cluster replication function and cannot be accessed during restart.

Prerequisites

- The **hadoop.rpc.protection** parameter of HDFS in the two clusters for data replication must use the same data transmission mode. The default value is **privacy**, indicating encrypted transmission. The value **authentication** indicates that transmission is not encrypted.
- For clusters in security mode, you need to configure mutual trust between clusters.

Procedure

Step 1 Log in to FusionInsight Manager of one of the two clusters.

Step 2 Choose **Cluster** > **Services** > **Yarn** and click **Configurations** then **All Configurations**.

Step 3 In the navigation pane, choose **Yarn** > **Distcp**.

Step 4 Modify **dfs.namenode.rpc-address**, set **haclusterX.remotenn1** to the service IP address and RPC port of one NameNode instance of the peer cluster, and set **haclusterX.remotenn2** to the service IP address and RPC port number of the other NameNode instance of the peer cluster.

haclusterX.remotenn1 and **haclusterX.remotenn2** do not distinguish active and standby NameNodes. The default NameNode RPC port is 8020 and cannot be modified on Manager.

Examples of modified parameter values: **10.1.1.1:8020** and **10.1.1.2:8020**.

NOTE

- If data of the current cluster needs to be backed up to the HDFS of multiple clusters, you can configure the corresponding NameNode RPC addresses to haclusterX1, haclusterX2, haclusterX3, and haclusterX4.

Step 5 Click **Save**. In the confirmation dialog box, click **OK**.

Step 6 Restart the Yarn service.

Step 7 Log in to FusionInsight Manager of the other cluster and repeat **Step 2** to **Step 6**.

----End

8.10.5 Managing Local Quick Restoration Tasks

Scenario

When DistCp is used to back up data, the backup snapshot is saved to HDFS of the active cluster. FusionInsight Manager supports using the local snapshot for quick data restoration, requiring less time than restoring data from the standby cluster.

Use FusionInsight Manager and the snapshots on HDFS of the active cluster to create a local quick restoration task and execute the task.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the backup task list, locate a created task and click **Restore** in the **Operation** column.

Step 3 Check whether the system displays "No data is available for quick restoration. Create a task on the restoration management page to restore data".

- If yes, click **OK** to close the dialog box. No backup data snapshot is created in the active cluster, and no further action is required.
- If no, go to **Step 4** to create a local quick restoration task.

NOTE

Metadata does not support quick restoration.

Step 4 Set **Name** to the name of the local quick restoration task.

Step 5 Set **Configuration** to a data source.

Step 6 Set **Recovery Point List** to a recovery point that contains the backup data.

Step 7 Set **Queue Name** to the name of the Yarn queue used in the task execution. The name must be the same as the name of the queue that is running properly in the cluster.

Step 8 Set **Data Configuration** to the object to be recovered.

Step 9 Click **Verify**, and wait for the system to display "The restoration task configuration is verified successfully."

Step 10 Click **OK**.

Step 11 In the restoration task list, locate a created task and click **Start** in the **Operation** column to execute the restoration task.

After the task is complete, **Task Status** of the task is displayed as **Successful**.

----End

8.10.6 Modifying a Backup Task

Scenario

This section describes how to modify the parameters of a created backup task on FusionInsight Manager to meet changing service requirements. The parameters of restoration tasks can only be viewed but cannot be modified.

Impact on the System

After a backup task is modified, the new parameters take effect when the task is executed next time.

Prerequisites

- A backup task has been created.
- A new backup task policy has been planned based on the actual situation.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

Step 2 In the task list, locate a specified task, click **Configure** in the **Operation** column to go to the configuration modification page.

On the displayed page, modify the following parameters:

- Started
- Period
- Destination NameService Name
- Target NameNode IP Address
- Target Path
- Max Number of Backup Copies
- Maximum Number of Recovery Points
- Maximum Number of Maps
- Maximum Bandwidth of a Map

NOTE

After the **Target Path** parameter of a backup task is modified, this task will be performed as a full backup task for the first time by default.

Step 3 Click **OK** to save the settings.

----End

8.10.7 Viewing Backup and Restoration Tasks

Scenario

This section describes how to view created backup and recovery tasks and check their running status on FusionInsight Manager.

Prerequisites

You have logged in to FusionInsight Manager.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Backup and Restoration**.

Step 2 Click **Backup Management** or **Restoration Management**.

Step 3 In the task list, obtain the previous execution result in the **Task Status** and **Task Progress** column. Green indicates that the task is executed successfully, and red indicates that the execution fails.

Step 4 In the **Operation** column of a specified task in the task list, choose **More** > **View History** or click **View History** to view the historical record of backup and restoration task execution.

In the displayed window, click before a specified record to display log information about the execution.

----End

Related Tasks

- Starting a backup or restoration task

  In the task list, locate a specified task and choose **More** > **Back Up Now** or click **Start** in the **Operation** column to start a backup or restoration task that is ready or fails to be executed. Executed restoration tasks cannot be repeatedly executed.

- Stopping a backup or restoration task

  In the task list, locate a specified task and choose **More** > **Stop** or click **Stop** in the **Operation** column to stop a backup or restoration task that is running. After the task is successfully stopped, its **Task Status** changes to **Stopped**.

- Deleting a backup or restoration task

  In the **Operation** column of a specified task, choose **More** > **Delete** or click **Delete** and then click **OK** in the displayed dialog box to delete the backup and restoration task. Backup data will be reserved by default after a task is deleted.

- Suspending a backup task

  In the task list, locate a specified task and choose **More** > **Suspend** in the **Operation** column to suspend a backup task. Only periodic backup tasks can be suspended. Suspended backup tasks are no longer executed automatically. When you suspend a backup task that is being executed, the task execution stops. To resume a task, choose **More** > **Resume**.

8.10.8 How Do I Configure the Environment When I Create a ClickHouse Backup Task on FusionInsight Manager and Set the Path Type to RemoteHDFS?

NOTE

This section applies only to MRS 3.1.0.

Question

How do I configure the environment when I create a ClickHouse backup task on FusionInsight Manager and set the path type to RemoteHDFS?

Answer

Step 1 Log in to FusionInsight Manager of the standby cluster.

Step 2 Choose **Cluster** > **Services** > **HDFS** and choose **More** > **Download Client**. Set **Select Client Type** to **Configuration Files Only**, select **x86_64** for x86 or **aarch64**

for ARM based on the type of the node where the client is to be installed, and click **OK**.

Step 3 After the client file package is generated, download the client to the local PC as prompted and decompress the package.

For example, if the client file package is **FusionInsight_Cluster_1_HDFS_Client.tar**, decompress it to obtain **FusionInsight_Cluster_1_HDFS_ClientConfig_ConfigFiles.tar**, and then decompress **FusionInsight_Cluster_1_HDFS_ClientConfig_ConfigFiles.tar** to the **D:\FusionInsight_Cluster_1_HDFS_ClientConfig_ConfigFiles** directory on the local PC. The directory name cannot contain spaces.

Step 4 Go to the **FusionInsight_Cluster_1_HDFS_ClientConfig_ConfigFiles** client directory and obtain the **hosts** file.

Step 5 Go to **FusionInsight_Cluster_1_HDFS_ClientConfig_ConfigFiles\HDFS\config** to obtain the **core-site.xml** and **hdfs-site.xml** files.

Step 6 Log in to FusionInsight Manager of the source cluster.

Step 7 Choose **Cluster** > **Services** > **ClickHouse**, choose **Configurations** > **All Configurations**, and select **backup** under **ClickHouse(Service)**.

For **remote_connect_core_config_file**, click **Upload File** and select the **core-site.xml** file prepared in **Step 5**.

For **remote_connect_hdfs_config_file**, click **Upload File** and select the **hdfs-site.xml** file prepared in **Step 5**.

Step 8 Click **Save**, confirm the information, and click **OK** to save the configuration. After saving the configurations, click **Finish**.

Step 9 Choose **Cluster** > **Services** > **ClickHouse**, click **Instance**, and view the instance IP address of **ClickHouseServer**.

Step 10 Log in to the host nodes of the ClickHouseServer instances as user **root** and check whether the **/etc/hosts** file contains the host information in **Step 4**. If not, add the host information in **Step 4** to the **/etc/hosts** file.

----End

8.11 SQL Inspector

8.11.1 Overview

SQL engines in the big data field are emerging one after another. In addition to a wide range of solutions, some problems are exposed. For example, the quality of

SQL input statements is uneven, SQL problems are difficult to locate, and large SQL statements consume too many resources.

Low-quality SQL statements pose unexpected impacts on the data analysis platform, degrading system performance or platform stability.

NOTE

This function is supported only by MRS 3.3.0 or later.

Function Description

MRS allows you to configure inspection rules for mainstream SQL engines (Hive, Spark, HetuEngine, and ClickHouse). MRS can identify typical large SQL queries and low-quality SQL statements and intercepts them before execution or block them during execution. Users do not need to change how they submit SQL statements or change SQL syntax. Service modifications are not required and inspection is easy to implement.

- You can configure SQL inspection rules on the UI that also allows you to query and modify the rules.
- During query response and execution, each SQL engine proactively inspects SQL statements based on the rules.
- Administrators can select to display hints on, intercept, or block SQL statements. The system logs SQL inspection events in real time for SQL audit. O&M engineers can analyze the logs, evaluate SQL statement quality on the live network, detect target statements, and take effective measures.

SQL inspection rules are classified into the following types:

- Static interception: The system displays hints on or intercepts SQL statements based on SQL syntax rules.
- Dynamic interception: The system displays hints on or intercepts SQL statements based on rules of data table statistics and metadata information.
- Runtime Blocking: The system blocks SQL statements based on system states (such as CPU, memory, and I/O) during the runtime of the SQL statements.

SQL requests that meet the static and dynamic interception rules can be intercepted, and the system gives hints for processing the statements properly. If a SQL request meets the blocking rule, the system blocks the SQL task.

Rules and Restrictions

- A SQL inspection rule can be associated with multiple SQL engines, and different threshold parameters can be configured for each service.
- A SQL inspection rule can be associated with multiple tenants. A rule takes effect only for associated tenants.

8.11.2 Adding an SQL Inspection

Scenario

You can add rules for specified tenants and SQL engines on FusionInsight Manager. The system will display hints on, intercept, or block SQL requests matched by the rules.

NOTE

Exercise caution when you add or modify a SQL inspection rule for a cluster, enable a rule, and set the threshold. An improper rule may cause upper-layer service interruption.

Adding a Rule

Step 1 Log in to FusionInsight Manager as a user with the Manager administrator rights.

Step 2 Click **Cluster** and choose **SQL Inspector**. The **SQL Inspector** page is displayed.

You can click **View Supported Rules** to view all SQL inspection rules supported by the current cluster.

Step 3 Click **Add Rule**. After the password of the current user is verified, the **Add Rule** page is displayed.

Step 4 Set the required parameters and click **OK**.

Parameter	Description
Name	Name of a SQL inspection rule
ID	Rule ID For details about meaning of the rules corresponding to the IDs, see **Table 8-84**.
Tenant	Click **Add** to select the name of the tenant to which the current rule will be associated. If you need to add a new tenant, plan and create a cluster tenant by referring to **Tenant Resources**.

Parameter	Description
Services and Actions	Click **Add** to specify the SQL engine to which this rule will be associated with and set the threshold parameters of the rule. Each rule can be associated with one SQL engine. If you want to configure a rule for other SQL engines, add new rules. • **Service**: Select the SQL engine associated with the current rule. • If an SQL request meets the rule, the system performs the following operations: – **Hint**: Record logs and display a hint for handling the SQL request. If the rule has parameters, you need to configure the threshold. – **Intercept**: Intercept the SQL request that meets the rule. If the rule has parameters, you need to configure the threshold. – **Block**: Block the SQL request that meets the rule. If the rule has parameters, you need to configure the threshold. **NOTE** For static and dynamic interception rules, **Hint** and **Block** operations are supported. For blocking rules, only the **Block** operation is supported.

Step 5 View the added prevention rule on the **SQL Defense** page. The rule takes effect dynamically.

To adjust the current rule, click **Modify** in the **Operation** column of the row that contains the target rule. After the user password is verified, you can modify rule parameters.

Figure 8-60 Viewing SQL inspection rules

----End

MRS SQL Inspection Rules

Table 8-84 MRS SQL inspection rules

ID	Description	Engine	Threshold	Example SQL Statement
static_0001	Check whether the number of occurrences of count(distinct) in the SQL statement exceeds the limit.	• Hive • Spark • HetuEngine	Number of occurrences of count(distinct) Recommended value: **10**	SELECT COUNT(DISTINCT deviceId), COUNT(DISTINCT collDeviceId) FROM table GROUP BY deviceName, collDeviceName, collCurrentVersion;
static_0002	Check whether the not in <subquery> statement is used in the SQL statement.	• Hive • Spark • HetuEngine	N/A	SELECT * FROM Orders o WHERE Orders.Order_ID not in (Select Order_ID FROM HeldOrders h where h.order_id = o.order_id);
static_0003	Check whether the number of joins in the SQL statement exceeds the limit.	• Hive • Spark • HetuEngine	Number of joins Recommended value: **20**	N/A

ID	Description	Engine	Threshold	Example SQL Statement
static_0004	Check whether the number of union all times in the SQL statement exceeds the limit.	• Hive • Spark • HetuEngine	Number of union all times Recommended value: **20**	select * from tables t1 union all select * from tables t2 union all select * from tables t3 union all select * from tables t4 union all select * from tables t5 union all select * from tables t6 union all select * from tables t7 union all select * from tables t8 union all select * from tables t9;
static_0005	The number of subquery nesting layers exceeds the limit.	• Hive • Spark • HetuEngine	Maximum number of nested subqueries Recommended value: **20**	select * from (with temp1 as (select * from tables) select * from temp1);
static_0006	Check whether the length of the SQL statement string exceeds the upper limit.	• Hive • Spark • HetuEngine	Length of the SQL string, in KB Recommended value: **10**	N/A
static_0007	Check whether the Cartesian product exists when multiple tables are associated.	• Hive • Spark • HetuEngine	N/A	select * from A,B;
static_0008	Check whether alter table update operation is performed at the cluster level (on cluster).	ClickHouse	N/A	alter table testtb1 on cluster default_cluster update price=10.0 where id='100'

ID	Description	Engine	Threshold	Example SQL Statement
static_0009	Check whether alter table delete operation is performed at the cluster level (on cluster).	ClickHouse	N/A	alter table testtb1 on cluster default_cluster delete where id ='10'
static_0010	Check whether the alter table add column operation is performed at the cluster level (on cluster).	ClickHouse	N/A	alter table testtb1 on cluster default_cluster add column testc String
static_0011	Check whether the alter table drop column operation is performed at the cluster level (on cluster).	ClickHouse	N/A	alter table testtb1 on cluster default_cluster drop column testc
static_0012	Check whether the optimize final operation is performed at the cluster level (on cluster).	ClickHouse	N/A	optimize table testtb1 on cluster default_cluster final
static_0013	Check whether the drop table operation is performed at the cluster level (on cluster).	ClickHouse	N/A	drop table testtb1 on cluster default_cluster;
static_0014	Check whether the truncate table operation is performed at the cluster level (on cluster).	ClickHouse	N/A	truncate table testtb1 on cluster default_cluster;
dynamic_0001	Check whether the number of scanned files exceeds the limit.	• Hive • Spark • HetuEngine	Number of files that will be scanned or have been scanned Recommended value: **100,000**	SELECT ss_ticket_number FROM store_sales WHERE ss_ticket_number= 72291252 LIMIT 10;

ID	Description	Engine	Threshold	Example SQL Statement
dynamic_0002	Check whether the number of partitions involved in operations (select, delete, update, and alter) on a table exceeds the upper limit.	• Hive • Spark • HetuEngine • ClickHouse	Number of partitions involved in the delete or alter operation Recommended value: **10,000**	DELETE FROM table_name WHERE column_name = value
dynamic_0003	When the right table of a join is a distributed table, check whether the data volume of the right table exceeds the upper limit.	ClickHouse	Number of rows in the right table when the join operation is performed. Recommended value: **100,000,000**	SELECT name, text FROM table_1 JOIN table_2 ON table_1.Id = table_2.Id
running_0001	Check whether the number of result rows returned by the Select statement to the client exceeds the upper limit.	• Hive • Spark • HetuEngine • ClickHouse	Number of rows in the query result Recommended value: **100,000**	select * from table
running_0002	Check whether the peak memory usage of the SQL statement exceeds the absolute value limit.	• Hive • Spark • HetuEngine • ClickHouse	Memory occupied by SQL running, in MB	N/A
running_0003	Check whether the running duration of the SQL statement exceeds the upper limit.	• Hive • Spark • HetuEngine • ClickHouse	SQL running duration threshold, in seconds	N/A

ID	Description	Engine	Threshold	Example SQL Statement
running_0004	The amount of data scanned by the SQL statements.	• Hive • Spark • HetuEngine • ClickHouse	Amount of data scanned by the SQL statement, in GB Recommended value: **10,240**	N/A

8.11.3 Configuring Hive SQL Inspection

Scenario

You can configure rules for Hive SQL inspection on FusionInsight Manager and configure rule parameters as you need.

Prerequisites

- The cluster client that contains the Hive service has been installed in the **/opt/hadoopclient** directory.
- The Hive service of the cluster is running properly.
- For a cluster with Kerberos authentication enabled, a user with Hive operation permissions has been created.

Constraints

- By default, SQL inspection rules need 5 seconds to take effect dynamically. After the queue is modified, it takes 10 minutes for Hive inspection rules to be reloaded.
- Interception and blocking rules will interrupt SQL tasks, so you need to set parameters of these rules properly based on the site requirements.
- For the rule dynamic_0001 (the number of files scanned by SQL statements exceeds the threshold), when the Spark and Tez engines reach the threshold, interception logs are printed in Yarn task logs and cannot be output on the Beeline client.
- Blocking rules have execution latency. For example, if the running_0004 rule is used and the threshold of the scanned data volume is 10 GB, the statement may be blocked when the data volume is 15 GB or higher due to the determination period and task concurrency.

Procedure

Step 1 Log in to FusionInsight Manager, click **Cluster**, and choose **SQL Inspector**. The **SQL Inspector** page is displayed.

Step 2 Add rules for Hive by referring to **Adding an SQL Inspection**.

For details about the rules supported by the Hive SQL engine, see **MRS SQL Inspection Rules**.

For example, add a rule whose ID is **static_0001** to check whether count distinct appears more than two times in the SQL statement. If so, the system displays a hint.

Figure 8-61 Adding a Hive SQL inspection rule

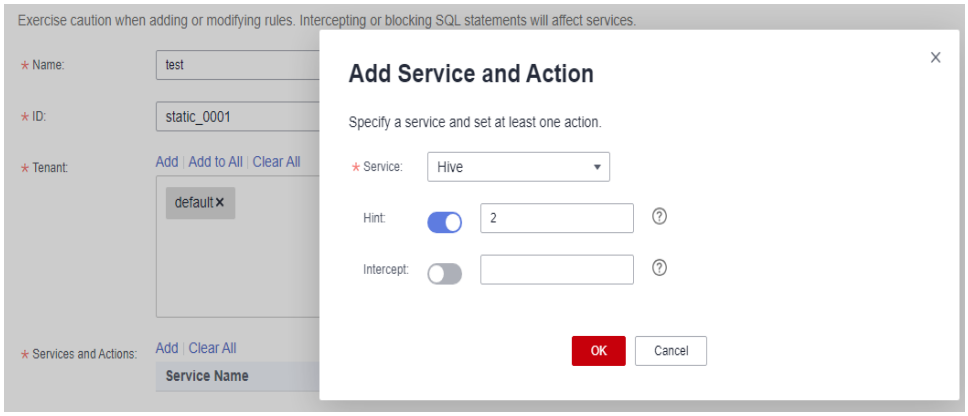

Step 3 Log in to the node where the Hive client is installed and run the following command to switch to the client installation directory.

cd /opt/hadoopclient

Run the following command to set environment variables:

source bigdata_env

Run the following command to authenticate the current user. Skip this step if Kerberos authentication is disabled for the cluster (the cluster is in normal mode).

kinit *Component service user who has the Hive operation permission*

Step 4 Run the following command to log in to the Hive client:

beeline

Step 5 Run the following commands to create a table and import data to the table.

drop table if exists hivetb;

create table hivetb(a int,b int);

insert into hivetb select 1,11;

insert into hivetb select 2,22;

Step 6 Run the following SQL statement to check whether the current rule takes effect:

select count(distinct a),count(distinct b) from hivetb;

If the number of times count distinct appears in the statement exceeds the threshold configured in **Step 2**, the following information is displayed:

```
...
WARN  : STATIC_0001 The count(distinct X) times exceeds the limit : 2, current count distinct times : 2
...
```


If the operation set in the rule is **Block**, the statement fails to be executed and the following information is displayed:

```
...
Error: Error while compiling statement: FAILED: RuleException STATIC_0001 The count(distinct X) times exceeds the limit : 2, current count distinct times : 2 (state=42000,code=40000)
...
```

NOTE

- For more Hive SQL inspection rules, see **MRS SQL Inspection Rules**.
- You can also obtain the SQL inspection rules via logs which are stored in **/var/log/Bigdata/audit/hive/hiveserver/queryinfo.log**.

----End

8.11.4 Configuring ClickHouse SQL Inspection

Scenario

You can configure rules for ClickHouse SQL inspection on FusionInsight Manager and configure rule parameters as you need.

Prerequisites

- The cluster client that contains the ClickHouse service has been installed in the **/opt/hadoopclient** directory.
- The ClickHouse logical cluster is running properly.
- For clusters with Kerberos authentication enabled, you need to create a service user who has the permission to operate the ClickHouse table. For example, create a human-machine user **clickhouseuser**.
- A tenant associated with the ClickHouse service has been created and associated with the ClickHouse service user. For details, see **Creating Tenants**

Constraints

- The default dynamic validity period of a rule is 1 minute.
- Interception and blocking rules will interrupt SQL queries, so you need to set parameters of these rules properly based on the site requirements.
- After configuring ClickHouse rules, you need to log in to the client again for the rules to take effect.

Procedure

Step 1 Log in to FusionInsight Manager, click **Cluster**, and choose **SQL Inspector**. The **SQL Inspector** page is displayed.

Step 2 Add rules for ClickHouse by referring to **Adding an SQL Inspection**.

For details about the rules supported by the ClickHouse SQL engine, see **MRS SQL Inspection Rules**.

For example, add a rule whose ID is **static_0008** and checks whether a SQL statement executes the cluster-level table update operation. If so, the system displays a hint.

Figure 8-62 Adding a ClickHouse SQL inspection rule

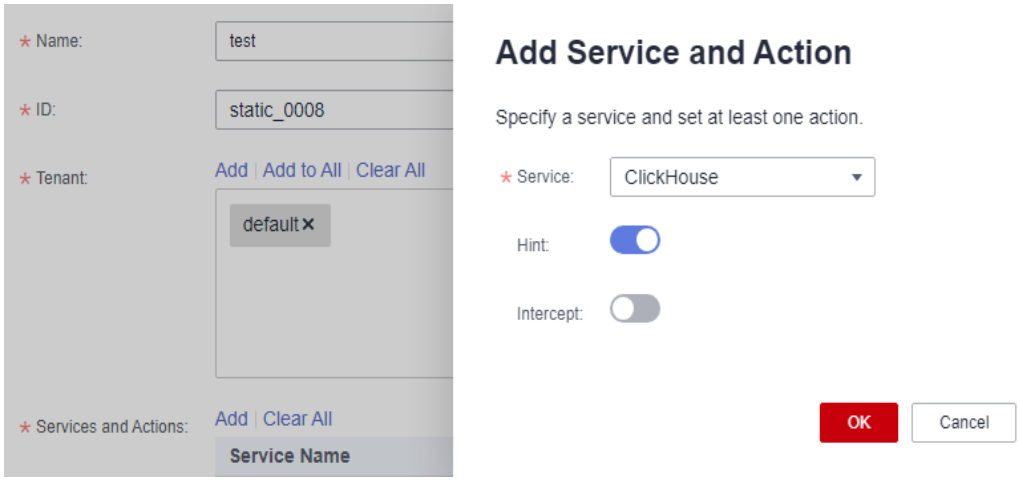

Step 3 Log in to the node where the ClickHouse client is installed and run the following command to switch to the client installation directory.

cd /opt/hadoopclient

Run the following command to set environment variables:

source bigdata_env

Step 4 If the current cluster is in security mode (Kerberos authentication is enabled), run the following command to authenticate the current user. The current user must have the permission to create ClickHouse tables. If the current cluster is in normal mode (Kerberos authentication is disabled), skip this step.

kinit *Component service user*

Example: **kinit clickhouseuser**

Step 5 Use the ClickHouse client to connect to the ClickHouse server.

Security mode

clickhouse client --host *IP address of the ClickHouseServer instance* **--port 9440 --secure**

Normal clusters:

clickhouse client --host *IP address of the ClickHouseServer instance*--**user** *Username* **--password --port 9000**

Enter the password.

Step 6 Run the following statements to create a data table:

CREATE DATABASE *cktest* **ON CLUSTER** *default_cluster*;

CREATE TABLE *cktest.test2* **ON CLUSTER** *default_cluster* **(`EventDate` DateTime, `CounterID` UInt32, `UserID` UInt32, `ver` UInt16) ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/***cktest/test2***', '{replica}') PARTITION BY toYYYYMM(EventDate) ORDER BY (EventDate, intHash32(UserID));**

CREATE TABLE *cktest.test2_dir* **ON CLUSTER** *default_cluster* **as** *cktest.test2* **ENGINE = Distributed(***default_cluster, cktest, test2***, rand());**

Step 7 Run the following command to insert data to the table:

insert into *cktest.test2* **values('2023-08-01',111,111,111);**

insert into *cktest.test2* **values('2023-08-02',222,111,111);**

Step 8 Run the following SQL statement for the created table to check whether the rule takes effect:

alter table *cktest.test2* **on cluster** *default_cluster* **update CounterID = toUInt32(222) where EventDate='2023-08-01' ;**

```
...
<Warning> SQLDefender: Distributed DDL ALTER UPDATE queries are undesirable.
...
```

If the operation set in the rule is **Intercept**, the statement fails to be executed and the following information is displayed:

```
...
DB::Exception: Distributed DDL ALTER TABLE UPDATE queries are undesirable..(QUERY_IS_PROHIBITED)
...
```

NOTE

For more ClickHouse SQL inspection rules, see **MRS SQL Inspection Rules**.

----End

8.11.5 Configuring HetuEngine SQL Inspection

Scenario

You can configure rules for HetuEngine SQL inspection on FusionInsight Manager and configure rule parameters as you need.

Prerequisites

- The cluster client that contains the HetuEngine service has been installed in the **/opt/hadoopclient** directory.
- The HetuEngine service and compute instances are running properly.
- If Kerberos authentication has been enabled for the cluster, you need to create a HetuEngine user and grant related permissions to the user. In addition, you need to use Ranger to assign the user the permission to manage databases, tables, and columns of the data source.

Constraints

- The default dynamic validity period of a rule is 5 minutes.
- Interception and blocking rules will interrupt SQL queries, so you need to set parameters of these rules properly based on the site requirements.
- Blocking rules are controlled by session-level parameters of the system. To configure blocking rules, service users must have the set session permission.
- For static rule **static_0003**, the total number of joins in queries does not include Semi joins and Anti joins.

- When prompt rules are configured for dynamic_0001 and dynamic_0002, prompt messages are recorded only in logs and are not displayed on the client.
- The client and server send asynchronous requests. For blocking rule **Running_0001**, after the server blocks the requests, the message "Query is gone " may be displayed on the client. In this case, you can view logs to check whether the requests are blocked.

Procedure

Step 1 Log in to FusionInsight Manager, click **Cluster**, and choose **SQL Inspector**. The **SQL Inspector** page is displayed.

Step 2 Add rules for HetuEngine by referring to **Adding an SQL Inspection**.

For details about the rules supported by the HetuEngine SQL engine, see **MRS SQL Inspection Rules**.

For example, add a rule whose ID is **static_0001** to check whether count distinct appears more than two times in the SQL statement. If so, the system displays a hint.

Figure 8-63 Adding a HetuEngine SQL inspection rule

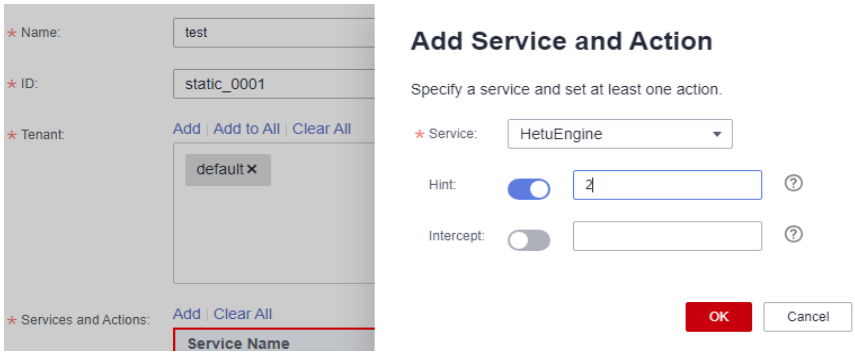

Step 3 Log in to the node where the HetuEngine client is installed and run the following command to switch to the client installation directory:

cd /opt/hadoopclient

Run the following command to set environment variables:

source bigdata_env

Step 4 Log in to the HetuEngine client based on the cluster authentication mode.

- In security mode, run the following command to authenticate the user and log in to the HetuEngine client:

  kinit hetu_test

  hetu-cli --catalog hive --tenant default --schema default

- In normal mode, run the following command to log in to the HetuEngine client:

  hetu-cli --catalog hive --tenant default --schema default --user hetu_test

NOTE

hetu_test is a service user who has at least the tenant role specified by **--tenant** and cannot be an OS user.

Step 5 Check whether the current rule takes effect.

Run the following statement to create a table:

CREATE TABLE table1(id int, name varchar,rank int);

INSERT INTO table1 VALUES(10,'sachin',1),(45,'rohit',2),(46,'rohit',3),(18,'virat',4),(25,'dhawan',5);

Run the following statement to query data:

select count(distinct id),count(distinct id),count(distinct id),count(distinct id),count(distinct id),count(distinct id) from table1;

If the number of times count distinct appears in the statement exceeds the threshold configured in **Step 2**, the following information is displayed:

```
WARNING: Occurrence number of 'COUNT(DISTINCT XX)' (6) reaches the hint limitation (2)
```

NOTE

- If the action set in the rule is **Intercept** or **Block**, the following information may be displayed:
  ```
  Intercepted. Reason: Occurrence number of 'COUNT(DISTINCT XX)' (6) reaches the interception limitation (2)
  ```
- You can query HetuEngine SQL inspection details in logs stored in **hdfs://hacluster/hetuserverhistory/tenant/coordinator/application_ID/container_ID/yyyyMMdd/server.log**.
- If warning information is required for JDBC secondary development, add the following configuration for the JDBC application:
  ```
  statement = connection.prepareStatement(sql.trim());
  resultSet = statement.executeQuery();
  SQLWarning sqlWarning = statement.getWarnings();
  ```

----End

8.11.6 Configuring Spark SQL Inspection

Scenario

You can configure rules for Spark SQL inspection on FusionInsight Manager and configure rule parameters as you need.

Prerequisites

- The cluster client that contains the Spark service has been installed in the **/opt/hadoopclient** directory.
- Spark is running properly.
- A tenant, for example, **sparkstatic1** has been added to **Tenant Resources**. For details, see **Creating Tenants**.
- For a cluster with Kerberos authentication enabled, a service user has been created, for example, user **sparkuser**. The user belongs to groups **hive**, **hadoop**, and **supergroup**, the primary group is **hive**, and the bound role is set to **sparkstatic1**.

Constraints

- The default dynamic validity period of a rule is 6 minutes.
- Only SQL jobs are supported.
- Interception and blocking rules will interrupt SQL queries, so you need to set parameters of these rules properly based on the site requirements.
- The static rule static_0007 is not required for Spark because it has a Cartesian product restriction (controlled by **spark.sql.crossJoin.enabled**, which is **true** by default). If the spark parameter is **false**, the static_0007 rule will not take effect.
- Dynamic rules do work on carbon tables.
- The dynamic rule dynamic_0002 supports SELECT, ALTER TABLE ADD PARTITION and ALTER TABLE DROP PARTITION. If you use a batch deletion statement that contains judgment conditions, for example, ALTER TABLE DROP PARTITION (pt < 10), the statement will be intercepted before the dynamic_0002 rule because the number of partitions is limited by **spark.sql.dropPartitionsInBatch.limit**, which is defaulted to **1000**.
- Blocking rules has execution latency. For example, if the running_0004 rule is used and the threshold of the scanned data volume is 10 GB, the statement may be blocked when the data volume is 15 GB or higher due to the determination period and task concurrency.
- If a job does not exceed the threshold until the last several tasks are complete before the block action is triggered, the job cannot be canceled.
- Blocking rule running_0004: The SQL execution duration includes the execution duration on the Driver side and the job running duration. When the SQL execution is blocked on the Driver side, the job cannot be canceled even if the execution duration exceeds the value of interruption threshold. This problem may occur when INSERT OVERWRITE is performed on a large number of partitions where storage and compute are decoupled.

Procedure

Step 1 Log in to FusionInsight Manager, click **Cluster**, and choose **SQL Inspector**. The **SQL Inspector** page is displayed.

Step 2 Add rules for Spark by referring to **Adding an SQL Inspection**.

For details about the rules supported by the Spark SQL engine, see **MRS SQL Inspection Rules**.

For example, add a rule whose ID is **static_0001** to check whether count distinct appears more than two times in the SQL statement. If so, the system displays a hint.

Figure 8-64 Adding a Spark SQL inspection rule

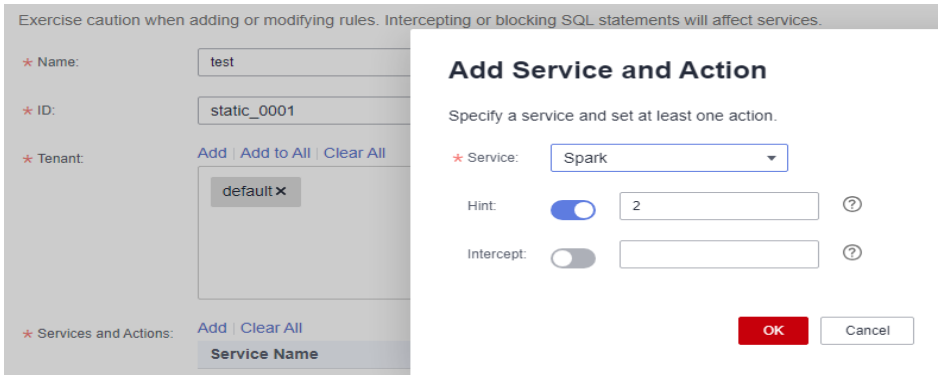

Step 3 Log in to the node where the Spark client is installed and run the following command to switch to the client installation directory.

cd /opt/hadoopclient

Run the following command to set environment variables:

source bigdata_env

source Spark/component_env

Step 4 Perform user authentication for clusters in security mode (Kerberos authentication enabled). Skip this step for clusters in normal mode (Kerberos authentication disabled).

kinit *Spark operation user*

Example:

kinit sparkuser

Enter the password as prompted and change the password upon your first login.

Step 5 Run the following command to log in to the spark-sql client:

cd opt/client/Spark/spark/bin

./spark-sql

Step 6 Run the following SQL statement on the client to check whether the current rule takes effect:

Run the following statement to create a table:

create table table1(id int, name string) stored as parquet

Run the following statement to query data:

select count(distinct id),count(distinct id),count(distinct id),count(distinct id),count(distinct id),count(distinct id) from table1;

If the number of times count distinct appears in the statement exceeds the threshold configured in **Step 2**, the following information is displayed:

```
WARNING:  static_0001 Occurrence num of 'COUNT(DISTINCT)'(6) reaches the hint threshold(2)
```

If the action set in the rule is **Intercept**, the following information is displayed:

```
Error in query: static_0001 Occurrence num of 'COUNT(DISTINCT)'(6) reaches the intercept threshold(2)
```


In Spark Beeline, you can obtain SQL inspection details from logs.

1. Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Yarn**. On the **Dashboard** page, click the link next to **ResourceManager WebUI** to enter the Yarn web UI.
2. Click the ID of the target application on the **All Applications** page. The application details page is displayed.

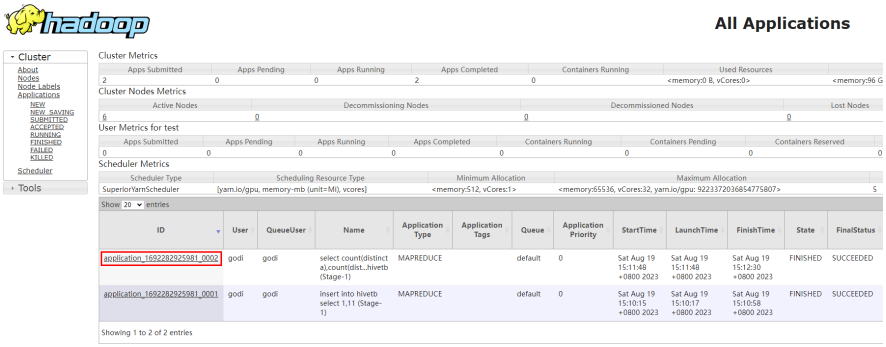

3. Click **Logs** of the application. On the displayed page, click **stdout** logs to view SQL inspection details.

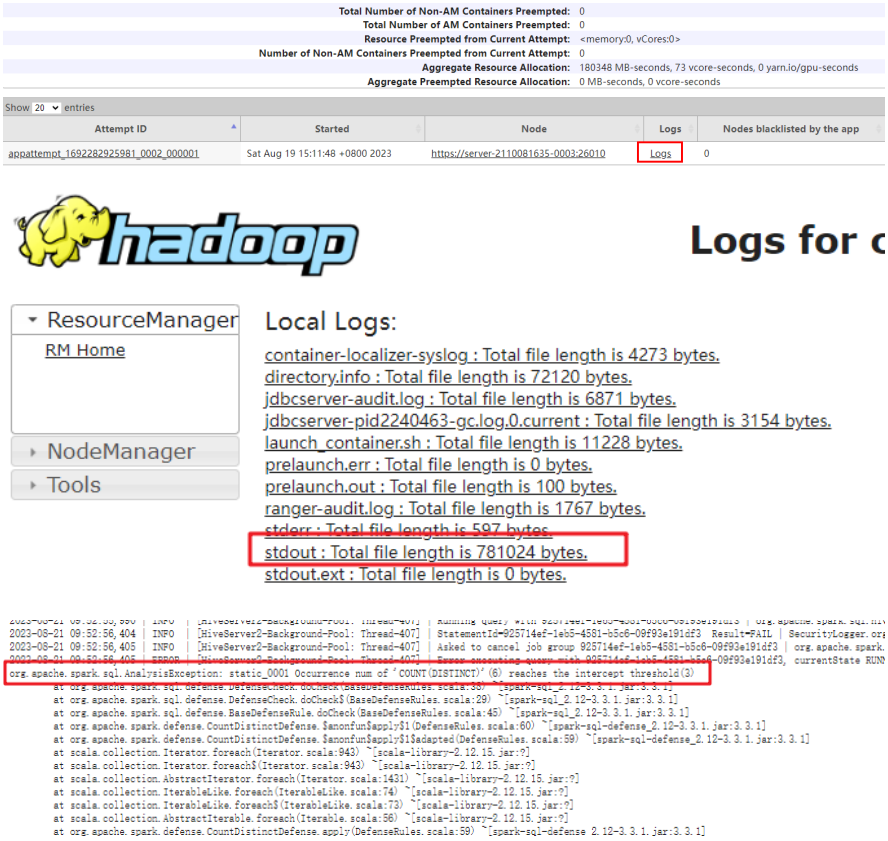

NOTE

1. For more Spark SQL inspection rules, see **MRS SQL Inspection Rules**.
2. You can view query info in the **/opt/hadoopclient/Spark/spark/audit/query.log** path if you are using the Spark client. The log contains detailed running information and the corresponding SQL inspection information.

----End

8.12 Security Management

8.12.1 Security Overview

8.12.1.1 Right Model

Role-based Access Control

FusionInsight adopts the role-based access control (RBAC) mode to manage rights on the big data system. It integrates the right management functions of the components to centrally manage rights. Common users are shielded from internal right management details, and the right management operations are simplified for administrators, improving right management usability and user experience.

The right model of FusionInsight consists four parts, that is users, user groups, roles, and rights.

Figure 8-65 Right model

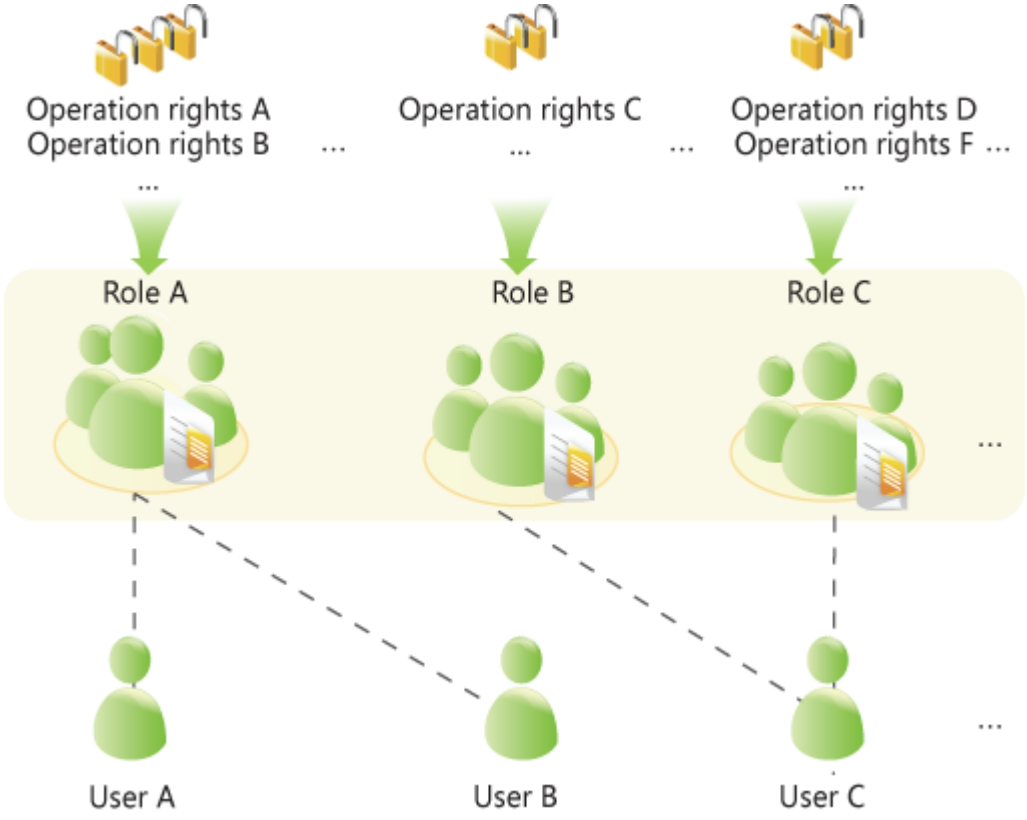

- **Right**

  Right, which is defined by components, allows users to access a certain resource of one component. Different components have different rights for their resources.

  For example:

  - HDFS provides read, write, and execute permissions on files.
  - HBase provides create, read, and write permissions on tables.

- **Role**

  Role is a collection of component rights. Each role can have multiple rights of multiple components. Different roles can have the rights of a resource of one component.

- **User group**

  User group is a collection of users. When a user group is bound to a role, users in this group obtain the rights defined by the role.

  Different user groups can be associated with the same role. A user group can also be associated with no role, and this user group does not have the rights of any component resources.

  NOTE

  In some components, the system grants related rights to specific user groups by default.

- **User**

  A user is a visitor to the system. Each user has the rights of the user group and role associated with the user. Users need to be added to the user group or associated with roles to obtain the corresponding rights.

Policy-based Access Control

The Ranger component uses policy-based access control (PBAC) to manage rights and implement fine-grained data access control on components such as HDFS, Hive, and HBase.

NOTE

The component supports only one right control mechanism. After the Ranger right control policy is enabled for the component, the right on the component in the role created on FusionInsight Manager becomes invalid (The ACL rules of HDFS and Yarn still take effect). You need to add a policy on the Ranger management page to grant rights on resources.

The Ranger right model consists of multiple right policies. A right policy consists of the following parts:

- Resource

  Resources are provided by components and can be accessed by users, such as HDFS files or folders, queues in Yarn, and databases, tables, and columns in Hive.

- User

  A User is a visitor to the system. The rights of each user are obtained based on the policy associated with the user. Information about users, user groups, and roles in the LDAP is periodically synchronized to the Ranger.

- Permission

  In a policy, you can configure various access conditions for resources, such as file read and write, permission conditions, rejection conditions, and exception conditions.

8.12.1.2 Right Mechanism

FusionInsight adopts the Lightweight Directory Access Protocol (LDAP) to store data of users and user groups. Information about role definitions is stored in the relational database and the mapping between roles and rights is saved in components.

FusionInsight uses Kerberos for unified authentication.

The verification process of user rights is as follows:

1. A client (a user terminal or FusionInsight component service) invokes the FusionInsight authentication interface.
2. FusionInsight uses the login username and password for Kerberos authentication.
3. If the authentication succeeds, the client sends a request for accessing the server (a FusionInsight component service).
4. The server finds the user group and role to which the login user belongs.
5. The server obtains all rights of the user group and the role.
6. The server checks whether the client has the right to access the resources it applies for.

Example (RBAC):

There are three files in HDFS, that is, fileA, fileB, and fileC.

- roleA has read and write right for fileA, and roleB has the read right for fileB.
- groupA is bound to roleA, and groupB is bound to roleB.
- userA belongs to groupA and roleB, and userB belongs to groupB.

When userA successfully logs in to the system and accesses the HDFS:

1. HDFS obtains the role (roleB) to which userA is bound.
2. HDFS also obtains the role (roleA) to which the user group of userA is bound.
3. In this case, userA has all the rights of roleA and roleB.
4. As a result, userA has read and write rights for fileA, has the read right on fileB, and has no right for fileC.

Similarly, when userB successfully logs in to the system and accesses the HDFS:

1. userB only has the rights of roleB.
2. As a result, userB has the read right on fileB, and has no rights for fileA and fileC.

8.12.1.3 Authentication Policies

The big data platform performs user identity authentication to prevent invalid users from accessing the cluster. The cluster provides authentication capabilities in both security mode and normal mode.

Security Mode

The clusters in security mode use the Kerberos authentication protocol for security authentication. The Kerberos protocol supports mutual authentication between

clients and servers. This eliminates the risks incurred by sending user credentials over the network for simulated authentication. In clusters, KrbServer provides the Kerberos authentication support.

Kerberos user object

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

Kerberos authentication

Kerberos supports password and keytab authentication. The validity period of authentication is 24 hours by default.

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

Normal Mode

Different components in a normal cluster use the native open-source authentication mode and do not support the **kinit** authentication command. FusionInsight Manager (including DBService, KrbServer, and LdapServer) uses the username and password for authentication. **Table 8-85** lists the authentication modes used by components.

Table 8-85 Component authentication modes

Service	Authentication Mode
IoTDB	Simple authentication
CDL	No authentication
ClickHouse	Simple authentication
Flume	No authentication
HBase	• Web UI: No authentication • Client: simple authentication
HDFS	• Web UI: no authentication • Client: simple authentication

Service	Authentication Mode
HetuEngine	• Web UI: no authentication • Client: no authentication
Hive	Simple authentication
Hue	Username and password authentication
Kafka	No authentication
Loader	• Web UI: username and password authentication • Client: no authentication
MapReduce	• Web UI: no authentication • Client: no authentication
Oozie	• Web UI: username and password authentication • Client: simple authentication
Spark2x	• Web UI: no authentication • Client: simple authentication
Storm	No authentication
YARN	• Web UI: no authentication • Client: simple authentication
ZooKeeper	Simple authentication

The authentication modes are as follows:

- Simple authentication: When the client connects to the server, the client automatically authenticates the user (for example, the OS user **root** or **omm**) by default. The authentication is imperceptible to the administrator or service user, which does not require **kinit**.
- Username and password authentication: Use the username and password of human-machine users in the cluster for authentication.
- No authentication: Any user can access the server by default.

8.12.1.4 Permission Verification Policies

Security Mode

After a user is authenticated by the big data platform, the system determines whether to verify the user's permission based on the actual permission management configuration to ensure that the user has limited or all permission on resources. If the user does not have the permission for accessing cluster resources, the system administrator must grant the required permission to the user. Otherwise, the user fails to access the resources. The cluster provides permission verification capabilities in both security mode and normal mode. The specific permission items of the components are the same in the two modes.

By default, the Ranger service is installed and Ranger authentication is enabled for a newly installed cluster in security mode. You can set fine-grained security access policies for accessing component resources through the permission plug-in of the component. If Ranger authentication is not required, administrators can manually disable it on the service page. After Ranger authentication is disabled, the system continues to perform permission control based on the role model of FusionInsight Manager when accessing component resources.

In a cluster in security mode, the following components support Ranger authentication: HDFS, YARN, Kafka, Hive, HBase, Storm, Impala, HetuEngine, CDL, and Spark2x.

For a cluster upgraded from an earlier version, Ranger authentication is not used by default when users access component resources. The administrator can manually enable Ranger authentication after installing Ranger.

By default, all components in the cluster of the security edition authenticate access. The authentication function cannot be disabled.

Normal Mode

Different components in a normal cluster use their own native open-source authentication behavior. **Table 8-86** lists detailed permission verification modes.

In a normal cluster, Ranger supports permission control on component resources based on OS users. The following components support Ranger authentication: HBase, HDFS, Hive, Spark2x, and YARN.

Table 8-86 Component permission verification modes in normal clusters

Service	Permission Verification	Permission Verification Enabling and Disabling
IoTDB	Required	Not supported
ClickHouse	Required	Not supported
Flume	Not required	Not supported
HBase	Not required	Supported
HDFS	Required	Supported
HetuEngine	Not required	Not supported
Hive	Not required	Not supported
Hue	Not required	Not supported
Kafka	Not required	Not supported
Loader	Not required	Not supported
MapReduce	Not required	Not supported
Oozie	Required	Not supported
Spark2x	Not required	Not supported

Service	Permission Verification	Permission Verification Enabling and Disabling
Storm	Not required	Not supported
YARN	Not required	Supported
ZooKeeper	Required	Supported
CDL	Not required	Not supported

Condition Priorities of the Ranger Permission Policy

When configuring a permission policy for a resource, you can configure Allow Conditions, Exclude from Allow Conditions, Deny Conditions, and Exclude from Deny Conditions for the resource, to meet unexpected requirements in different scenarios.

The priorities of different conditions are listed in descending order: Exclude from Deny Conditions > Deny Conditions > Exclude from Allow Conditions > Allow Conditions

The following figure shows the process of determining condition priorities. If the component resource request does not match the permission policy in Ranger, the system rejects the access by default. However, for HDFS and Yarn, the system delivers the decision to the access control layer of the component for determination.

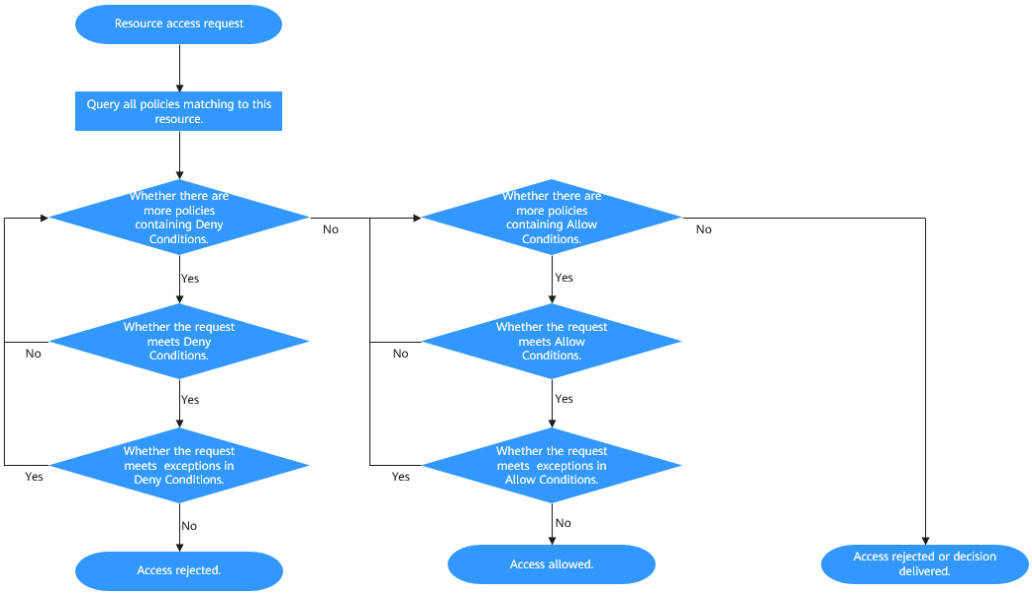

For example, if you want to grant the read and write permissions of the **FileA** folder to the **groupA** user group, but the user in the group is not **UserA**, you can add an allowed condition and an exception condition.

8.12.1.5 User Account List

User Classification

The MRS cluster provides the following three types of users. The system administrator needs to periodically change the passwords. It is not recommended to use the default passwords.

NOTE

This section describes information about default users in MRS clusters.

User Type	Description
System users	• User created on FusionInsight Manager for O&M and service scenarios. There are two types of users: – **Human-machine** user: used in scenarios such as FusionInsight Manager O&M and operations on a component client. When creating a user of this type, you need to set password and confirm password by referring to **Creating a User**. – **Machine-machine** user: used for system application development. • User who runs OMS processes
Internal system users	Internal user to perform Kerberos authentication, process communications, save user group information, and associate user permissions. It is recommended that internal system users not be used in O&M scenarios. Operations can be performed as user **admin** or another user created by the system administrator based on service requirements.
Database users	• User who manages OMS database and accesses data • User who runs service components (Hue, Hive, HetuEngine, Loader, Oozie, Ranger, and DBService) in the database.

System Users

NOTE

- User **root** of the OS is required, the password of user **root** on all nodes must be the same.
- User **ldap** of the OS is required. Do not delete this account. Otherwise, the cluster may not work properly. The OS administrator maintains the password management policies.

User Type	Username	Initial Password	Description	Password Change Method
System administrator	admin	User-defined password	FusionInsight Manager administrator. **NOTE** By default, user **admin** does not have the management permission on other components. For example, when accessing the native UI of a component, the user fails to access the complete component information due to insufficient management permission on the component.	For details, see **Changing the Password for User admin**.
Node OS user	ommdba	Random password	User that creates the system database. This user is an OS user generated on the management node and does not require a unified password. This account cannot be used for remote login.	For details, see **Changing the Password for an OS User**.
	omm	Random password	Internal running user of the system. This user is an OS user generated on all nodes and does not require a unified password.	

Internal System Users

User Type	Default User	Initial Password	Description	Password Change Method
Kerberos administrator	kadmin/admin	Admin@123	Used to add, delete, modify, and query user accounts on Kerberos.	For details, see **Changing the Password for the Kerberos Administrator**.

User Type	Default User	Initial Password	Description	Password Change Method
OMS Kerberos administrator	kadmin/admin	Admin@123	Used to add, delete, modify, and query user accounts on OMS Kerberos.	For details, see **Changing the Password for the OMS Kerberos Administrator**.
LDAP administrator	cn=root,dc=hadoop,dc=com	• Versions earlier than MRS 3.1.2: LdapChangeMe@123 • MRS 3.1.2 or later: randomly generated by the system	Used to add, delete, modify, and query the user account information on LDAP.	• For versions earlier than MRS 3.1.2, see **Changing the Passwords of the LDAP Administrator and the LDAP User (Including OMS LDAP)**. • For MRS 3.1.2 or later, see **Modifying OMS Service Configuration Parameters**.

User Type	Default User	Initial Password	Description	Password Change Method
OMS LDAP administrator	cn=root,dc=hadoop,dc=com	• Versions earlier than MRS 3.1.2: LdapChangeMe@123 • MRS 3.1.2 or later: randomly generated by the system	Used to add, delete, modify, and query the user account information on OMS LDAP.	
LDAP user	cn=pg_search_dn,ou=Users,dc=hadoop,dc=com	Randomly generated by the system	Used to query information about users and user groups on LDAP.	
OMS LDAP user	cn=pg_search_dn,ou=Users,dc=hadoop,dc=com	Randomly generated by the system	Used to query information about users and user groups on OMS LDAP.	

User Type	Default User	Initial Password	Description	Password Change Method
LDAP administrator account	cn=krbkdc,ou=Users,dc=hadoop,dc=com	• Versions earlier than MRS 3.1.2: LdapChangeMe@123 • MRS 3.1.2 or later: randomly generated by the system	Used to query Kerberos component authentication account information.	• For versions earlier than MRS 3.1.2, see **Changing the Password for the LDAP Administrator**. • For MRS 3.1.2 or later, see **Modifying OMS Service Configuration Parameters**.
	cn=krbadmin,ou=Users,dc=hadoop,dc=com	• Versions earlier than MRS 3.1.2: LdapChangeMe@123 • MRS 3.1.2 or later: randomly generated by the system	Used to add, delete, modify, and query Kerberos component authentication account information.	

User Type	Default User	Initial Password	Description	Password Change Method
Component running user	hdfs	Hdfs@123	This user is the HDFS system administrator and has the following permissions: 1. File system operation permissions: • Views, modifies, and creates files. • Views and creates directories. • Views and modifies the groups where files belong. • Views and sets disk quotas for users. 2. HDFS management operation permissions: • Views the web UI status. • Views and sets the active and standby HDFS status. • Enters and exits the HDFS in security mode. • Checks the HDFS file system. 3. Logs in to the FTP service page.	For details, see **Changing the Password for a Component Running User**.

User Type	Default User	Initial Password	Description	Password Change Method
	hbase	Hbase@123	This user is the HBase and HBase1 to HBase4 system administrator and has the following permissions: ● Cluster management permission: Performs **Enable** and **Disable** operations on tables to trigger MajorCompact and ACL operations. ● Grants and revokes permissions, and shuts down the cluster. ● Table management permission: Creates, modifies, and deletes tables. ● Data management permission: Reads data in tables, column families, and columns. ● Logs in to the HMaster web UI. ● Logs in to the FTP service page.	
	cdl	CDCUser123!	System administrator of the CDL Currently, CDL does not involve user permissions.	

User Type	Default User	Initial Password	Description	Password Change Method
	iotdb	Iotdb@123	This user is the IoTDB system administrator and has the following user permissions: 1. IoTDB administrator permissions: ● Creates or deletes a storage group. ● Uses TTL. 2. IoTDB data operation permissions: ● Creates, modifies, and deletes a time sequence. ● Writes, reads, and deletes data in a time sequence. 3. Views user or role permission information. 4. Grants or revokes permissions to or from a user or role.	

User Type	Default User	Initial Password	Description	Password Change Method
			NOTE In a normal cluster, the IoTDB service retains the features of open-source versions. The default username is **root**, and the default password is **root** (for MRS 3.3.0 or later, the default password is **Iotdb@123**). This user is an administrator and has all permissions, which cannot be assigned, revoked, or deleted.	
	mapred	Mapred@123	This user is the MapReduce system administrator and has the following permissions: ● Submits, stops, and views the MapReduce tasks. ● Modifies the Yarn configuration parameters. ● Logs in to the FTP service page. ● Logs in to the Yarn web UI.	

User Type	Default User	Initial Password	Description	Password Change Method
	zookeeper	ZooKeeper@123	This user is the ZooKeeper system administrator and has the following permissions: ● Adds, deletes, modifies, and queries all nodes in ZooKeeper. ● Modifies and queries quotas of all nodes in ZooKeeper.	
	rangeradmin	Rangeradmin@123	This user has the Ranger system management permissions and user permissions: ● Ranger web UI management permission ● Management permission of each component that uses Ranger authentication	
	rangeraudit or	Rangerauditor@123	Default audit user of the Ranger system.	

User Type	Default User	Initial Password	Description	Password Change Method
	hive	Hive@123	This user is the Hive system administrator and has the following permissions: 1. Hive administrator permissions: • Creates, deletes, and modifies a database. • Creates, queries, modifies, and deletes a table. • Queries, inserts, and uploads data. 2. HDFS file operation permissions: • Views, modifies, and creates files. • Views and creates directories. • Views and modifies the groups where files belong. 3. Submits and stops the MapReduce tasks. 4. Ranger policy management permission	

User Type	Default User	Initial Password	Description	Password Change Method
	kafka	Kafka@123	This user is the Kafka system administrator and has the following permissions: ● Creates, deletes, produces, and consumes the topic; modifies the topic configuration. ● Controls the cluster metadata, modifies the configuration, migrates the replica, elects the leader, and manages ACL. ● Submits, queries, and deletes the consumer group offset. ● Queries the delegation token. ● Queries and submits the transaction.	
	storm	Admin@123	Storm system administrator User permission: Submits Storm tasks.	
	rangerusersync	Randomly generated by the system	Synchronizes users and internal users of user groups.	
	rangertagsync	Randomly generated by the system	Internal user for synchronizing tags.	

User Type	Default User	Initial Password	Description	Password Change Method
	oms/ manager	Randomly generated by the system	Controller and NodeAgent authentication user. The user has the permission on the **supergroup** group.	
	backup/ manager	Randomly generated by the system	User for running backup and restoration tasks. The user has the permission on the **supergroup**, **wheel**, and **ficommon** groups. After cross-system mutual trust is configured, the user has the permission to access data in the HDFS, HBase, Hive, and ZooKeeper systems.	

User Type	Default User	Initial Password	Description	Password Change Method
	hdfs/ hadoop.*<System domain name>*	Randomly generated by the system	This user is used to start the HDFS and has the following permissions: 1. File system operation permissions: ● Views, modifies, and creates files. ● Views and creates directories. ● Views and modifies the groups where files belong. ● Views and sets disk quotas for users. 2. HDFS management operation permissions: ● Views the web UI status. ● Views and sets the active and standby HDFS status. ● Enters and exits the HDFS in security mode. ● Checks the HDFS file system. 3. Logs in to the FTP service page.	

User Type	Default User	Initial Password	Description	Password Change Method
	hetuserver/hadoop.<*System domain name*>	Randomly generated by the system	This user is used to start HetuEngine and has the following permissions: ● Accesses KrbServer and HDFS files in the cluster from HetuEngine. ● Used for communication between HetuEngine internal nodes.	
	mapred/hadoop.<*System domain name*>	Randomly generated by the system	This user is used to start the MapReduce and has the following permissions: ● Submits, stops, and views the MapReduce tasks. ● Modifies the Yarn configuration parameters. ● Logs in to the FTP service page. ● Logs in to the Yarn web UI.	
	mr_zk/hadoop.<*System domain name*>	Randomly generated by the system	Used for MapReduce to access ZooKeeper.	
	hbase/hadoop.<*System domain name*>	Randomly generated by the system	User for the authentication between internal components during the HBase system startup.	

User Type	Default User	Initial Password	Description	Password Change Method
	hbase/zkclient.<*System domain name*>	Randomly generated by the system	User for HBase to perform ZooKeeper authentication in a security mode cluster.	
	thrift/hadoop.<*System domain name*>	Randomly generated by the system	ThriftServer system startup user.	
	thrift/<*hostname*>	Randomly generated by the system	User for the ThriftServer system to access HBase. This user has the read, write, execution, creation, and administration permission on all NameSpaces and tables of HBase. <*hostname*> indicates the name of the host where the ThriftServer node is installed in the cluster.	

User Type	Default User	Initial Password	Description	Password Change Method
	hive/hadoop.<*System domain name*>	Randomly generated by the system	User for the authentication between internal components during the Hive system startup. The user permissions are as follows: 1. Hive administrator permissions: ● Creates, deletes, and modifies a database. ● Creates, queries, modifies, and deletes a table. ● Queries, inserts, and uploads data. 2. HDFS file operation permissions: ● Views, modifies, and creates files. ● Views and creates directories. ● Views and modifies the groups where files belong. 3. Submits and stops the MapReduce tasks.	
	loader/hadoop.<*System domain name*>	Randomly generated by the system	User for Loader system startup and Kerberos authentication	

User Type	Default User	Initial Password	Description	Password Change Method
	HTTP/*<hostname>*	Randomly generated by the system	Used to connect to the HTTP interface of each component. *<hostname>* indicates the host name of a node in the cluster.	
	hue	Randomly generated by the system	User for Hue system startup, Kerberos authentication, and HDFS and Hive access	
	flume	Randomly generated by the system	User for Flume system startup and HDFS and Kafka access. The user has read and write permission of the HDFS directory **/flume**.	
	flume_server	Randomly generated by the system	User for Flume system startup and HDFS and Kafka access. The user has read and write permission of the HDFS directory **/flume**.	
	spark2x/hadoop.*<System domain name>*	Randomly generated by the system	This user is the Spark2x system administrator and has the following user permissions: 1. Starts the Spark2x service. 2. Submits Spark2x tasks.	
	spark_zk/hadoop.*<System domain name>*	Randomly generated by the system	Used for Spark2x to access ZooKeeper.	

User Type	Default User	Initial Password	Description	Password Change Method
	zookeeper/hadoop.*<System domain name>*	Randomly generated by the system	ZooKeeper system startup user.	
	zkcli/hadoop.*<System domain name>*	Randomly generated by the system	ZooKeeper server login user.	
	oozie	Randomly generated by the system	User for Oozie system startup and Kerberos authentication.	
	kafka/hadoop.*<System domain name>*	Randomly generated by the system	Used for security authentication of Kafka.	
	storm/hadoop.*<System domain name>*	Randomly generated by the system	Storm system startup user.	
	storm_zk/hadoop.*<System domain name>*	Randomly generated by the system	Used for the Worker process to access ZooKeeper.	
	flink/hadoop.*<System domain name>*	Randomly generated by the system	Internal user of the Flink service.	
	check_ker_M	Randomly generated by the system	User who performs a system internal test about whether the Kerberos service is normal.	

User Type	Default User	Initial Password	Description	Password Change Method
	cdl/hadoop.*<System domain name>*	Randomly generated by the system	Internal user of the CDL service.	
	clickhouse/hadoop.*<System domain name>*	Randomly generated by the system	Used for security authentication of ClickHouse. This user is an internal user and can be used only in the cluster.	
	default	None	ClickHouse internal user, which is an administrator user that can be used only in non-security mode.	
	rangeradmin/hadoop.*<System domain name>*	Randomly generated by the system	Ranger system startup user, which is used for authentication between internal components.	
	tez	Randomly generated by the system	User for TezUI system startup, Kerberos authentication, and access to Yarn	
	K/M	Randomly generated by the system	Kerberos internal functional user. It cannot be deleted, and its password cannot be changed. This internal account can only be used on nodes where Kerberos service is installed.	None
	kadmin/changepw	Randomly generated by the system		
	kadmin/history	Randomly generated by the system		

User Type	Default User	Initial Password	Description	Password Change Method
	krbtgt<*System domain name*>	Randomly generated by the system		
LDAP user	admin	None	FusionInsight Manager administrator. The primary group is **compcommon**, which does not have the group permission but has the permission of the **Manager_administrator** role.	The LDAP user cannot log in to the system, and the password cannot be changed.
	backup		The primary group is **compcommon**.	
	backup/manager		The primary group is **compcommon**.	
	oms		The primary group is **compcommon**.	
	oms/manager		The primary group is **compcommon**.	
	clientregister		The primary group is **compcommon**.	
	zookeeper		The primary group is **hadoop**.	
	zookeeper/hadoop.<*System domain name*>		The primary group is **hadoop**.	
	zkcli		The primary group is **hadoop**.	
	zkcli/hadoop.<*System domain name*>		The primary group is **hadoop**.	

User Type	Default User	Initial Password	Description	Password Change Method
	flume		The primary group is **hadoop**.	
	flume_server		The primary group is **hadoop**.	
	hdfs		The primary group is **hadoop**.	
	hdfs/hadoop.*<System domain name>*		The primary group is **hadoop**.	
	mapred		The primary group is **hadoop**.	
	mapred/hadoop.*<System domain name>*		The primary group is **hadoop**.	
	mr_zk		The primary group is **hadoop**.	
	mr_zk/hadoop.*<System domain name>*		The primary group is **hadoop**.	
	hue		The primary group is **supergroup**.	
	hive		The primary group is **hive**.	
	hive/hadoop.*<System domain name>*		The primary group is **hive**.	
	hbase		The primary group is **hadoop**.	

User Type	Default User	Initial Password	Description	Password Change Method
	hbase/hadoop.*<System domain name>*		The primary group is **hadoop**.	
	thrift		The primary group is **hadoop**.	
	thrift/hadoop.*<System domain name>*		The primary group is **hadoop**.	
	oozie		The primary group is **hadoop**.	
	hbase/zkclient.*<System domain name>*		The primary group is **hadoop**.	
	loader		The primary group is **hadoop**.	
	loader/hadoop.*<System domain name>*		The primary group is **hadoop**.	
	spark2x		The primary group is **hadoop**.	
	spark2x/hadoop.*<System domain name>*		The primary group is **hadoop**.	
	spark_zk		The primary group is **hadoop**.	
	kafka		The primary group is **kafkaadmin**.	

User Type	Default User	Initial Password	Description	Password Change Method
	kafka/hadoop.*<System domain name>*		The primary group is **kafkaadmin**.	
	storm		The primary group is **stormadmin**.	
	storm/hadoop.*<System domain name>*		The primary group is **stormadmin**.	
	storm_zk		The primary group is **storm**.	
	storm_zk/hadoop.*<System domain name>*		The primary group is **storm**.	
	kms/hadoop		The primary group is **kmsadmin**.	
	knox		The primary group is **compcommon**.	
	executor		The primary group is **compcommon**.	
	rangeradmin		The primary group is **supergroup**.	
	rangeradmin/hadoop.*<System domain name>*		The primary group is **supergroup**.	
	rangerusersync		The primary group is **supergroup**.	
	rangertagsync		The primary group is **supergroup**.	
	rangerauditor		The primary group is **compcommon**.	

NOTE

Log in to FusionInsight Manager, choose **System** > **Permission** > **Domain and Mutual Trust**, and check the value of **Local Domain**. In the preceding table, all letters in the system domain name contained in the username of the system internal user are lowercase letters.

For example, if **Local Domain** is set to **9427068F-6EFA-4833-B43E-60CB641E5B6C.COM**, the username of default HDFS startup user is **hdfs/hadoop.9427068f-6efa-4833-b43e-60cb641e5b6c.com**.

Database Users

The system database users include OMS database users and DBService database users.

Database Type	Default User	Initial Password	Description	Password Change Method
OMS database	ommdba	• Versions earlier than MRS 3.2.0: **dbChange@123456** • MRS 3.2.0 o	OMS database administrator who performs maintenance operations, such as creating, starting, and stopping.	For details, see **Changing the Password of the OMS Database Administrator**.

Database Type	Default User	Initial Password	Description	Password Change Method
		rlatter:randompassword		

Database Type	Default User	Initial Password	Description	Password Change Method
	omm	• Versions earlier than MRS 3.2.0: **ChangeMe@123456** • MRS 3.2.0 or l	User for accessing OMS database data	For details, see **Changing the Password for the Data Access User of the OMS Database**.

Database Type	Default User	Initial Password	Description	Password Change Method
		ater:randompassword		

Database Type	Default User	Initial Password	Description	Password Change Method
DBService database	omm	• Versions earlier than MRS 3.2.0: **dbserverAdmin@123** • MRS 3.2.0 o	Administrator of the GaussDB database in the DBService component	• For versions earlier than MRS 3.1.2, see **Changing the Password for a Component Database User**. • The initial password cannot be changed in MRS 3.1.2 or later.

Database Type	Default User	Initial Password	Description	Password Change Method
		rlatter:randompassword		
	compdb user	Random password	MRS 3.1.2 or later: Administrator of the GaussDB database in the DBService component. It is used in service O&M scenarios. If the password of this account has expired, you need to reset the password upon your first login.	For details, see **Changing the Password for User compdbuser of the DBService Database**.
	hetu	Random password	User for HetuEngine to connect to the DBService database **hetumeta**. This user exists only in MRS 3.1.2 or later.	• For versions earlier than MRS 3.1.2, see **Changing the Password for a Component Database User**. • For MRS 3.1.2 or later, see **Resetting the Component Database User Password**.

Database Type	Default User	Initial Password	Description	Password Change Method
	hive	• Versions earlier than MRS 3.1.2: HiveUser@ • MRS 3.1.2 or later: ran	User for Hive to connect to the DBService database **hivemeta**.	

Database Type	Default User	Initial Password	Description	Password Change Method
		dompassword		

Database Type	Default User	Initial Password	Description	Password Change Method
	hue	• Versions earlier than MRS 3.1.2: HueUser@123 • MRS 3.1.2 or later:	User for Hue to connect to the DBService database **hue**.	

Database Type	Default User	Initial Password	Description	Password Change Method
		randompassword		

Database Type	Default User	Initial Password	Description	Password Change Method
	sqoop	• Versions earlier than MRS 3.1.2: SqoopUser@ • MRS 3.1.2 or later: r	User for Loader to connect to the DBService database **sqoop**.	

Database Type	Default User	Initial Password	Description	Password Change Method
		randompassword		

Database Type	Default User	Initial Password	Description	Password Change Method
	oozie	• Versions earlier than MRS 3.1.2: OozieUser@ • MRS 3.1.2 or later: ra	User for Oozie to connect to the DBService database **oozie**.	

Database Type	Default User	Initial Password	Description	Password Change Method
		ndompassword		

Database Type	Default User	Initial Password	Description	Password Change Method
	rangeradmin	• Versionsearlierthan MRS 3.1.2: Admin12! • MRS 3.1.2 orlater: ran	User for Ranger to connect to the DBService database.	

Database Type	Default User	Initial Password	Description	Password Change Method
		dompassword		
	kafkaui	Random password	User for Kafka UI to connect to the DBService database. This user exists only in MRS 3.1.2 or later.	
	flink	Random password	User for Flink to connect to the DBService database. This user exists only in MRS 3.1.2 or later.	
	cdl	Random password	User for CDL to connect to the DBService database **cdl**. This user exists only in MRS 3.2.0 or later.	

8.12.1.6 Default Permission Information

Role

Default Role	Description
Manager_administrator	Manager administrator who has all permissions for Manager. Manager administrators can create first-level tenants, create and modify user groups, and specify user permissions.
Manager_operator	Manager operator who has all the permissions on the **Homepage**, **Cluster**, **Hosts**, and **O&M** tab pages.

Default Role	Description
Manager_auditor	Manager auditor who has all permissions on the **Audit** tab page. Manager auditors can view and manage Manager system audit logs.
Manager_viewer	Manager viewer who has the permission to view information about **Homepage**, **Cluster**, **Hosts**, **Alarm**, **Events**, and **System** > **Permission**, and download clients. (Only MRS 3.2.0 or later supports client download.)
Manager_tenant	Manager tenant administrator. This role can create and manage sub-tenants for the non-leaf tenants to which the current user belongs. It has the permission to view alarms and events on **O&M > Alarm**.
System_administrator	System administrator, this role has Manager system administrator rights and all services administrator rights.
default	This role is the default role created for the **default** tenant. It has the management permissions on the Yarn component and the default queue. The default role of the default tenant that is not the first cluster to be installed is **c***<cluster ID>*_**default**.
Manager_administrator_180	FusionInsight Manager System administrator group. Internal system user group, which is used only between components.
Manager_auditor_181	FusionInsight Manager system auditor group. Internal system user group, which is used only between components.
Manager_operator_182	FusionInsight Manager system operator group. Internal system user group, which is used only between components.
Manager_viewer_183	FusionInsight Manager system viewer group. Internal system user group, which is used only between components.
System_administrator_186	System administrator group. Internal system user group, which is used only between components.
Manager_tenant_187	Tenant system user group. Internal system user group, which is used only between components.
default_1000	This group is created for tenant. Internal system user group, which is used only between components.

User group

Type	Default User Group	Description
OS User Group	hadoop	Users added to this group are granted the permission to submit all Yarn queue tasks.
	hadoopmanager	Users added to this user group can have the O&M manager rights of HDFS and Yarn. The O&M manager of HDFS can access the NameNode WebUI and perform active to standby switchover manually. The O&M manager of Yarn can access the ResourceManager WebUI, operate NodeManager nodes, refresh queues, and set node labels, but cannot submit tasks.
	hetuadmin	HetuEngine administrator group. Users in this group have the permission to perform operations on HSConsole.
	hive	Common user group. Hive users must belong to this user group.
	iotdbgroup	Users added to this user group have the administrator rights of the IoTDB component.
	kafka	Kafka common user group. A user in this group can access a topic only when a user in the kafkaadmin group grants the read and write permission of the topic to the user.
	kafkaadmin	Kafka administrator group. Users in this group have the rights to create, delete, authorize, read, and write all topics.
	kafkasuperuser	Topic read/write user group of Kafka. Users added to this group have the read and write permissions on all topics.
	cdladmin	CDL administrator group. Only users in this group can access CDL APIs.
	cdl	Common user group of CDL. Users in this group can create and query CDL jobs.
	storm	Users who are added to the storm user group can submit topologies and manage their own topologies.
	stormadmin	Users who are added to the stormadmin user group can have the storm administrator rights and can submit topologies and manage all topologies.
	supergroup	Users added to this user group have the administrator rights of HBase, HDFS, and Yarn and can use Hive.
	yarnviewgroup	Indicates the read-only user group of the Yarn task. Users in this user group can have the view permission on Yarn and MapReduce tasks.

Type	Default User Group	Description
	check_sec_ldap	Perform internal test on the active LDAP to see whether it works properly. This user group is generated randomly in a test and automatically deleted after the test is complete. Internal system user group, which is used only between components.
	compcommon	System internal group for accessing cluster system resources. All system users and system running users are added to this user group by default.
OS User Group	wheel	Primary group of the FusionInsight internal running user omm.
	ficommon	System common group that corresponds to **compcommon** for accessing cluster common resource files stored in the OS.

NOTE

If the current cluster is not the cluster that is installed for the first time in FusionInsight Manager, the default user group name of all components except Manager in the cluster is **c***<cluster ID>_ default user group name*, for example, **c2_hadoop**.

User

For details, see **User Account List**.

Service-related User Security Parameters

- **HDFS**

  The **dfs.permissions.superusergroup** parameter specifies the administrator group with the highest permission on the HDFS. The default value is **supergroup**.

- **Spark2x**

  The **spark.admin.acls** parameter specifies the administrator list of the Spark2x. Members in the list are authorized to manage all Spark tasks. Users not added in the list cannot manage all Spark tasks. The default value is **admin**.

8.12.1.7 FusionInsight Manager Security Functions

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

- User management: Users can be added, deleted, modified, queried, bound to user groups, and assigned with roles.

  For details, see **Managing Users**.

- User group management: User groups can be added, deleted, modified, queried, and bound to roles.

  For details, see **Managing User Groups**.
- Role management: Roles can be added, deleted, modified, queried, and assigned with the resource access rights of one or multiple components.

  For details, see **Managing Roles**.
- Tenant management: Tenants can be added, deleted, modified, queried, and bound to component resources. FusionInsight generates a role for each tenant to facilitate management. If a tenant is assigned with the rights of some resources, its corresponding role also has these rights.

  For details, see **Tenant Resources**.

8.12.2 Account Management

8.12.2.1 Account Security Settings

8.12.2.1.1 Unlocking LDAP Users and Management Accounts

Scenario

If the LDAP user **cn=pg_search_dn,ou=Users,dc=hadoop,dc=com** and LDAP management accounts **cn=krbkdc,ou=Users,dc=hadoop,dc=com** and **cn=krbadmin,ou=Users,dc=hadoop,dc=com** are locked, the administrator must unlock these accounts.

NOTE

If you input an incorrect password for the LDAP user or management account for five consecutive times, the LDAP user or management account is locked. The account is automatically unlocked after 5 minutes.

Procedure

Step 1 Log in to the active management node as user **omm**.

Step 2 Run the following command to go to the related directory:

cd ${BIGDATA_HOME}/om-server/om/ldapserver/ldapserver/local/script

Step 3 Run the following command to unlock the LDAP user or management account:

./ldapserver_unlockUsers.sh *USER_NAME*

In the command, *USER_NAME* indicates the name of the user to be unlocked.

For example, to unlock the LDAP management **account cn=krbkdc,ou=Users,dc=hadoop,dc=com**, run the following command:

./ldapserver_unlockUsers.sh krbkdc

After the script is executed, enter the password of user **krbkdc** after **ROOT_DN_PASSWORD**. If the following information is displayed, the account is successfully unlocked.


```
Unlock user krbkdc successfully.
```

----End

8.12.2.1.2 Internal an Internal System User

Scenario

If the service is abnormal, the internal user of the system may be locked. Unlock the user promptly, or the cluster cannot run properly. For the list of system internal users, see **User Account List** in . The internal user of the system cannot be unlocked using FusionInsight Manager.

Prerequisites

Obtain the default password of the LDAP administrator **cn=root,dc=hadoop,dc=com** by referring to **User Account List** in .

Procedure

Step 1 Use the following method to confirm whether the internal system username is locked:

1. OLdap port number obtaining method:
   a. Log in to FusionInsight Manager, choose **System** > **OMS** > **oldap** > **Modify Configuration**.
   b. The **LDAP Listening Port** parameter value is **oldap port**.
2. Domain name obtaining method:
   a. Log in to FusionInsight Manager, choose **System** > **Permission** > **Domain and Mutual Trust**.
   b. The **Local Domain** parameter value is the domain name.
      For example, the domain name of the current system is **9427068F-6EFA-4833-B43E-60CB641E5B6C.COM**.
3. Run the following command on each node in the cluster as user **omm** to query the number of password authentication failures:

   ldapsearch -H ldaps://*OMS Floating IP Address***:***OLdap port* **-LLL -x -D cn=root,dc=hadoop,dc=com -b krbPrincipalName=***Internal system username***@***Domain name***,cn=***Domain name***,cn=krbcontainer,dc=hadoop,dc=com -w** *Password of LDAP administrator* **-e ppolicy | grep krbLoginFailedCount**

   NOTE

   To obtain the floating IP address of OMS, log in to the Master2 node remotely, and run the **ifconfig** command. In the command output, **eth0:wsom** indicates the floating IP address of OMS. Record the value of **inet**. If the floating IP address of OMS cannot be queried on the Master2 node, switch to the Master1 node to query and record the floating IP address. If there is only one Master node, query and record the IP address on the Master node.

   For example, run the following command to check the number of password authentication failures for user **oms/manager**:

   ldapsearch -H ldaps://10.5.146.118:21750 -LLL -x -D cn=root,dc=hadoop,dc=com -b krbPrincipalName=oms/

manager@9427068F-6EFA-4833-B43E-60CB641E5B6C.COM,cn=9427068F-6EFA-4833-B43E-60CB641E5B6C.COM,cn=krbcontainer,dc=hadoop,dc=com -w *Password of user cn=root,dc=hadoop,dc=com* **-e ppolicy | grep krbLoginFailedCount**

```
krbLoginFailedCount: 5
```

4. Log in to FusionInsight Manager, choose **System** > **Permission** > **Security Policy** > **Password Policy**.
5. Check the value of the **Password Retries** parameter. If the value is less than or equal to the value of **krbLoginFailedCount**, the user is locked.

   NOTE

   You can also check whether internal users are locked by viewing operations logs.

Step 2 Log in to the active management node as user **omm** and run the following command to unlock the user:

sh ${BIGDATA_HOME}/om-server/om/share/om/acs/config/unlockuser.sh --userName *Internal system username*

Example: **sh ${BIGDATA_HOME}/om-server/om/share/om/acs/config/unlockuser.sh --userName oms/manager**

----End

8.12.2.1.3 Enabling and Disabling Permission Verification on Cluster Components

Scenario

HDFS and ZooKeeper verify the permission of users who attempt to access the services in both security and normal clusters by default. Users without related permission cannot access resources in HDFS and ZooKeeper. When the cluster is deployed in normal mode, HBase and YARN do not verify the permission of users who attempt to access the services by default. All users can access resources in HBase and YARN.

Based on actual service requirements, administrators can enable permission verification on HBase and YARN or disable permission verification on HDFS and ZooKeeper in normal clusters.

Impact on the System

After the enabling and disabling operations, the service configuration will expire. You need to restart the corresponding service for the configuration to take effect.

Enabling Permission Verification on HBase

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Cluster**, click the name of the desired cluster, choose **Services** > **HBase**, and click **Configurations**.

Step 3 Click **All Configurations**.

Step 4 Search for parameters **hbase.coprocessor.region.classes**, **hbase.coprocessor.master.classes**, and **hbase.coprocessor.regionserver.classes**.

Add the coprocessor parameter **org.apache.hadoop.hbase.security.access.AccessController** to the end of the values of the preceding parameters, and use a comma (,) to separate the values from those of the original coprocessors. The **hbase.coprocessor.region.classes** parameter is used as an example.

Figure 8-66 hbase.coprocessor.region.classes

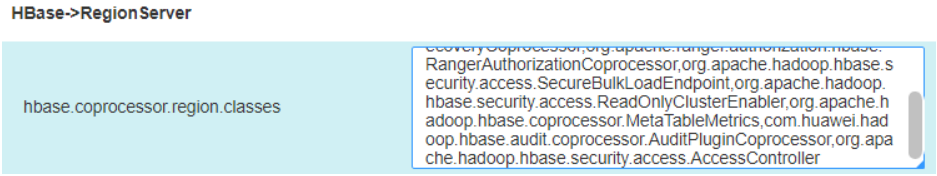

Step 5 Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

Disabling Permission Verification on HBase

NOTE

After HBase permission verification is disabled, the existing permission data will be retained. If you want to delete permission information, disable permission verification, enter the HBase shell, and delete table **hbase:acl**.

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Cluster**, click the name of the desired cluster, choose **Services** > **HBase**, and click **Configurations**.

Step 3 Click **All Configurations**.

Step 4 Search for parameters **hbase.coprocessor.region.classes**, **hbase.coprocessor.master.classes**, and **hbase.coprocessor.regionserver.classes**.

Delete the coprocessor parameter **org.apache.hadoop.hbase.security.access.AccessController**.

Step 5 Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

Disabling Permission Verification on HDFS

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Cluster**, click the name of the desired cluster, choose **Services** > **HDFS**, and click **Configurations**.

Step 3 Click **All Configurations**.

Step 4 Search for parameters **dfs.namenode.acls.enabled** and **dfs.permissions.enabled**.

- **dfs.namenode.acls.enabled** indicates whether to enable HDFS ACL. The default value is **true**, indicating that the ACL is enabled. Change the value to **false**.
- **dfs.permissions.enabled** indicates whether to enable permission check for HDFS. The default value is **true**, indicating that permission check is enabled. Change the value to **false**. After the modification, the owner, owner group, and permission of the directories and files in HDFS remain unchanged.

Step 5 Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

Enabling Permission Verification on YARN

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Cluster**, click the name of the desired cluster, choose **Services** > **Yarn**, and click **Configurations**.

Step 3 Click **All Configurations**.

Step 4 Search for parameter **yarn.acl.enable**.

yarn.acl.enable indicates whether to enable the permission check for YARN.

- In normal clusters, the value is set to **false** by default to disable permission check. To enable permission check, change the value to **true**.
- In security clusters, the value is set to **true** by default to enable authentication.

Figure 8-67 Setting the yarn.acl.enable parameter

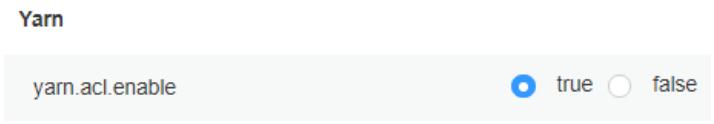

Step 5 Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

Disabling Permission Verification on ZooKeeper

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Cluster**, click the name of the desired cluster, choose **Services** > **ZooKeeper**, and click **Configurations**.

Step 3 Click **All Configurations**.

Step 4 Search for parameter **skipACL**.

skipACL indicates whether to skip the ZooKeeper permission check. The default value is **no**, indicating that permission check is enabled. Change the value to **yes**.

Figure 8-68 Setting the skipACL parameter

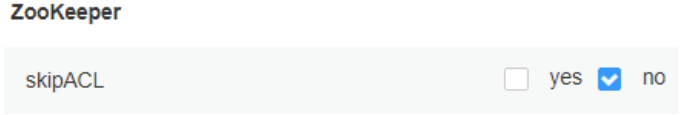

Step 5 Click **Save**, click **OK**, and wait for message "Operation successful" to display.

----End

8.12.2.1.4 Logging In to a Non-Cluster Node Using a Cluster User in Normal Mode

Scenario

When the cluster is installed in normal mode, the component clients do not support security authentication and cannot use the **kinit** command. Therefore, nodes outside the cluster cannot use users in the cluster by default. This may result in a user authentication failure when one of these nodes access a component server.

The node administrator can configure a user who has the same name as that of a user for a node outside the cluster, allow the user to log in to the node using the SSH protocol, and connect to the servers of components in the cluster by using the user who logs in to the OS.

Prerequisites

- Nodes outside the cluster can connect to the service plane of the cluster.
- The KrbServer service of the cluster is running properly.
- You have obtained the password of user **root** of the node outside the cluster.
- A human-machine user has been planned and added to the cluster, and you have obtained the authentication credential file. For details, see **Creating a User** and **Exporting an Authentication Credential File**.

Procedure

Step 1 Log in to the node where a user is to be added as user **root**.

Step 2 Run the following command:

rpm -qa | grep pam and **rpm -qa| grep krb5-client**

The following RPM packages are displayed:

```
pam_krb5-32bit-2.3.1-47.12.1
pam-modules-32bit-11-1.22.1
yast2-pam-2.17.3-0.5.211
pam-32bit-1.1.5-0.10.17
pam_mount-32bit-0.47-13.16.1
pam-config-0.79-2.5.58
pam_krb5-2.3.1-47.12.1
pam-doc-1.1.5-0.10.17
pam-modules-11-1.22.1
pam_mount-0.47-13.16.1
pam_ldap-184-147.20
pam-1.1.5-0.10.17
krb5-client-1.6.3
```

Step 3 Check whether the RPM packages in the list are installed in the OS.

- If yes, go to **Step 5**.
- If no, go to **Step 4**.

Step 4 Obtain the lacked RPM packages from the OS image, upload the files to the current directory, and run the following command to install the RPM package:

rpm -ivh *.rpm

NOTE

The RPM packages to be installed may bring security risks. The risks that may be brought by the installation of these RPM packages must be taken into consideration during OS hardening.

After the RPM packages are installed, go to **Step 5**.

Step 5 Run the following command to configure Kerberos authentication on PAM:

pam-config --add --krb5

NOTE

If you need to cancel Kerberos authentication and system user login on a non-cluster node, run the **pam-config --delete --krb5** command as user **root**.

Step 6 Decompress the authentication credential file to obtain **krb5.conf**, use WinSCP to upload this configuration file to the **/etc** directory on the node outside the cluster, and run the following command to configure related permission to enable other users to access the file, such as permission **604**:

chmod 604 /etc/krb5.conf

Step 7 Run the following command in the connection session as user **root** to add the corresponding OS user to the human-machine user, and specify **root** as the primary group.

The OS user password is the same as the initial password when the human-machine user is created on Manager.

useradd *User name* **-m -d /home/admin_test -g root -s /bin/bash**

For example, if the name of the human-machine user is **admin_test**, run the following command:

useradd admin_test -m -d /home/admin_test -g root -s /bin/bash

NOTE

When you use the newly added OS user to log in to the node by using the SSH protocol for the first time, the system prompts that the password has expired after you enter the user password, and the system prompts that the password needs to be changed after you enter the user password again. You need to enter a new password that meets the password complexity requirements of both the node OS and the cluster.

----End

8.12.2.2 Changing the Password for a System User

8.12.2.2.1 Changing the Password for User admin

Scenario

User **admin** is the system administrator account of FusionInsight Manager. You are advised to periodically change the password on FusionInsight Manager to improve system security.

Procedure

Step 1 Log in to FusionInsight Manager.

User **admin** is required for login.

Step 2 Move the cursor to **Hello, admin** in the upper right corner of the page.

In the displayed menu, click **Change Password**.

Step 3 Set **Old Password**, **New Password**, and **Confirm Password**, and click **OK**.

The password must meet the following complexity requirements:

- Contains 8 to 64 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~`!?,.;-_'(){}[]/<>@#$%^&*+|\=).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be a common easily-cracked password.
- Cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in **Configuring Password Policies**.

----End

8.12.2.2.2 Changing the Password for an OS User

Scenario

During FusionInsight Manager installation, the system automatically creates user **omm** and **ommdba** on each node in the cluster. Periodically change the login passwords of the OS users **omm** and **ommdba** of the cluster node to improve the system O&M security.

The passwords of users **omm** and **ommdba** of the nodes can be different.

Prerequisites

- You have obtained the IP address of the node where the passwords of users **omm** and **ommdba** are to be changed.
- You have obtained the password of user **root** before changing the passwords of users **omm** and **ommdba**.

Changing the Password of an OS User

Step 1 Log in to the node where the password is to be changed as user **root**.

Step 2 Run the following command to change the user password:

passwd *ommdba*

The command output in Red Hat is as follows:

```
Changing password for user ommdba.
New password:
```


Step 3 Enter a new password. The policy for changing the password of an OS user varies according to the OS that is actually used.

```
Retype New Password:
Password changed.
```

----End

8.12.2.3 Changing the Password for a System Internal User

8.12.2.3.1 Changing the Password for the Kerberos Administrator

Scenario

It is recommended that the administrator periodically change the password of Kerberos administrator **kadmin** to improve the system O&M security.

If the user password is changed, the OMS Kerberos administrator password is changed as well.

Prerequisites

You have installed the client on any node in the cluster and obtained the IP address of the node.

Procedure

Step 1 Log in to the node where the client is installed as user **root**.

Step 2 Run the following command to go to the client directory, for example, **/opt/hadoopclient**:

cd /opt/hadoopclient

Step 3 Run the following command to set environment variables:

source bigdata_env

Step 4 Run the following command to change the password for **kadmin/admin**. The password change takes effect on all servers. Keep the password secure because it cannot be retrieved once lost.

kpasswd kadmin/admin

Enter the password (default password: **Admin@123**) and set a new password. The new password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~`!?,.;-_'(){}[]/<>@#$%^&*+|\=).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be a common easily-cracked password, for example, **Admin@12345**.
- Cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in **Configuring Password Policies**.

----End

8.12.2.3.2 Changing the Password for the OMS Kerberos Administrator

Scenario

It is recommended that the administrator periodically change the password of OMS Kerberos administrator **kadmin** to improve the system O&M security.

If the user password is changed, the Kerberos administrator password is changed as well.

Procedure

Step 1 Log in to any management node in the cluster as user **omm**.

Step 2 Run the following command to go to the related directory:

cd ${BIGDATA_HOME}/om-server/om/meta-0.0.1-SNAPSHOT/kerberos/scripts

Step 3 Run the following command to set environment variables:

source component_env

Step 4 Run the following command to change the password for **kadmin/admin**. This operation takes effect for all servers. Keep the password secure because it cannot be retrieved once lost.

kpasswd kadmin/admin

Enter the password (default password: **Admin@123**) and set a new password. The new password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!?,.;-_'(){}[]/<>@#$%^&*+|\=).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be a common easily-cracked password, for example, **Admin@12345**.
- Cannot be the same as the password used in the last *N* times. *N* indicates the value of **Repetition Rule** in **Configuring Password Policies**.

----End

8.12.2.3.3 Changing the Passwords of the LDAP Administrator and the LDAP User (Including OMS LDAP)

NOTE

This section applies only to MRS 3.1.0. For later versions, see **Modifying OMS Service Configuration Parameters**.

Scenario

It is recommended that the administrator periodically changes the passwords of LDAP administrator **cn=root,dc=hadoop,dc=com** and LDAP user **cn=pg_search_dn,ou=Users,dc=hadoop,dc=com** to improve the system O&M security.

If the passwords are changed, the password of the OMS LDAP administrator or user is changed as well.

NOTE

If the cluster is upgraded from an early version to a latest version, the LDAP administrator password will inherit the password policy of the old cluster. To ensure system security, you are advised to change the password after the cluster upgrade.

Impact on the System

- Changing the user password of the LdapServer service is a high-risk operation and requires restarting the KrbServer and LdapServer services. If KrbServer is restarted, users may fail to be queried by running the **id** command on nodes in the cluster temporarily. Therefore, exercise caution when restarting KrbServer.
- After the password of LDAP user **cn=pg_search_dn,ou=Users,dc=hadoop,dc=com** is changed, the user may be locked in the LDAP component. Therefore, you are advised to unlock the user after changing the password. For details about how to unlock the user, see **Unlocking LDAP Users and Management Accounts**.

Prerequisites

Before changing the password of LDAP user **cn=pg_search_dn,ou=Users,dc=hadoop,dc=com**, ensure that the user is not locked by running the following command on the active management node of the cluster:

NOTE

To query the OLdap port number, perform the following steps:

1. Log in to FusionInsight Manager, choose **System** > **OMS** > **oldap** > **Modify Configuration**:
2. The value of **LDAP Service Listening Port** is the OLDAP port.

ldapsearch -H ldaps://*Floating IP address of OMS:OLDAP port***-LLL -x -D cn=pg_search_dn,ou=Users,dc=hadoop,dc=com -W -b cn=pg_search_dn,ou=Users,dc=hadoop,dc=com -e ppolicy**

Enter the password of the LDAP user **pg_search_dn**. If the following information is displayed, the user is locked. In this case, unlock the user. For details, see **Unlocking LDAP Users and Management Accounts**.

NOTE

The password of the LDAP user **pg_search_dn** is randomly generated by the system. You can obtain the password from the **/etc/sssd/sssd.conf or /etc/ldap.conf** file on the active node.

```
ldap_bind: Invalid credentials (49); Account locked
```

Procedure

Step 1 Log in to FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Service** > **LdapServer**.

Step 2 Choose **More** > **Change Database Password**. In the displayed dialog box, enter the password of the current login user and click **OK**.

Step 3 In the **Change Password** dialog box, select the user whose password to be modified in the **User Information** drop-down box.

Step 4 Enter the old password in the **Old Password** text box, and enter the new password in the **New Password** and **Confirm Password** text boxes.

The password must meet the following complexity requirements:

- Contains 16 to 32 characters.
- Contains at least three types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~!@#$%^&*()-_=+|[{}];,<.>/?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be the same as the current password.

Step 5 Select **I have read the information and understood the impact** and click **OK** to confirm the modification and restart the service.

----End

8.12.2.3.4 Changing the Password for the LDAP Administrator

NOTE

This section applies only to MRS 3.1.0. For later versions, see **Modifying OMS Service Configuration Parameters**.

Scenario

It is recommended that the administrator periodically changes the passwords of LDAP administrator accounts **cn=krbkdc,ou=Users,dc=hadoop,dc=com** and **cn=krbadmin,ou=Users,dc=hadoop,dc=com** to improve the system O&M security.

Impact on the System

- You need to restart the KrbServer service after changing the password.
- After the password is changed, check whether the LDAP administrator accounts **cn=krbkdc,ou=Users,dc=hadoop,dc=com** and **cn=krbadmin,ou=Users,dc=hadoop,dc=com** are locked, run the following command on the active management node of the cluster to check whether **krbkdc** is locked (the method for user **krbadmin** is similar):

  NOTE

  OLdap port number obtaining method:

  1. Log in to FusionInsight Manager, choose **System** > **OMS** > **oldap** > **Modify Configuration**:
  2. The **LDAP Listening Port** parameter value is **oldap port**.

  ldapsearch -H ldaps://*OMS_FLOAT_ IP address:OLdap port* **-LLL -x -D cn=krbkdc,ou=Users,dc=hadoop,dc=com -W -b cn=krbkdc,ou=Users,dc=hadoop,dc=com -e ppolicy**

Enter the password of the LDAP administrator account **krbkdc**. The default password is **LdapChangeMe@123**. If the following message is displayed, the account is locked. For details about how to unlock the account, see **Unlocking LDAP Users and Management Accounts**.

```
ldap_bind: Invalid credentials (49); Account locked
```

Prerequisites

You have obtained the management node IP address.

Procedure

Step 1 Log in to the active management node as user **omm** with the IP address of the active management node.

Step 2 Run the following command to go to the related directory:

cd ${BIGDATA_HOME}/om-server/om/meta-0.0.1-SNAPSHOT/kerberos/scripts

Step 3 Run the following command to change the password of the LDAP administrator account:

./okerberos_modpwd.sh

Enter the old password and then enter a new password twice.

The password must meet the following complexity requirements:

- Contains 16 to 32 characters.
- Contains at least three types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (`~!@#$%^&*()-_=+|[{}];,<.>/?).
- Cannot be the same as the current password.

If the following information is displayed, the password is changed.

```
Modify kerberos server password successfully.
```

Step 4 Log in to FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Services** > **KrbServer**. On the displayed page, choose **More** > **Restart Service**.

Enter the password and do not select **Restart upper-layer services**. Click **OK** to restart the KrbServer service.

----End

8.12.2.3.5 Changing the Password for a Component Running User

Scenario

The administrator is advised to periodically change the password for each component running user to improve the system O&M security.

Component running users can be classified into the following two types depending on whether their initial passwords are randomly generated by the system:

- If the initial password of a component running user is randomly generated by the system, the user is of the machine-machine type.
- If the initial password of a component running user is not randomly generated by the system, the user is of the human-machine type.

Impact on the System

If the initial password is randomly generated by the system, the cluster needs to be restarted for the password changing to take effect. Services are unavailable during the restart.

Prerequisites

You have installed the client on any node in the cluster and obtained the IP address of the node.

Procedure

Step 1 Log in to the node where the client is installed as the client installation user

Step 2 Run the following command to switch to the client directory, for example, **/opt/client**:

cd /opt/client

Step 3 Run the following command to set environment variables:

source bigdata_env

Step 4 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**. The password will expire upon your first login. Change the password as prompted. Keep the password secure because it cannot be retrieved once lost.

Step 5 Run the following command to change the password of an internal component running user.

cpw *Internal system username*

Example: **cpw hdfs**

User **hdfs** is an example. Replace it with the actual username.

The password must meet the following complexity requirements:

- Contains at least 8 characters.
- Contains at least four types of the following: uppercase letters, lowercase letters, numbers, spaces, and special characters (~`!?,.;-_'(){}[]/<>@#$%^&*+|\=).
- Cannot be the same as the username or the username spelled backwards.

- Cannot be a common easily-cracked password, for example, **Admin@12345**.
- Cannot be the same as the password used in latest *N* times. *N* indicates the value of **Number of Historical Passwords** configured in **Configuring Password Policies**. This policy applies to only human-machine accounts.

NOTE

Run the following command to check user information:

getprinc *Internal system username*

Example: **getprinc hdfs**

Step 6 Determine the type of the user whose password needs to be changed.

- If the user is a machine-machine user, go to **Step 7**.
- If the user is a human-machine user, the password is changed successfully and no further action is required.

Step 7 Log in to FusionInsight Manager.

Step 8 On the home page, click ··· or **More** and click **Restart**.

Step 9 In the displayed window, enter the password of the current login user and click **OK**.

Step 10 In the displayed restart confirmation dialog box, click **OK**.

Step 11 Wait for message "Operation successful" to display.

----End

8.12.2.4 Changing the Password for a Database User

8.12.2.4.1 Changing the Password of the OMS Database Administrator

Scenario

It is recommended that the administrator periodically change the password of the OMS database administrator to improve the system O&M security.

Procedure

Step 1 Log in to the active management node as user **root**.

NOTE

The password of user **ommdba** cannot be changed on the standby management node. Otherwise, the cluster may not work properly. Change the password on the active management node only.

Step 2 Run the following command to switch to another user:

su - omm

Step 3 Run the following command to go to the related directory:

cd $OMS_RUN_PATH/tools

Step 4 Run the following command to change the password for user **ommdba**:

mod_db_passwd ommdba

Step 5 Enter the old password of user **ommdba** and enter a new password twice.

The password must meet the following complexity requirements:

- Contains 16 to 32 characters.
- Contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#$%^&*()-+_=\|[{}];",<.>/?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be the same as the last 20 historical passwords.

If the following information is displayed, the password is changed.

```
Congratulations, update [ommdba] password successfully.
```

----End

8.12.2.4.2 Changing the Password for the Data Access User of the OMS Database

Scenario

It is recommended that the administrator periodically change the password of the user accessing the OMS database to improve the system O&M security.

Impact on the System

The OMS service needs to be restarted for the new password to take effect. The service is unavailable during the restart.

Procedure

Step 1 On FusionInsight Manager, choose **System** > **OMS** > **gaussDB** > **Change Password**.

Step 2 Locate the row where user **omm** is located and click **Change Password** in the **Operation** column.

Step 3 In the displayed window, enter the password of the current login user and click **OK**.

Step 4 Enter the old and new passwords as prompted.

The password must meet the following complexity requirements:

- Contains 8 to 32 characters.
- Contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#$%^&*()-+_=\|[{}];",<.>/?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be the same as the last 20 historical passwords.

Step 5 Click **OK**. Wait until the system displays a message indicating that the operation is successful.

Step 6 Locate the row where user **omm** is located and click **Restart OMS Service** in the **Operation** column.

Step 7 In the displayed window, enter the password of the current login user and click **OK**.

Step 8 In the displayed restart confirmation dialog box, click **OK** to restart the OMS service.

----End

8.12.2.4.3 Changing the Password for a Component Database User

Scenario

It is recommended that the administrator periodically change the password for each component database user to improve the system O&M security.

NOTE

This section applies only to MRS 3.1.0. For versions later than MRS 3.1.0, see **Resetting the Component Database User Password**.

Impact on the System

The services need to be restarted for the new password to take effect. The services are unavailable during the restart.

Procedure

Step 1 On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and click **Services**.

Step 2 Click the name of the service whose database user password is to be reset. On the **Dashboard** page displayed, click **Stop Service**.

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

After confirming the impact of stopping the service, wait until the service is stopped.

Step 3 Click the service whose database user password is to be changed, and choose **More** > **Change Database Password**. On the displayed page, enter the password of the current login user and click **OK**.

Step 4 Enter the old and new passwords as prompted.

The password must meet the following complexity requirements:

- Contains 8 to 32 characters.
- Contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#$%^&*()-+_=\|[{}];",<.>/?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be the same as the last 20 historical passwords.

Step 5 Select **I have read the information and understand the impact** and click **OK**.

Step 6 After the password is changed, choose **More** > **Restart Service**. In the displayed dialog box, enter the password of the current login user, click **OK**, and select **Restart the upper-layer services**. Click **OK** to restart the services.

----End

8.12.2.4.4 Resetting the Component Database User Password

Scenario

Default passwords for components in the MRS cluster to connect to the DBService database are random. You are advised to periodically reset the passwords of component database users to improve system O&M security.

NOTE

This section applies only to MRS 3.1.2 or later. For versions earlier than MRS 3.1.2, see **Changing the Password for a Component Database User**.

Impact on the System

To reset passwords, you need to stop and then restart services, during which services are unavailable.

Procedure

Step 1 On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and click **Services**.

Step 2 Click the name of the service whose database user password is to be reset, for example, **Kafka**, and click **Stop Service** on the **Dashboard** page.

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

After confirming the impact of stopping the service, wait until the service is stopped.

Step 3 On the **Dashboard** page, choose **More** > **Reset Database Password**.

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

Select "I have read the information and understand the impact", and click **OK**.

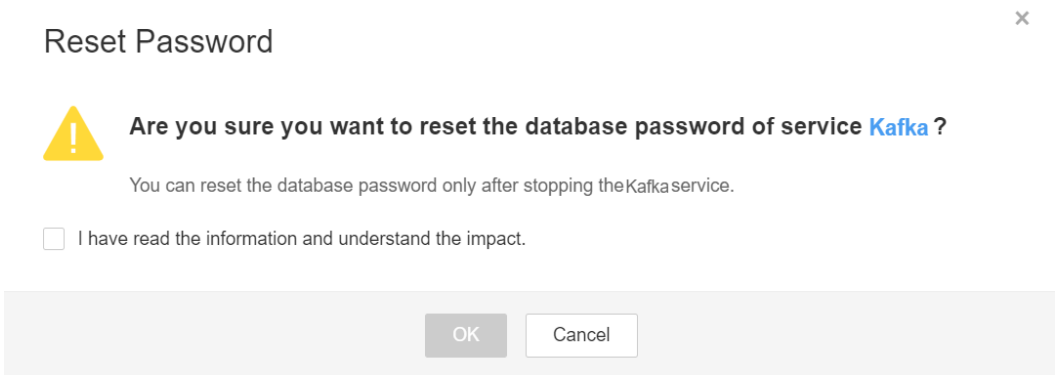

Step 4 After the password is reset, click **Start Service** on the **Dashboard** page.

Step 5 In the displayed dialog box, click **OK** and wait until the service is started.

----End

8.12.2.4.5 Changing the Password for User omm in DBService

Step 1 Log in to the active DBService node as user **root**.

NOTE

The password of user **omm** for the DBService database cannot be changed on the standby DBService node. Change the password on the active DBService node only.

Step 2 Run the following command to switch to another user:

su - omm

Step 3 Run the following command to go to the related directory:

source $DBSERVER_HOME/.dbservice_profile

cd ${DBSERVICE_SOFTWARE_DIR}/sbin/

Step 4 Run the following command to change the password of user **omm**:

sh modifyDBPwd.sh

Step 5 Enter the old password of user **omm** and enter a new password twice.

The password must meet the following complexity requirements:

- Contains 8 to 32 characters.
- Contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#$%^&*()-+_=\|[{}];'',<.>/?).
- Cannot be the same as the username or the username spelled backwards.
- Cannot be the same as the last 20 historical passwords.

If the following information is displayed, the password is changed.

```
Successful to modify password.
```

----End

8.12.2.4.6 Changing the Password for User compdbuser of the DBService Database

Scenario

It is recommended that the administrator periodically change the password of the OMS database administrator to improve the system O&M security.

Procedure

Step 1 Log in to FusionInsight Manager, choose **Cluster** > **Services** > **DBService**, click **Instance**, and view the IP address of the active DBService node.

Step 2 Log in to the active DBService node as user **root**.

NOTE

The password of user **compuserdb** cannot be changed on the standby DBService node. Change the password on the active management node only.

Step 3 Switch to the **$DBSERVER_HOME** directory and configure environment variables:

su - omm

cd $DBSERVER_HOME

source .dbservice_profile

Step 4 Run the following command to change the password of user **compdbuser** as user **omm** of the DBService database:

gsql -U omm -W omm*Password of user omm of the DBService database* **-d postgres -p 20051 -c "alter user compdbuser identified by '***New password***' valid until '***Expiration time***';"**

NOTE

- The new password must meet the following complexity requirements:
  - Contains 16 to 32 characters.
  - Contains at least three types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!@#$%^&*()-+_=\|[{}];:",<.>/?).
  - Cannot be the same as the username or the username spelled backwards.
  - Cannot be the same as the last 20 historical passwords.
- The expiration time format is xxxx-xx-xx, for example, **2020-10-31**.

If the following information is displayed, the modification is successful:

```
ALTER ROLE
```

----End

8.12.2.5 Changing or Resetting the Password for User admin of Manager

User **admin** is the system administrator account of Manager. You are advised to periodically change the password on Manager to improve system security.

If the password is lost, reset the password by referring to **Resetting the Password for User admin**.

Changing the Password for User admin

You can change the password of user **admin** on Manager only for clusters with Kerberos authentication enabled and clusters with Kerberos authentication disabled but the EIP function enabled.

Step 1 Log in to FusionInsight Manager.

User **admin** is required for login.

Step 2 Move your cursor over **Hello, admin** in the upper right corner of the page.

On the user account drop-down menu, choose **Change Password**.

Step 3 Set **Old Password**, **New Password**, and **Confirm Password**, and click **OK**.

----End

Resetting the Password for User admin

Step 1 Log in to the **Master1** node.

Step 2 (Optional) To change the password as user **omm**, run the following command to switch the user:

sudo su - omm

Step 3 Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

cd /opt/Bigdata/client

Step 4 Run the following command to configure environment variables:

source bigdata_env

Step 5 Run the following command to log in to the console as user **kadmin/admin**:

kadmin -p kadmin/admin

NOTE

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

Step 6 Run the following command to reset the password of user **admin**:

cpw admin

----End

8.12.3 Certificate Management

8.12.3.1 Replacing the CA Certificate

Scenario

The MRS CA certificate is used for data encryption during the communication between the client and the server of a component to ensure communication security. You can replace the CA certificate on FusionInsight Manager to ensure product security. This operation is applicable to the following scenarios:

- After the cluster is installed for the first time, import an enterprise certificate.
- If the enterprise certificate has expired or security hardening is required, replace it with a new certificate.

After the CA certificate is replaced, the certificates that are used by HDFS, YARN, MapReduce, HBase, Loader, Hue, Flink (MRS 3.2.0 or later)Oozie, Hive, Tomcat, CAS, HTTPD, and LDAP in MRS will be automatically updated.

The certificate file and key file can be applied for from the enterprise certificate center or generated by the cluster user.

NOTE

- Only CA certificates that can be issued and in **X.509** format can be imported in FusionInsight.
- FusionInsight requires that the OS encoding format be **en_US.UTF-8** or **POSIX**. Otherwise, the certificate function will be abnormal.
- If an isolated faulty node exists in the current cluster, the CA certificate of the node will not be replaced. After the node is de-isolated, you need to reinstall the services running on the node to ensure that the node and the cluster use the same CA certificate.

Impact on the System

The MRS system must be restarted during the replacement and cannot be accessed or provide services.

Prerequisites

- You have obtained the files to be imported to the MRS cluster, including the CA certificate file (*.crt), key file (*.key), and file (password.property) that saves the key file password. The certificate name and key name support letters and digits.
- You have prepared a password for accessing the key file, for example, **Userpwd@123**.

  To avoid potential security risks, the password must meet the following complexity requirements:

  – Contains at least 8 characters.

  – Contains at least four types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!?,.;-_'(){}[]/<>@#$%^&*+|\=).
- When applying for a certificate from the certificate center, provide the password for accessing the key file and apply for the certificate files in CRT, CER, CERT, and PEM formats and the key files in KEY and PEM formats. The applied certificate must have the issuing function.

Procedure

Step 1 Log in to any management node in the cluster as user **omm**.

Step 2 Select a method for generating certificate files and key files.

- If the certificate is generated by the certificate center, save the certificate file and key file to the **omm** user directory on the management node.

  NOTE

  If the obtained certificate file is not in the **.crt** format and the key file is not in the **.key** format, run the following commands to change the file formats:

  mv *Certificate name.Certificate formatCertificate name***.crt**

  mv *Key name.Key format Key name***.key**

  For example, run the following commands to name the certificate file **ca.crt** and name the key file **ca.key**:

  mv server.cer ca.crt

  mv server_key.pem ca.key

- If the certificate is generated by the cluster user, run the following commands to generate the certificate file and key file in the **omm** user directory on the management node:

  a. Generate the key file.

     Run the following command to check whether the OpenSSL version is 1.1.1 or later:

     /usr/bin/openssl version

     - If yes, run the following command:

       openssl genrsa -out *Key name*.**key -aes256 3072**

     - If no, run the following command:

       openssl genrsa -out *Key name*.**key -aes256 3072 -sha256**

     For example, to generate the key file **ca.key**, run the following command:

     openssl genrsa -out ca.key -aes256 3072 -sha256

     Enter the password twice as prompted, and press **Enter**.

     ```
     Enter pass phrase for ca.key:
     Verifying - Enter pass phrase for ca.key:
     ```

  b. Generate the certificate file.

     openssl req -new -x509 -days 1825 -key *Key name*.**key -out** *Certificate name*.**crt -subj "/C=cn/ST=guangdong/L=shenzhen/O=huawei/OU=huawei/CN=huawei" -sha256**

     For example, to generate the certificate file **ca.crt**, run the following command:

     openssl req -new -x509 -days 1825 -key ca.key -out ca.crt -subj "/C=cn/ST=guangdong/L=shenzhen/O=huawei/OU=huawei/CN=huawei" -sha256

     Enter the password for the key file as prompted, and press **Enter**.

     ```
     Enter pass phrase for ca.key:
     ```

Step 3 Run the following command in the **omm** user directory on the management node to save the password for accessing the key file.

sh ${BIGDATA_HOME}/om-server/om/sbin/genPwFile.sh

Enter the password twice as prompted, and press **Enter**. After being encrypted, the password is saved in **password.property**.

```
Please input key password:
Please Confirm password:
```

NOTE

- The **password.property** file generated on the node you have logged is available only for the current cluster and cannot be used for other clusters. The file contains security information. Keep it secure and control the access permission.
- In active/standby DR scenarios, the **genPwFile.sh** script must be executed on both the active and DR cluster nodes, and the same password must be entered for the two clusters.

Step 4 Compress the three files in the **.tar** format and save them to the local computer.

tar -cvf *Package name Certificate name* **.crt** *Key name* **.key password.property**

For example, **tar -cvf test.tar ca.crt ca.key password.property**

NOTE

In active/standby DR scenarios, run this command on each cluster node.

Step 5 Log in to FusionInsight Manager and choose **System** > **Certificate**.

Step 6 In the **Upload Certificate** area, click the file selection button. In the window for selecting files, select the obtained **.tar** certificate file packages and open them and click **Upload**. The system automatically imports the certificate.

Step 7 After the certificate is imported, the system prompts you to synchronize the cluster configuration and restart the web service for the new certificate to take effect. After you complete these operations, click **OK**.

Step 8 In the dialog box that is displayed, enter the password and click **OK** to automatically synchronize the cluster configuration and restart the web service.

Step 9 After the cluster is restarted, enter the URL for accessing FusionInsight Manager in the address box of the browser and check whether the FusionInsight Manager web UI can be successfully displayed.

NOTE

The enterprise certificate has expired or security is hardened. After replacing the MRS certificate, replace the local certificate as well.

Step 10 Click the wanted cluster from the **Cluster** drop-down list.

Step 11 Choose **More** > **Restart**. In the displayed dialog box, enter the password of the current login user and click **OK**.

NOTE

After the CA certificate is replaced, you need to restart the cluster offline to make the certificate take effect. Rolling restart is not supported.

Step 12 In the displayed restart confirmation dialog box, click **OK**.

----End

8.12.3.2 Replacing HA Certificates

Scenario

HA certificates are used to encrypt the communication between active/standby processes and high availability processes to ensure security. Replace the HA certificates on active and standby management nodes on FusionInsight Manager to ensure product security. This operation is applicable to the following scenarios:

- After the cluster is installed for the first time, import an enterprise certificate.
- If the enterprise certificate has expired or security hardening is required, replace it with a new certificate.

NOTE

but is not applicable to scenarios where active and standby management nodes are not installed.

The certificate file and key file can be applied for from the enterprise certificate center or generated by the cluster user.

Impact on the System

FusionInsight Manager must be restarted during the replacement and cannot be accessed or provide services.

Prerequisites

- You have obtained the **root-ca.crt** root file and the **root-ca.pem** key file of the certificate to be replaced.
- You have prepared a password, for example, **Userpwd@123**, for accessing the key file.

  To avoid potential security risks, the password must meet the following complexity requirements:

  – Contains at least 8 characters.

  – Contains at least four types of the following: uppercase letters, lowercase letters, numbers, and special characters (~`!?,.;-_'(){}[]/<>@#$%^&*+|\=).
- When applying for a certificate from the certificate center, provide the password for accessing the key file and apply for the certificate files in CRT, CER, CERT, and PEM formats and the key files in KEY and PEM formats. The applied certificate must have the issuing function.

Procedure

Step 1 Log in to the active management node as user **omm** using the IP address of the active management node.

Step 2 Select a method for generating certificate files and key files.

- If the certificate is generated by the certificate center, save the certificate file and key file to the **${OMS_RUN_PATH}/workspace0/ha/local/cert** directory on the active and standby management nodes.

  NOTE

  If the obtained certificate file is not in the **.crt** format and the key file is not in the **.pem** format, run the following commands to change the file formats:

  mv *Certificate name.Certificate format* **root-ca.crt**

  mv *Key name.Key format* **root-ca.pem**

  For example, run the following commands to name the certificate file **root-ca.crt** and the key file **root-ca.pem**:

  mv server.cer root-ca.crt

  mv server_key.key root-ca.pem
- If the certificate is generated by the cluster user, run the following command to generate **root-ca.crt** and **root-ca.pem** in the **${OMS_RUN_PATH}/workspace0/ha/local/cert** directory:

  sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --root-ca --country=CN --state=*state* **--city=***city* **--company=***company* **--organize=***organize* **--common-name=***commonname* **--email=***Cluster user email address*

NOTE

The validity period of the generated certificate file is 10 years. When the system certificate file is about to expire, the system generates the "ALM-12055 Certificate File Is About to Expire" alarm.

For example, run the following command:

sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --root-ca --country=CN --state=guangdong --city=shenzhen --company=huawei --organize=IT --common-name=HADOOP.COM --email=abc@*xxx*.com

Enter the password as prompted and press **Enter**.

```
Enter pass phrase for /opt/huawei/Bigdata/om-server/OMS/workspace/ha/local/cert/root-ca.pem:
```

The command is executed if the following information is displayed:

```
Generate root-ca pair success.
```

Step 3 On the active management node, run the following command as user **omm** to copy **root-ca.crt** and **root-ca.pem** to the **${BIGDATA_HOME}/om-server/om/security/certHA** directory:

cp -arp ${OMS_RUN_PATH}/workspace0/ha/local/cert/root-ca.* ${BIGDATA_HOME}/om-server/om/security/certHA

Step 4 Copy **root-ca.crt** and **root-ca.pem** generated on the active management node to the **${BIGDATA_HOME}/om-server/om/security/certHA** directory on the standby management node as user **omm**.

scp ${OMS_RUN_PATH}/workspace0/ha/local/cert/root-ca.* omm@***IP address of the standby management node***:${BIGDATA_HOME}/om-server/om/security/certHA

Step 5 Run the following command to generate an HA certificate and perform the automatic replacement:

sh ${BIGDATA_HOME}/om-server/om/sbin/replacehaSSLCert.sh

Enter the password as prompted and press **Enter**.

```
Please input ha ssl cert password:
```

The DBService HA certificate is replaced successfully if the following information is displayed:

```
[INFO] Succeed to replace ha ssl cert.
```

NOTE

If the user wants to update the package for encrypting the HA password, add the **-u** parameter.

Step 6 Run the following command to restart the OMS:

sh ${BIGDATA_HOME}/om-server/om/sbin/restart-oms.sh

The following information is displayed:

```
start HA successfully.
```

Step 7 Log in to the standby management node as user **omm** using the IP address of the standby management node, and repeat steps **Step 5** and **Step 6**.

Run **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** to check whether **HAAllResOK** of the management node is **Normal** and whether FusionInsight Manager can be logged in to again. If yes, the operation is successful.

----End

8.12.4 Security Hardening

8.12.4.1 Hardening Policies

Hardening Tomcat

Tomcat is hardened as follows based on open-source software during FusionInsight Manager software installation and use:

- The Tomcat version is upgraded to the official version.
- Permissions on the directories under applications are set to **500**, and the write permission on some directories is supported.
- The Tomcat installation package is automatically deleted after the system software is installed.
- The automatic deployment function is disabled for projects in application directories. Only the **web**, **cas**, and **client** projects are deployed.
- Some unused **http** methods are disabled, preventing attacks by using the **http** methods.
- The default shutdown port and command of the Tomcat server are changed to prevent hackers from shutting down the server and attacking servers and applications.
- To ensure security, the value of **maxHttpHeaderSize** is changed, which enables server administrators to control abnormal requests of clients.
- The Tomcat version description file is modified after Tomcat is installed.
- To prevent disclosure of Tomcat information, the Server attributes of Connector are modified so that attackers cannot obtain information about the server.
- Permissions on files and directories of Tomcat, such as the configuration files, executable files, log directories, and temporary folders, are under control.
- Session facade recycling is disabled to prevent request leakage.
- LegacyCookieProcessor is used as CookieProcessor to prevent the leakage of sensitive data in cookies.

Hardening LDAP

LDAP is hardened as follows after a cluster is installed:

- In the LDAP configuration file, the password of the administrator account is encrypted using SHA. After the OpenLDAP is upgraded to 2.4.39 or later, data is automatically synchronized between the active and standby LDAP nodes using the SASL External mechanism, which prevents disclosure of the password.

- The LDAP service in the cluster supports the SSLv3 protocol by default, which can be used safely. When the OpenLDAP is upgraded to 2.4.39 or later, the LDAP automatically uses TLS1.0 or later to prevent unknown security risks.

Hardening JDK

- If the client process uses the AES256 encryption algorithm, JDK security hardening is required. The operations are as follows:

  Obtain the Java Cryptography Extension (JCE) package whose version matches that of JDK. The JCE package contains **local_policy.jar** and **US_export_policy.jar**. Copy the JAR files to the following directory and replace the files in the directory.

  - Linux: *JDK installation directory*/**jre/lib/security**
  - Windows: *JDK installation directory***jre\lib\security**

  NOTE

  Access the Open JDK open-source community to obtain the JCE file.

- If the client process uses the SM4 encryption algorithm, the JAR package needs to be updated.

  Obtain **SMS4JA.jar** in the *client installation directory*/**JDK/jdk/jre/lib/ext/** directory, and copy the JAR package to the following directory:

  - Linux: *JDK installation directory*/**jre/lib/ext/**
  - Windows: *JDK installation directory***jre\lib\ext**

8.12.4.2 Configuring a Trusted IP Address to Access LDAP

Scenario

By default, the LDAP service deployed in the OMS and cluster can be accessed by any IP address. To enable the LDAP service to be accessed by only trusted IP addresses, you can configure the INPUT policy in the iptables filtering list.

Impact on the System

After the configuration, the LDAP service cannot be accessed by IP addresses that are not configured. Before the expansion, the added IP addresses need to be configured as trusted IP addresses.

Prerequisites

- You have collected the management plane IP addresses and service plane IP addresses of all nodes in the cluster and all floating IP addresses.
- You have obtained the **root** user account for all nodes in the cluster.

Procedure

Configuring trusted IP addresses for the LDAP service on the OMS

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **OMS** and choose **oldap** > **Modify Configuration** to view the OMS LDAP port number, that is, the value of **LDAP Listening Port**. The default port number is **21750**.

Step 3 Log in to the active management node as user **root** using the IP address of the active management node.

Step 4 Run the following command to check the INPUT policy in the iptables filtering list:

iptables -L

For example, if no rule is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
```

Step 5 Run the following command to configure all IP addresses used by the cluster as trusted IP addresses. Each IP address needs to be added independently.

iptables -A INPUT -s *Trusted IP address* **-p tcp --dport** *Port number* **-j ACCEPT**

For example, to configure **10.0.0.1** as a trusted IP address and enable it to access port **21750**, you need to run the following command:

iptables -A INPUT -s 10.0.0.1 -p tcp --dport 21750 -j ACCEPT

Step 6 Run the following command to configure all IP addresses as untrusted IP addresses. The trusted IP addresses will not be affected by this rule.

iptables -A INPUT -p tcp --dport *Port number* **-j DROP**

For example, to disable all IP addresses to access port **21750**, run the following command:

iptables -A INPUT -p tcp --dport 21750 -j DROP

Step 7 Run the following command to view the modified INPUT policy in the iptables filtering list:

iptables -L

For example, after a trusted IP address is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
ACCEPT     tcp  --  10.0.0.1            anywhere            tcp dpt:21750
DROP       tcp  --  anywhere            anywhere            tcp dpt:21750
```

Step 8 Run the following command to view the rules and rule numbers in the iptables filtering list:

iptables -L -n --line-number

```
Chain INPUT (policy ACCEPT)
num target     prot opt source               destination
1   DROP       tcp  --  0.0.0.0/0           0.0.0.0/0           tcp dpt:21750
```

Step 9 Run the following command to delete the desired rule from the iptables filtering list based on site requirement:

iptables -D INPUT *Number of the rule to be deleted*

For example, to delete rule 1, run the following command:

iptables -D INPUT 1

Step 10 Log in to the standby management node as user **root** using the standby IP address. Repeat **Step 4** to **Step 9**.

Configuring trusted IP addresses for the LDAP service in the cluster

Step 11 Log in to FusionInsight Manager.

Step 12 Click **Cluster**, click the name of the desired cluster, and choose **Service** > **LdapServer**. On the displayed page, click **Instance** to view the nodes where the LDAP services locate.

Step 13 Go to the **Configurations** page, and view the LDAP port number of the cluster, that is, the value of **LDAP_SERVER_PORT**. The default value is **21780**.

Step 14 Log in to the LDAP node as user **root** using the LDAP service IP address.

Step 15 Run the following command to view the INPUT policy in the iptables filtering list:

iptables -L

For example, if no rule is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
```

Step 16 Run the following command to configure all IP addresses used by the cluster as trusted IP addresses. Each IP address needs to be added independently.

iptables -A INPUT -s *Trusted IP address* **-p tcp --dport** *Port number* **-j ACCEPT**

For example, to configure **10.0.0.1** as a trusted IP address and enable it to access port **21780**, you need to run the following command:

iptables -A INPUT -s 10.0.0.1 -p tcp --dport 21780 -j ACCEPT

Step 17 Run the following command to configure all IP addresses as untrusted IP addresses. The trusted IP addresses will not be affected by this rule.

iptables -A INPUT -p tcp --dport *Port number* **-j DROP**

For example, to disable all IP addresses to access port **21780**, run the following command:

iptables -A INPUT -p tcp --dport 21780 -j DROP

Step 18 Run the following command to view the modified INPUT policy in the iptables filtering list:

iptables -L

For example, after a trusted IP address is configured, the INPUT policy is displayed as follows:

```
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
ACCEPT     tcp  --  10.0.0.1             anywhere            tcp dpt:21780
DROP       tcp  --  anywhere             anywhere            tcp dpt:21780
```

Step 19 Run the following command to view the rules and rule numbers in the iptables filtering list:

iptables -L -n --line-number


```
Chain INPUT (policy ACCEPT)
num target     prot opt source               destination
1   DROP       tcp  --  0.0.0.0/0            0.0.0.0/0           tcp dpt:21780
```

Step 20 Run the following command to delete the desired rule from the iptables filtering list based on site requirement:

iptables -D INPUT *Number of the rule to be deleted*

For example, to delete rule 1, run the following command:

iptables -D INPUT 1

Step 21 Log in to the LDAP node as user **root** using the IP address of another LDAP service, and repeat **Step 15** to **Step 20**.

----End

8.12.4.3 HFile and WAL Encryption

HFile and WAL Encryption

NOTICE

- Setting the HFile and WAL encryption mode to SMS4 or AES has a great impact on the system and will cause data loss in case of any misoperation. Therefore, this operation is not recommended.
- Batch data import using Bulkload does not support data encryption.

HFile and Write ahead log (WAL) in HBase are not encrypted by default. To encrypt them, perform the following operations.

Step 1 On any HBase node, run the following commands to create a key file as user **omm**:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh *<path>/hbase.jks <type> <length> <alias>*

- */<path>/hbase.jks* indicates the path for storing the generated JKS file.
- *<type>* indicates the encryption type, which can be SMS4 or AES.
- *<length>* indicates the key length. SMS4 supports 16-bit and AES supports 128-bit.
- *<alias>* indicate the alias of the key file. When you create the key file for the first time, retain the default value **omm**.

For example, to generate an SMS4 encryption key, run the following command:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks SMS4 16 omm

To generate an AES encryption key, run the following command:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks AES 128 omm

NOTE

- To ensure operations can be successfully performed, the **<path>/hbase.jks** directory needs to be created in advance, and the cluster operation user must have the **rw** permission of this directory.
- After running the command, enter the same *<password>* four times. The password encrypted in **Step 3** is the same as the password in this step.

Step 2 Distribute the generated key files to the same directory on all nodes in the cluster and assign read and write permission to user **omm**.

NOTE

- Administrators need to select a safe procedure to distribute keys based on the enterprise security requirements.
- If the key files of some nodes are lost, repeat the step to copy the key files from other nodes.

Step 3 On FusionInsight Manager, set **hbase.crypto.keyprovider.parameters.encryptedtext** to the encrypted password. Set **hbase.crypto.keyprovider.parameters.uri** to the path and name of the key file.

- The format of **hbase.crypto.keyprovider.parameters.uri** is **jceks://** *<key_Path_Name>*.

  <key_Path_Name> indicates the path of the key file. For example, if the path of the key file is **/home/hbase/conf/hbase.jks**, set this parameter to **jceks:///home/hbase/conf/hbase.jks**.

- The format of **hbase.crypto.keyprovider.parameters.encryptedtext** is *<encrypted_password>*.

  <encrypted_password> indicates the encrypted password generated during the key file creation. The parameter value is displayed in ciphertext. Run the following command as user **omm** to obtain the related encrypted password on the nodes where HBase service is installed:

  sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh

  NOTE

  After running the command, you need to enter **<password>**. The password is the same as that entered in **Step 1**.

Step 4 On FusionInsight Manager, set **hbase.crypto.key.algorithm** to **SMS4** or **AES** to use SMS4 or AES for HFile encryption.

Step 5 On FusionInsight Manager, set **hbase.crypto.wal.algorithm** to **SMS4** or **AES** to use SMS4 or AES for WAL encryption.

Step 6 On FusionInsight Manager, set **hbase.regionserver.wal.encryption** to **true**.

Step 7 Save the settings and restart the HBase service for the settings to take effect.

Step 8 Create an HBase table through CLI or code and configure the encryption mode to enable encryption. **<type>** indicates the encryption type, and **d** indicates the column family.

- When you create an HBase table through CLI, set the encryption mode to SMS4 or AES for the column family.

  create '*<table name>*', {*NAME => 'd'*, ***ENCRYPTION => '****<type>***'}

- When you create an HBase table using code, set the encryption mode to SMS4 or AES by adding the following information to the code:

```
public void testCreateTable()
{
    String tableName = "user";
    Configuration conf = getConfiguration();
    HTableDescriptor htd = new HTableDescriptor(TableName.valueOf(tableName));

    HColumnDescriptor hcd = new HColumnDescriptor("d");
    //Set the encryption mode to SMS4 or AES.
    hcd.setEncryptionType("<type>");
    htd.addFamily(hcd);

    HBaseAdmin admin = null;
    try
    {
        admin = new HBaseAdmin(conf);

        if(!admin.tableExists(tableName))
        {
            admin.createTable(htd);
        }
    }
    catch (IOException e)
    {
        e.printStackTrace();
    }
    finally
    {
        if(admin != null)
        {
            try
            {
                admin.close();
            }
            catch (IOException e)
            {
                e.printStackTrace();
            }
        }
    }
}
```

Step 9 You can check whether the encryption configuration is successful by referring to **Verifying the Encryption Configuration**.

Step 10 If you have configured SMS4 or AES encryption by performing **Step 1** to **Step 7**, but do not set the related encryption parameter when creating the table in **Step 8**, the inserted data is not encrypted.

In this case, you can perform the following steps to encrypt the inserted data:

1. Run the **flush** command for the table to import the data in the memory to the HFile.

   flush *'<table_name>'*

2. Run the following commands to modify the table properties:

   disable *'<table_name>'*

alter *'<table_name>'*,**NAME=>** *'<column_name>'*,**ENCRYPTION =>** *'<type>'*

enable *'<table_name>'*

3. Insert a new data record and flush the table.

   NOTE

   A new data record must be inserted so that the HFile will generate a new HFile and the unencrypted data inserted previously will be rewritten and encrypted.

   put *'<table_name>'*, **'id2','f1:c1','value22222222222222222222222222222222'**

   flush *'<table_name>'*

4. Perform the following step to rewrite the HFile:

   major_compact'*<table_name>*'

   NOTICE

   During this step, the HBase table is disabled and cannot provide services. Exercise caution when you perform this step.

5. You can perform **Step 6** to check whether the encryption configuration is successful.

----End

Verifying the Encryption Configuration

NOTE

This operation can be performed only when test data can be written to an empty table.

Step 1 Log in to the node where the client is installed as the client installation user. Switch to the client installation directory, for example, **/opt/client**.

cd /opt/client

Step 2 Run the following command to set environment variables:

source bigdata_env

Step 3 Run the following command to authenticate the current user if Kerberos authentication has been enabled for the current security cluster. The current user must have the permission to read and write HBase tables and the HDFS operation permission.

kinit *Component service user*

Run the following command to set the Hadoop username if Kerberos authentication is not enabled for the current normal cluster:

export HADOOP_USER_NAME=hbase

Step 4 Run the following command to log in to the HBase client:

hbase shell

Run the following command to insert a new data record and flush the table to generate an HFile:

put *'<table_name>'*, **'id2','***d***:c1','value22222222222222222222222222222222'**

flush *'<table_name>'*

NOTE

- *<table_name>* indicates the table configured with SMS4 or AES encryption. For details about how to configure SMS4 or AES encryption, go to **Step 8**.
- *d* indicates the column family configured with SMS4 or AES encryption. For details about how to configure SMS4 or AES encryption, go to **Step 8**.

Step 5 Press **Ctrl**+**C** to exit the HBase client.

Step 6 Run the following command to view the directory where the HFile file generated in **Step 4** is stored:

hdfs dfs -ls

The file directory format is **/hbase/data/***<namespace_name>***/***<table_name>***/***<region_name>***/***<columnfamily_name>***/***<HFile_name>*.

NOTE

If *<namespace_name>* is not specified during HBase table creation, **default** is used by default.

Example:

```
/hbase/data/default/create_table/dd61b81b1ba1aad6513b9bdcfd8f871c/d/
aa6fe387b27443afaba40f5b584c1fa7
```

Step 7 Run the following command to view the HFile content:

hbase hfile -f *<HFile path>* **-p**

NOTE

<HFile path> indicates the directory where the HFile file is located.

The error message "com. huawei.hadoop.hbase.io.crypto.CryptoRuntimeException" will be displayed in the command output. However, the **HBase shell** can still read the table data, indicating that the encryption configuration is successful.

----End

Modifying a Key File

NOTICE

Modifying a key file has a great impact on the system and will cause data loss in case of any misoperation. Therefore, this operation is not recommended.

During the **HFile and WAL Encryption** operation, the related key file must be generated and its password must be set to ensure system security. After a period of running, you can replace the key file with a new one to encrypt HFile and WAL.

Step 1 Run the following command to generate a new key file as user **omm**:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh *<path>/hbase.jks <type> <length> <alias-new>*

- *<path>/hbase.jks*: indicates the path for storing the generated **hbase.jks** file. The path and file name must be consistent with those of the key file generated in **HFile and WAL Encryption**.
- *<alias-new>*: indicates the alias of the key file. The alias must be different with that of the old key file.
- *<type>*: indicates the encryption type, which can be SMS4 or AES.
- *<length>* indicates the key length. SMS4 supports 16-bit and AES supports 128-bit.

For example, to generate an SMS4 encryption key, run the following command:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks SMS4 16 omm_new

To generate an AES encryption key, run the following command:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks AES 128 omm_new

NOTE

- To ensure operations can be successfully performed, the **<path>/hbase.jks** directory needs to be created in advance, and the cluster operation user must have the **rw** permission of this directory.
- After running the command, you need to enter the same *<password>* for three times. This password is the password of the key file. You can use the password of the old file without any security risk.

Step 2 Distribute the generated key files to the same directory on all nodes in the cluster and assign read and write permission to user **omm**.

NOTE

Administrators need to select a safe procedure to distribute keys based on the enterprise security requirements.

Step 3 On the HBase service configuration page of FusionInsight Manager, add custom configuration items, set **hbase.crypto.master.key.name** to **omm_new**, set **hbase.crypto.master.alternate.key.name** to **omm**, and save the settings.

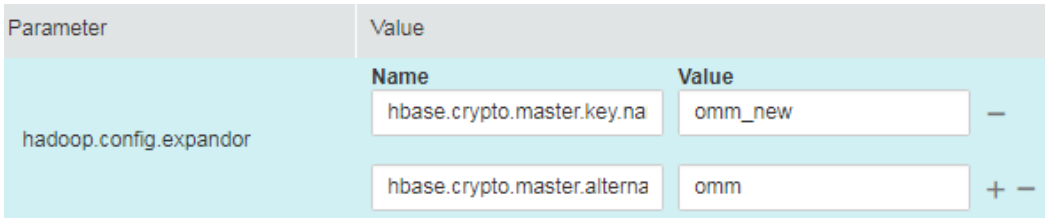

Step 4 Restart the HBase service for the configuration to take effect.

Step 5 In HBase shell, run the **major compact** command to generate the HFile file based on the new encryption algorithm.

major_compact *'<table_name>'*

Step 6 You can view the major compact progress from the HMaster web page.

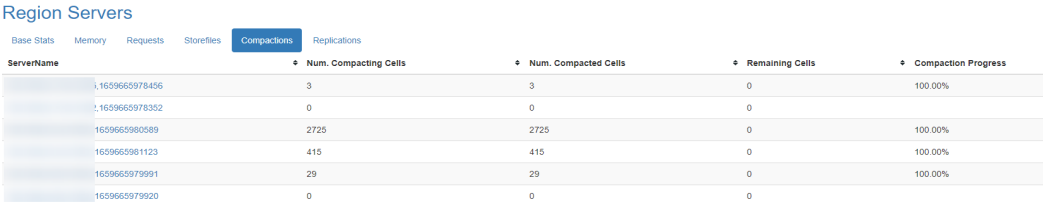

Step 7 When all items in **Compaction Progress** reach **100%** and those in **Remaining KVs** are **0**, run the following command as user **omm** to destroy the old key file:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh *<path>/hbase.jks <alias-old>*

- *<path>/hbase.jks*: indicates the path for storing the generated **hbase.jks** file. The path and file name must be consistent with those of the key file generated in **HFile and WAL Encryption**.
- *<alias-old>*: indicates the alias of the old key file to be deleted.

For example:

sh ${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/install/FusionInsight-HBase-2.2.3/hbase/bin/hbase-encrypt.sh /home/hbase/conf/hbase.jks omm

NOTE

To ensure operations can be successfully performed, the **<path>/hbase.jks** directory needs to be created in advance, and the cluster operation user must have the **rw** permission of this directory.

Step 8 Repeat **Step 2** and distribute the updated key files again.

Step 9 Delete the HBase self-defined configuration item **hbase.crypto.master.alternate.key.name** added in **Step 3** from FusionInsight Manager.

Step 10 Repeat **Step 4** for the configuration take effect.

----End

8.12.4.4 Configuring Hadoop Security Parameters

Configuring Security Channel Encryption

The channels between components are not encrypted by default. You can set the following parameters to configure security channel encryption.

Page access for setting parameters: On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, click **Services**, and click the target service. On the displayed page, click **Configuration** and click **All Configurations**. Enter a parameter name in the search box.

NOTE

Restart corresponding services for the modification to take effect after you modify configuration parameters.

Table 8-87 Parameter description

Service	Parameter	Description	Default Value
HBase	hbase.rpc.protection	Indicates whether the HBase channels, including the remote procedure call (RPC) channels for HBase clients to access the HBase server and the RPC channels between the HMaster and RegionServer, are encrypted. If this parameter is set to **privacy**, the channels are encrypted and the authentication, integrity, and privacy functions are enabled. If this parameter is set to **integrity**, the channels are not encrypted and only the authentication and integrity functions are enabled. If this parameter is set to **authentication**, the channels are not encrypted, only packets are authenticated, and integrity and privacy are not required. **NOTE** The privacy mode encrypts transmitted content, including sensitive information such as user tokens, to ensure the security of the transmitted content. However, this mode has great impact on performance. Compared with the other two modes, this mode reduces read/write performance by about 60%. Modify the configuration based on the enterprise security requirements. The configuration items on the client and server must be the same.	• Security mode: **privacy** • Normal mode: **authentication**
HDFS	dfs.encrypt.data.transfer	Indicates whether the HDFS data transfer channels and the channels for clients to access HDFS are encrypted. The HDFS data transfer channels include the data transfer channels between DataNodes and the Data Transfer (DT) channels for clients to access DataNodes. The value **true** indicates that the channels are encrypted. The channels are not encrypted by default.	false

Service	Parameter	Description	Default Value
HDFS	dfs.encrypt.data.transfer.algorithm	Indicates whether the HDFS data transfer channels and the channels for clients to access HDFS are encrypted. This parameter is valid only when **dfs.encrypt.data.transfer** is set to **true**. The default value is **3des**, indicating that 3DES algorithm is used to encrypt data. The value can also be set to **rc4**. However, to avoid security risks, you are not advised to set the parameter to this value.	3des
HDFS	hadoop.rpc.protection	Indicates whether the RPC channels of each module in Hadoop are encrypted. The channels include: • RPC channels for clients to access HDFS • RPC channels between modules in HDFS, for example, between DataNode and NameNode • RPC channels for clients to access YARN • RPC channels between NodeManager and ResourceManager • RPC channels for Spark to access YARN and HDFS • RPC channels for MapReduce to access YARN and HDFS • RPC channels for HBase to access HDFS The default value is **privacy**, indicating encrypted transmission. The value **authentication** indicates that transmission is not encrypted. **NOTE** You can set this parameter on the HDFS component configuration page. The parameter setting is valid globally, that is, the setting of whether the RPC channel is encrypted takes effect on all modules in Hadoop.	• Security mode: **privacy** • Normal mode: **authentication**

Setting the Maximum Number of Concurrent Web Connections

To ensure web server reliability, new connections are rejected when the number of user connections reaches a specific threshold. This prevents DDOS attacks and

service unavailability caused by too many users accessing the web server at the same time.

Page access for setting parameters: On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, click **Services**, and click the target service. On the displayed page, click **Configuration** and click **All Configurations**. Enter a parameter name in the search box.

Table 8-88 Parameter description

Service	Parameter	Description	Default Value
HDFS/Yarn	hadoop.http.server.MaxRequests	Specifies the maximum number of concurrent web connections of each component.	2000
Spark2x	spark.connection.maxRequest	Specifies the maximum number of request connections of JobHistory.	5000

8.12.4.5 Configuring an IP Address Whitelist for Modification Allowed by HBase

If the Replication function is enabled for HBase clusters, a protection mechanism for data modification is added on the standby HBase cluster to ensure data consistency between the active and standby clusters. Upon receiving an RPC request for data modification, the standby HBase cluster checks the permission of the user who sends the request (only HBase manage users have the modification permission). Then it checks the validity of the source IP address of the request. Only modification requests from IP addresses in the white list are accepted. The IP address white list is configured by the **hbase.replication.allowedIPs** item.

Log in to FusionInsight Manager and choose **Cluster** > **Services** > **HBase**. Click **Configurations** and enter the parameter name in the search box.

Table 8-89 Parameter description

Parameter	Description	Default Value
hbase.replication.allowedIPs	Allows replication request processing from configured IP addresses only. It supports comma separated regex patterns. Each pattern can be any of the following: • Regex pattern Example: 10.18.40.*, 10.18.*, 10.18.40.11 • Range pattern (Range can be specified only in the last octet) Example: 10.18.40.[10-20] If this item is empty (default value), the white list contains only the IP address of the RegionServer of the cluster, indicating that only modification requests from the RegionServer of the standby HBase cluster are accepted.	N/A

8.12.4.6 Updating a Key for a Cluster

Scenario

When a cluster is installed, an encryption key is generated automatically by the system so that the security information in the cluster (such as all database user passwords and key file access passwords) can be stored in encryption mode. After the cluster is installed, if the original key is accidentally disclosed or a new key is required, you can manually update the key.

Impact on the System

- After a cluster key is updated, a new key is generated randomly in the cluster. This key is used to encrypt and decrypt the newly stored data. The old key is not deleted, and it is used to decrypt data encrypted using the old key. After security information is modified, for example, a database user password is changed, the new password is encrypted using the new key.
- When a key is updated for a cluster, the cluster must be stopped and cannot be accessed.

Prerequisites

- You have obtained the IP addresses of the active and standby management nodes.
- You have stopped the upper-layer service applications that depend on the cluster.

Procedure

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **Cluster** > *Name of the desired cluster* and click **Stop**. In the dialog box that is displayed, enter the password of the current user

and click **OK**. Wait for a while until a message indicating that the operation is successful is displayed.

Step 3 Log in to the active management node as user **omm**.

Step 4 Run the following command to disable logout upon timeout:

TMOUT=0

NOTE

After the operations in this section are complete, run the **TMOUT=***Timeout interval* command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.

Step 5 Run the following command to go to the related directory:

cd ${BIGDATA_HOME}/om-server/om/tools

Step 6 Run the following command to update the cluster key:

sh updateRootKey.sh

Enter **y** as prompted.

```
The root key update is a critical operation.
Do you want to continue?(y/n):
```

If the following information is displayed, the key is updated successfully.

```
Step 4-1: The key save path is obtained successfully.
...
Step 4-4: The root key is sent successfully.
```

Step 7 On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and click **Start**.

In the displayed dialog box, click **OK**. Wait until a message is displayed, indicating that the startup is successful.

----End

8.12.4.7 Hardening the LDAP

Configuring the LDAP Firewall Policy

In the cluster adopting the dual-plane networking, the LDAP is deployed on the service plane. To ensure the LDAP data security, you are advised to configure the firewall policy in the cluster to disable relevant LDAP ports.

Step 1 Log in to FusionInsight Manager.

Step 2 Click **Cluster**, click the name of the desired cluster, choose **Services** > **LdapServer**, and click **Configurations**.

Step 3 Check the value of **LDAP_SERVER_PORT**, which is the service port of LdapServer.

Step 4 To ensure data security, configure the firewall policy for the whole cluster to disable the LdapServer port based on the customer's firewall environment.

----End

Enabling the LDAP Audit Log Output

Users can set the audit log output level of the LDAP service and output audit logs in a specified directory, for example, **/var/log/messages**. The logs output can be used to check user activities and operation commands.

NOTE

If the function of LDAP audit log output is enabled, massive logs are generated, affecting the cluster performance. Exercise caution when enabling this function.

Step 1 Log in to any LdapServer node.

Step 2 Run the following command to edit the **slapd.conf.consumer** file, and set the value of **loglevel** to **256** (you can run the **man slapd.conf** command on the OS to view the log level definition).

cd ${BIGDATA_HOME}/FusionInsight_BASE_8.1.0.1/install/FusionInsight-ldapserver-2.7.0/ldapserver/local/template

vi slapd.conf.consumer

```
...
pidfile         [PID_FILE_SLAPD_PID]
argsfile        [PID_FILE_SLAPD_ARGS]
loglevel   256
...
```

Step 3 Log in to FusionInsight Manager, click **Cluster**, click the name of the desired cluster, choose **Services** > **LdapServer**. On the displayed page, choose **More** > **Restart Service**. Enter the administrator password and restart the service.

----End

8.12.4.8 Configuring Kafka Data Encryption During Transmission

Scenario

Data between the Kafka client and the broker is transmitted in plain text. The Kafka client may be deployed in an untrusted network, exposing the transmitting data to leakage and tampering risks.

Procedure

The channel between components is not encrypted by default. You can set the following parameters to enable security channel encryption.

Page access for setting parameters: On FusionInsight Manager, click **Cluster**, click the name of the desired cluster, and choose **Services** > **Kafka**. On the displayed page, click **Configuration** and click **All Configurations**. Enter a parameter name in the search box.

NOTE

After the configuration, restart the corresponding service for the settings to take effect.

Table 8-90 describes the parameters related to transmission encryption on the Kafka server.

Table 8-90 Parameters relevant to Kafka data encryption during transmission

Parameter	Description	Default Value
ssl.mode.enable	Indicates whether to enable the Secure Sockets Layer (SSL) protocol. If this parameter is set to **true**, services relevant to the SSL protocol are started during the broker startup.	false
security.inter.broker.protocol	Indicates communication protocol between brokers. The communication protocol can be PLAINTEXT, SSL, SASL_PLAINTEXT, or SASL_SSL.	SASL_PLAINTEXT

The SSL protocol can be configured for the server or client to encrypt transmission and communication only after **ssl.mode.enable** is set to **true** and broker enables the **SSL** and **SASL_SSL** protocols.

8.12.4.9 Configuring HDFS Data Encryption During Transmission

Configuring HDFS Security Channel Encryption

The channel between components is not encrypted by default. You can set parameters to enable security channel encryption.

Navigation path for setting parameters: On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations**. On the displayed page, click the **All Configurations** tab. Enter a parameter name in the search box.

NOTE

After the configuration, restart the corresponding service for the settings to take effect.

Table 8-91 Parameters

Configuration Item	Description	Default Value
hadoop.rpc.protection	**NOTICE** • The setting takes effect only after the service is restarted. Rolling restart is not supported. • After the setting, you need to download the client configuration file again. Otherwise, HDFS cannot provide the read and write services. • After the setting, you need to restart the executor. Otherwise, the job management and file management functions on the console become unavailable. Indicates whether the RPC channels of each module in Hadoop are encrypted. The channels include: • RPC channels for clients to access HDFS • RPC channels between modules in HDFS, for example, between DataNode and NameNode • RPC channels for clients to access Yarn • RPC channels between NodeManager and ResourceManager • RPC channels for Spark to access Yarn and HDFS • RPC channels for MapReduce to access Yarn and HDFS • RPC channels for HBase to access HDFS **NOTE** The setting takes effect globally, that is, the encryption attribute of the RPC channel of each module in the Hadoop takes effect.	• Security mode: privacy • Normal mode: authentication **NOTE** • **authentication**: indicates that only authentication is required. • **integrity**: indicates that authentication and consistency check need to be performed. • **privacy**: indicates that authentication, consistency check, and encryption need to be performed.

Configuration Item	Description	Default Value
dfs.encrypt.data.transfer	Indicates whether the HDFS data transfer channels and the channels for clients to access HDFS are encrypted. The HDFS data transfer channels include the data transfer channels between DataNodes and the Data Transfer (DT) channels for clients to access DataNodes. The value **true** indicates that the channels are encrypted. The channels are not encrypted by default. **NOTE** • This parameter is valid only when **hadoop.rpc.protection** is set to **privacy**. • If a large amount of service data is transmitted, enabling encryption by default severely affects system performance. • If data transmission encryption is configured for one cluster in the trusted cluster, the same data transmission encryption must be configured for the peer cluster.	false
dfs.encrypt.data.transfer.algorithm	Indicates the algorithm to encrypt the HDFS data transfer channels and the channels for clients to access HDFS. This parameter is valid only when **dfs.encrypt.data.transfer** is set to **true**. **NOTE** The default value is **3des**, indicating that 3DES algorithm is used to encrypt data. The value can also be set to **rc4**. However, to avoid security risks, you are not advised to set the parameter to this value.	3des
dfs.encrypt.data.transfer.cipher.suites	This parameter can be left empty or set to **AES/CTR/NoPadding** to specify the cipher suite for data encryption. If this parameter is not specified, the encryption algorithm specified by **dfs.encrypt.data.transfer.algorithm** is used for data encryption. The default value is **AES/CTR/NoPadding**.	AES/CTR/NoPadding

8.12.4.10 Configuring Spark2x Data Encryption During Transmission

Scenario

This section describes how to configure encryption for Spark2x security channels to enhance security.

Procedure

To modify parameters, log in to FusionInsight Manager, click **Cluster** and choose **Services** > **Spark2x**. On the displayed page, click **Configurations** and click **All Configurations**. Enter a parameter name in the search box.

NOTE

After the configuration, restart the corresponding service for the settings to take effect.

Table 8-92 Parameters

Parameter	Description	Default Value
spark.authenticate	Whether to enable Spark internal security authentication	Security mode: **true** Normal mode: **false**
spark.authenticate.enableSaslEncryption	Whether to enable encrypted communication based on Simple Authentication and Security Layer (SASL)	Security mode: **true** Normal mode: **false**
spark.network.crypto.enabled	Whether to enable RPC encryption based on Advanced Encryption Standard (AES)	Security mode: **true** Normal mode: **false**
spark.network.sasl.serverAlwaysEncrypt	Whether to disable unencrypted connections for ports with SASL authentication enabled	false
spark.network.crypto.keyLength	Length of the encryption key to be generated	256
spark.network.crypto.keyFactoryAlgorithm	Algorithm used to generate the encryption key	PBKDF2WithHmacSHA1

Parameter	Description	Default Value
spark.io.encryption.enabled	Whether to enable local disk I/O encryption	Security mode: **true** Normal mode: **false**
spark.io.encryption.keygen.algorithm	Algorithm used to generate the I/O encryption key	HmacSHA256
spark.io.encryption.keySizeBits	Size of an I/O encryption key, in bits	256
spark.ssl.ui.enabled	Whether to enable Secure Sockets Layer (SSL) authentication for the web UI connection	Security mode: **true** Normal mode: **false**

8.12.4.11 Configuring ZooKeeper SSL

Scenario

By default, SSL channel encryption transmission is disabled between the ZooKeeper client and server and between instances on the server. This section describes how to enable the ZooKeeper channel encryption transmission.

NOTE

This function is available only for MRS clusters of version 3.1.2 or later.

Impact on the System

- When SSL channel encryption transmission is enabled on the ZooKeeper server, the performance deteriorates.
- When SSL channel encryption transmission is enabled on the ZooKeeper server, ZooKeeper and dependent upper-layer components need to be restarted. During the restart, services are unavailable.
- To enable SSL channel encryption transmission on the ZooKeeper server, you need to download the client again.
- If SSL channel encryption transmission is enabled for ZooKeeper, rolling restart is not supported.

Procedure

Step 1 Log in to FusionInsight Manager, click **Cluster** and choose **Services** > **ZooKeeper**. On the displayed page, click **Configurations** and click **All Configurations**.

Step 2 Enter the parameter name in the search box, and change the value as follows:

Table 8-93 Security configuration item

Parameter	Description	Default Value	New Value
ssl.enabled	Whether to enable SSL communication encryption.	false	true

Step 3 After the modification is complete, click **Save** and then click **OK**.

Step 4 Click **Cluster** and choose **Services** > **ZooKeeper**. On the ZooKeeper service page, choose **More** > **Restart Service**, enter the password for authentication, and confirm the operation impact on the **Restart Service** page.

You can select **Restart upper-layer services**. During the restart of all affected components, services will be unavailable. Exercise caution when performing this operation.

Step 5 Click **OK** and wait until the services are restarted successfully.

Step 6 Choose **Cluster** > **Active/Standby Cluster DR** to check whether active/standby DR is configured for the current cluster.

- If yes, go to **Step 7**.
- If no, no further action is required.

Step 7 The **ssl.enabled** configuration of the ZooKeeper service in the active cluster must be the same as that in the DR cluster. Modify the **ssl.enabled** parameter in the cluster where no operation is performed by referring to the preceding steps.

Step 8 Log in to the active OMS node in the active cluster as user **root** and run the following commands to restart the DR management process:

su - omm

${BIGDATA_HOME}/om-server/om/share/om/disaster/sbin/restart-disaster.sh

If the following information is displayed, the operation is successful:

```
...
disaster start with process id : 23256
End into restart-disaster.sh
```

Step 9 Log in to the active OMS node in the DR cluster as user **root** and run the following commands to restart the DR management process:

su - omm

${BIGDATA_HOME}/om-server/om/share/om/disaster/sbin/restart-disaster.sh

----End

8.12.4.12 Encrypting the Communication Between the Controller and the Agent

Scenario

After a cluster is installed, Controller and Agent need to communicate with each other. The Kerberos authentication is used during the communication. By default, the communication is not encrypted during the communication for the sake of cluster performance. Users who have demanding security requirements can use the method described in this section for encryption.

Impact on the System

- Controller and all Agents automatically restart, which interrupts FusionInsight Manager.
- The performance of management nodes deteriorates in large clusters. You are advised to enable the encryption function for clusters with a maximum of 200 nodes.

Prerequisites

You have obtained the IP addresses of the active and standby management nodes.

Procedure

Step 1 Log in to the active management node as user **omm**.

Step 2 Run the following command to disable logout upon timeout:

TMOUT=0

NOTE

After the operations in this section are complete, run the **TMOUT=***Timeout interval* command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.

Step 3 Run the following command to go to the related directory:

cd ${CONTROLLER_HOME}/sbin

Step 4 Run the following command to enable communication encryption:

./enableRPCEncrypt.sh -t

Run the **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** command to check whether **ResHAStatus** of the active management node Controller is **Normal** and whether you can log in to FusionInsight Manager again. If yes, the enablement is successful.

Step 5 Run the following command to disable communication encryption when necessary:

./enableRPCEncrypt.sh -f

Run the **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** command to check whether **ResHAStatus** of the active management node Controller is **Normal**

and whether you can log in to FusionInsight Manager again. If yes, the enablement is successful.

----End

8.12.4.13 Updating SSH Keys for User omm

Scenario

During cluster installation, the system automatically generate the SSH public key and private key for user **omm** to establish the trust relationship between nodes. After the cluster is installed, if the original keys are accidentally disclosed or new keys are used, the system administrator can perform the following operations to manually change the keys.

Prerequisites

- The cluster has been stopped.
- No other management operations are being performed.

Procedure

Step 1 Log in as user **omm** to the node whose SSH keys need to be replaced.

If the node is a Manager management node, run the following command on the active management node.

Step 2 Run the following command to disable logout upon timeout:

TMOUT=0

NOTE

After the operations in this section are complete, run the **TMOUT=***Timeout interval* command to restore the timeout interval in a timely manner. For example, **TMOUT=600** indicates that a user is logged out if the user does not perform any operation within 600 seconds.

Step 3 Run the following command to generate a key for the node:

- If the node is a Manager management node, run the following command:

  sh ${CONTROLLER_HOME}/sbin/update-ssh-key.sh

- If the node is a non-Manager management node, run the following command:

  sh ${NODE_AGENT_HOME}/bin/update-ssh-key.sh

If "Succeed to update ssh private key." is displayed when the preceding command is executed, the SSH key is generated successfully.

Step 4 Run the following command to copy the public key of the node to the active management node:

scp ${HOME}/.ssh/id_rsa.pub** *oms_ip*:**${HOME}/.ssh/id_rsa.pub_bak

oms_ip: indicates the IP address of the active management node.

Enter the password of user **omm** to copy the files.

Step 5 Log in to the active management node as user **omm**.

Step 6 Run the following command to disable logout on system timeout:

TMOUT=0

Step 7 Run the following command to go to the related directory:

cd ${HOME}/.ssh

Step 8 Run the following command to add new public keys:

cat id_rsa.pub_bak >> authorized_keys

Step 9 Run the following command to move the temporary public key file, for example, **/tmp**.

mv -f id_rsa.pub_bak /tmp

Step 10 Copy the **authorized_keys** file of the active management node to the other nodes in the cluster:

scp authorized_keys *node_ip*:**/${HOME}/.ssh/authorized_keys**

node_ip: indicates the IP address of another node in the cluster. Multiple IP addresses are not supported.

Step 11 Run the following command to confirm private key replacement without entering the password:

ssh *node_ip*

node_ip: indicates the IP address of another node in the cluster. Multiple IP addresses are not supported.

Step 12 Log in to FusionInsight Manager. On **Homepage**, locate the desired cluster and choose > **Start** to start the cluster.

----End

8.12.4.14 Changing the Timeout Duration of the Manager Page

FusionInsight Manager allows you to configure the timeout duration of the Manager page based on service requirements. You must properly set the timeout duration to prevent information leakage in long-time exposure of the web page.

NOTE

This function is supported only by MRS 3.3.0 or later.

Changing the Timeout Duration of the Manager Page

Step 1 Log in to FusionInsight Manager.

Step 2 Choose **System** > **OMS**.

Step 3 In the list, locate the row that contains **tomcat** and click **Modify Configuration**.

Step 4 On the displayed page, set **Session Timeout** as required and click **OK**.

NOTICE

- Set the minimum session duration based on service requirements. Otherwise, there will be security risks.
- Currently, you cannot use the method described as follows to change the timeout duration of component web UIs.

----End

8.12.5 Security Maintenance

8.12.5.1 Account Maintenance Suggestions

It is recommended that the administrator conduct routine checks on the accounts. The check covers the following items:

- Check whether the accounts of the OS, FusionInsight Manager, and each component are necessary and whether temporary accounts have been deleted.
- Check whether the permissions of the accounts are appropriate. Different administrators have different rights.
- Check and audit the logins and operation records of all types of accounts.

8.12.5.2 Password Maintenance Suggestions

Accessing portal requires identity authentication. The complexity and validity period of an account password must meet your security requirements.

Refer to the following suggestions to maintain passwords:

1. Assign dedicated personnel to keep OS passwords.
2. Use passwords that meet certain strength requirements, such as minimum password length or mixing of letter cases.
3. Encrypt passwords before transferring them, and do not transfer them via email.
4. Encrypt passwords for storage.
5. Remind enterprise users to change passwords during system handover.
6. Change passwords periodically.

8.12.5.3 Log Maintenance Suggestions

Operation logs help discover exceptions such as illegal operations and login by unauthorized users. The system records important operations in logs. You can use operation logs to locate problems.

Checking Logs Regularly

Check system logs periodically and handle exceptions such as unauthorized operations or logins in a timely manner.

Backing Up Logs Regularly

The audit logs provided by FusionInsight Manager and cluster record the user activities and operations. You can export the audit logs on FusionInsight Manager. If there are too many audit logs in the system, you can configure dump parameters to dump audit logs to a specified server to ensure that the cluster nodes disk space is sufficient.

Maintenance Owner

Network monitoring engineers and system maintenance engineers

8.12.6 Security Statement

JDK Usage Statement

MRS MRS cluster is a big data cluster that provides users with distributed data analysis and computing capabilities. The built-in JDK of MRS MRS is OpenJDK, which is used in the following scenarios:

- Platform service running and maintenance
- Linux client operations, including service submission and application O&M

JDK Risk Description

The system performs permission control on the built-in JDK. Only users in the related group of the FusionInsight platform can access the JDK. In addition, the platform is deployed on a customer's intranet. Therefore, the security risk is low.

JDK Hardening

For details about how to harden the JDK, see "Hardening JDK" in **Hardening Policies**.

Public IP Addresses in Hue

Hue uses the test cases of third-party packages, such as **ipadrress**, **requests**, and **Django**, and uses the public IP addresses in the comments of the test cases. However, these public IP addresses are not involved when Hue provides services, and the Hue configuration file does not involve these public IP addresses.

9 MRS Manager Operation Guide (Applicable to 2.x and Earlier Versions)

9.1 Introduction to MRS Manager

Overview

MRS manages and analyzes massive data and helps you rapidly obtain desired data from structured and unstructured data. The structure of open-source components is complex. The installation, configuration, and management processes are time- and labor-consuming. MRS Manager is a unified enterprise-level cluster management platform and provides the following functions:

- Cluster monitoring enables you to quickly view the health status of hosts and services.
- Graphical metric monitoring and customization enable you to quickly obtain key information about the system.
- Service property configurations can meet service performance requirements.
- With cluster, service, and role instance functions, you can start or stop services and clusters in one click.

Introduction to the MRS Manager GUI

MRS Manager provides a unified cluster management platform, facilitating rapid and easy O&M for clusters. For details about how to access MRS Manager, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Table 9-1 describes the functions of each operation entry.

Figure 9-1 MRS Manager

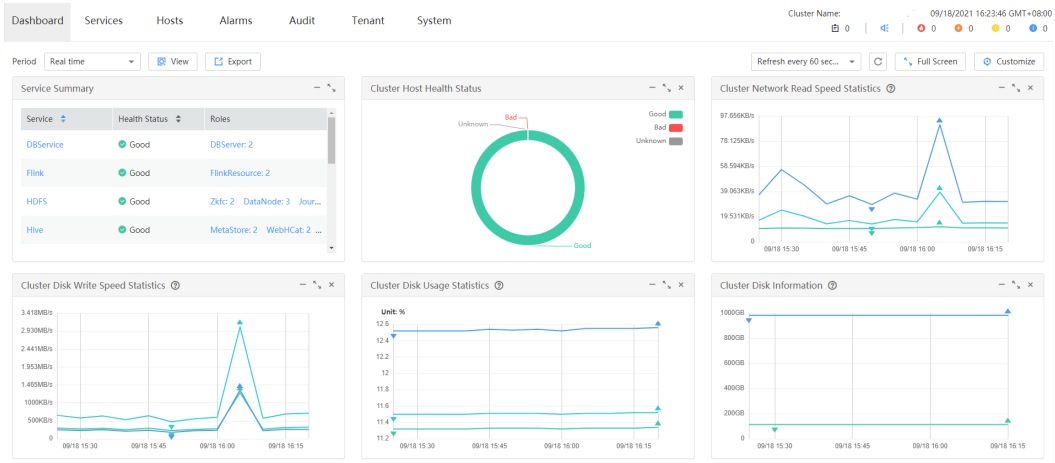

Table 9-1 Functions of each entry on the operation bar

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

Go to the **System** tab page, and switch to another function pages through shortcuts. See **Table 9-2**.

The following is an example of quick redirection through shortcuts:

Step 1 On MRS Manager, click **System**.

Step 2 On the **System** tab page, click a function link. The function page is displayed.

For example, in the **Backup and Restoration** area, click **Back Up Data**. The page for backing up data is displayed.

Step 3 Move the cursor to the left border of the browser window. The **System** black shortcut menu is displayed. After you move the cursor out of the menu, the menu is collapsed.

Step 4 In the shortcut menu that is displayed, you can click a function link to go to the corresponding function page.

For example, choose **Maintenance** > **Export Log**. The page for exporting logs is displayed.

----End

Table 9-2 Shortcut menus on the **System** tab page

Menu	Function Link
Backup and Restoration	Back Up Data
	Restore Data
Maintenance	Export Log
	Export Audit Log
	Check Health Status
Monitoring and Alarm	Configure Syslog
	Configure Alarm Threshold
	Configure SNMP
	Configure Monitoring Metric Dump
	Configure Resource Contribution Ranking
Permission	Manage User
	Manage User Group
	Manage Role
	Configure Password Policy
	Change OMS Database Password
Resource Management	Static service pool
Patch	Manage Patch

Reference

MapReduce Service (MRS) is a data analysis service on the public cloud. It is used to manage and analyze massive sets of data.

MRS uses MRS Manager to manage big data components, such as components in the Hadoop ecosystem. Therefore, some concepts on the MRS Console on the public cloud must be different from those on MRS Manager. For details, see **Table 9-3**.

Table 9-3 Difference Comparison

Concept	Public Cloud MRS	MRS Manager
MapReduce Service	Indicates the data analysis cloud service on the public cloud, called MRS. This service includes components such as Hive, Spark, Yarn, HDFS, and ZooKeeper.	Provides a unified management platform for big data components in tenant clusters.

9.2 Checking Running Tasks

Scenario

When you perform operations on MRS Manager to trigger a task, the task execution process and progress are displayed. After the task window is closed, you need to open the task window by using the task management function.

MRS Manager reserves 10 latest tasks by default, for example, restarting services, synchronizing service configurations, and performing health check.

Procedure

Step 1 On MRS Manager, click  to open the task list.

You can view the following information in the task list: **Name**, **Status**, **Progress**, **Start Time** and **End Time**.

Step 2 Click the target task name to view the detailed information about the running task.

----End

9.3 Monitoring Management

9.3.1 Dashboard

On MRS Manager, nodes in a cluster can be classified into management nodes, control nodes, and data nodes. The change trends of key host monitoring metrics

on each type of node can be calculated and displayed as curve charts in reports based on the customized periods. If a host belongs to multiple node types, the metric statistics will be repeatedly collected.

This section provides overview of MRS clusters and describes how to view, customize, and export node monitoring metrics on MRS Manager.

Procedure

Step 1 Log in to MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 Choose **Dashboard** on MRS Manager.

Step 3 In **Period**, you can specify a period to view monitoring data. The options are as follows:

- Real time
- Last 3 hours
- Last 6 hours
- Last 24 hours
- Last week
- Last month
- Last 3 months
- Last 6 months
- Customize. If you select this option, you can customize the period for viewing monitoring data.

Step 4 Click **View** to view monitoring data in a period.

- You can view **Health Status** and **Roles** of each service on the **Service Summary** page of MRS Manager.
- Click  above the curve chart to view details about a metric.

Step 5 Customize a monitoring report.

1. Click **Customize** and select monitoring metrics to be displayed on MRS Manager.

   MRS Manager supports a maximum of 14 monitoring metrics, but at most 12 customized monitoring metrics can be displayed on the page.

   - Cluster Host Health Status
   - Cluster Network Read Speed Statistics
   - Host Network Read Speed Distribution
   - Host Network Write Speed Distribution
   - Cluster Disk Write Speed Statistics
   - Cluster Disk Usage Statistics
   - Cluster Disk Information
   - Host Disk Usage Distribution
   - Cluster Disk Read Speed Statistics

- Cluster Memory Usage Statistics
- Host Memory Usage Distribution
- Cluster Network Write Speed Statistics
- Host CPU Usage Distribution
- Cluster CPU Usage Statistics

2. Click **OK** to save the selected monitoring metrics for display.

   NOTE

   Click **Clear** to cancel all the selected monitoring metrics in a batch.

Step 6 Set an automatic refresh interval or click for an immediate refresh.

The following refresh interval options are supported:

- Refresh every 60 seconds
- **Refresh every 120 seconds**
- Stop refreshing

  NOTE

  If you select **Full Screen**, the **Dashboard** window will be maximized.

Step 7 Export a monitoring report.

1. Select a period. The options are as follows:
   - Real time
   - Last 3 hours
   - Last 6 hours
   - Last 24 hours
   - Last week
   - Last month
   - Last 3 months
   - Last 6 months
   - Customize. If you select this option, you can customize a time of period to export a report.
2. Click **Export**. MRS Manager will generate a report about the selected monitoring metrics in a specified time of period. Save the report.

   NOTE

   To view the curve charts of monitoring metrics in a specified period, click **View**.

----End

9.3.2 Managing Services and Monitoring Hosts

You can manage the following status and indicators of all services (including role instances) and hosts on the MRS Manager:

- Status information: includes operation, health, configuration, and role instance status.

- Metric information: includes key monitoring metrics for services.
- Metric export: allows you to export monitoring reports.

NOTE

Set an automatic refresh interval or click  for an immediate refresh.

The following refresh interval options are supported:

- Refresh every 30 seconds
- Refresh every 60 seconds
- Stop refreshing

Managing Service Monitoring

Step 1 On MRS Manager, click **Services**.

The service list includes **Service**, **Operating Status**, **Health Status**, **Configuration Status**, **Roles**, and **Operation** are displayed in the component list.

- **Table 9-4** describes the service operating status.

Table 9-4 Service operating status

Status	Description
Started	The service is started.
Stopped	The service is stopped.
Failed to start	Failed to start the role instance.
Failed to stop	Failed to stop the role instance.
Unknown	Indicates initial service status after the background system restarts.

- **Table 9-5** describes the service health status.

Table 9-5 Service health status

Status	Description
Good	Indicates that all role instances in the service are running properly.
Bad	Indicates that the running status of at least one role instance is **Faulty** or the status of the service on which the current service depends is abnormal.
Unknown	Indicates that all role instances in the service are in the **Unknown** state.
Concerning	Indicates that the background system is restarting the service.

Status	Description
Partially Healthy	Indicates that the status of the service on which the service depends is abnormal, and APIs related to the abnormal service cannot be invoked by external systems.

- **Table 9-6** describes the service health status.

  Table 9-6 Service configuration status

Status	Description
Synchronized	The latest configuration takes effect.
Expired	The latest configuration does not take effect after the parameter modification. Related services need to be restarted.
Failed	The communication is incorrect or data cannot be read or written during the parameter configuration. Use **Synchronize Configuration** to rectify the fault.
Configuring	Parameters are being configured.
Unknown	Current configuration status cannot be obtained.

By default, the **Service** column is sorted in ascending order. You can click the icon next to **Service**, **Operating Status**, **Health Status**, or **Configuration Status** to change the sorting mode.

Step 2 Click a specified service in the list to view its status and metric information.

Step 3 Customize monitoring metrics and export customized monitoring information.

1. In the **Charts** area, click **Customize** to customize service monitoring metrics.
2. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.
3. Click **Export** to export the displayed metrics.

----End

Managing Role Instances

Step 1 On MRS Manager, click **Services** and click the target service name in the service list.

Step 2 Click **Instance** to view the role status.

The role instance list contains the **Role**, **Host Name**, **OM IP Address**, **Business IP Address**, **Rack**, **Operation Status**, **Health Status,** and **Configuration Status** of an instance.

- Table 9-7 shows the configuration status of a role instance.

  Table 9-7 Role instance status

Status	Description
Started	The role instance has been started.
Stopped	The role instance has been stopped.
Failed to start	Failed to start the role instance.
Failed to stop	Failed to stop the role instance.
Decommissioning	The role instance is being decommissioned.
Decommissioned	The role instance has been decommissioned.
Recommissioning	The role instance is being recommissioned.
Unknown	Indicates initial role instance status after the background system restarts.

- Table 9-8 shows the health status of a role instance.

  Table 9-8 Role instance health status

Status	Description
Good	The role instance is running properly.
Restoring	The background system is restarting a role instance.
Bad	The role instance is abnormal. For example, the port cannot be accessed if PID does not exist.
Unknown	The host where a role instance resides does not connect to the background system.
Partially Healthy	The role instance is partially running properly.

- Table 9-9 shows the configuration status of a role instance.

  Table 9-9 Role instance configuration status

Status	Description
Synchronized	The latest configuration takes effect.
Expired	The latest configuration does not take effect after the parameter modification. Related services need to be restarted.

Status	Description
Failed	The communication is incorrect or data cannot be read or written during the parameter configuration. Use **Synchronize Configuration** to rectify the fault.
Configuring	Parameters are being configured.
Unknown	Current configuration status cannot be obtained.

By default, the **Role** column is sorted in ascending order. You can click the sorting icon next to **Role**, **Host Name**, **OM IP Address**, **Business IP Address**, **Rack**, **Operating Status**, **Health Status**, or **Configuration Status** to change the sorting mode.

You can filter out all instances of the same role in the **Role** column.

You can set search criteria in the role search area by clicking **Advanced Search**, and click **Search** to view specified role information. Click **Reset** to clear the search criteria. Fuzzy search is supported.

Step 3 Click the target role instance to view its status and metric information.

Step 4 Customize monitoring metrics and export customized monitoring information.

1. In the **Charts** area, click **Customize** to customize service monitoring metrics.
2. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.
3. Click **Export** to export the displayed metrics.

----End

Managing Hosts

Step 1 On MRS Manager, click **Hosts** to view the status of all hosts.

The host list contains the host name, management IP address, service IP address, rack, network speed, operating status, health status, disk usage, memory usage, and CPU usage.

- **Table 9-10** shows the host operating status.

  Table 9-10 Host operating status

Status	Description
Normal	The host and service roles on the host are running properly.
Isolated	The host is isolated, and the service roles on the host stop running.

- **Table 9-11** describes the host health status.

Table 9-11 Host health status

Status	Description
Good	The host can properly send heartbeats.
Bad	The host fails to send heartbeats due to timeout.
Unknown	The host initial status is unknown during the operation of adding or deleting a host.

By default, the **Host Name** column is sorted by host name in ascending order. You can click the sorting icon next to **Host Name**, **OM IP Address**, **Business IP Address**, **Rack**, **Network Speed**, **Operating Status**, **Health Status**, **Disk Usage**, **Memory Usage**, or **CPU Usage** to change the sorting mode.

You can set search criteria in the role search area by clicking **Advanced Search**, and click **Search** to view specified role information. Click **Reset** to clear the search criteria. Fuzzy search is supported.

Step 2 Click the target host in the host list to view its status and metric information.

Step 3 Customize monitoring metrics and export customized monitoring information.

1. In the **Charts** area, click **Customize** to customize service monitoring metrics.
2. In **Period** area, select a time of period and click **View** to view the monitoring data within the time period.
3. Click **Export** to export the displayed metrics.

----End

9.3.3 Managing Resource Distribution

On MRS Manager, you can query the top value curves, bottom value curves, or average data curves of key service and host monitoring metrics, that is, the resource distribution information. MRS Manager allows you to view the monitoring data of the last hour.

You can also modify the resource distribution on MRS Manager to display both the top and bottom value curves in service and host resource distribution figures.

Resource distribution of some monitoring metrics is not recorded.

Procedure

- View the resource distribution of service monitoring metrics.
  a. On MRS Manager, click **Services**.
  b. Select the target service from the service list.
  c. Click **Resource Distribution**.
     Select key metrics of the service from **Metric**. MRS Manager displays the resource distribution of the metrics in the last hour.
- View the resource distribution of host monitoring metrics.

a. Click **Hosts**.
b. Click the name of the specified host in the host list.
c. Click **Resource Distribution**.

   Select key metrics of the host from **Metrics**. MRS Manager displays the resource distribution of the metrics in the last hour.

- Configure resource distribution.
  a. On MRS Manager, click **System**.
  b. In **Configuration**, click **Configure Resource Contribution Ranking** under **Monitoring and Alarm**.
  c. Change the number of resources to be displayed.
     - Set **Number of Top Resources** to the number of top values.
     - Set **Number of Bottom Resources** to the number of bottom values.

     NOTE

     The sum of the maximum value and minimum value of resource distribution cannot be greater than 5.
  d. Click **OK** to save the configurations.

     The message "Number of top and bottom resources saved successfully" is displayed in the upper right corner of the page.

9.3.4 Configuring Monitoring Metric Dumping

You can configure interconnection parameters on MRS Manager to save monitoring metric data to a specified FTP server using the FTP or SFTP protocol. In this way, MRS clusters can interconnect with third-party systems. The FTP protocol does not encrypt data, which brings potential security risks. Therefore, the SFTP protocol is recommended.

MRS Manager supports the collection of all the monitoring metric data in the managed clusters. The collection period is 30 seconds, 60 seconds, or 300 seconds. The monitoring metric data is stored to different monitoring files on the FTP server by collection period. The monitoring file naming rule is in the "*Cluster name*_**metric**_*Monitoring metric data collection period*_*File saving time*.**log**" format.

Prerequisites

The ECS corresponding to the dump server must be in the same VPC as the Master node of the MRS cluster, and the Master node can access the IP address and specified port of the dump server. The FTP service on the dump server is running properly.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In **Configuration**, click **Configure Monitoring Metric Dump** under **Monitoring and Alarm**.

Step 3 Table 9-12 describes dump parameters.

Table 9-12 Dump parameters

Parameter	Description	Mandatory
Dump Monitoring Metric	Whether to enable monitoring metric data interconnection. • : Enabled. • : Disabled.	Yes
FTP IP Address	FTP server for storing monitoring files after monitoring metric data is interconnected.	Yes
FTP Port	Port for connecting to the FTP server.	Yes
FTP Username	Username for logging in to the FTP server.	Yes
FTP Password	Password for logging in to the FTP server.	Yes
Save Path	Path for storing monitoring files on the FTP server.	Yes
Dump Interval (s)	Interval at which monitoring files are periodically stored on the FTP server, in seconds.	Yes
Dump Mode	Protocol used for sending monitoring files. The options are **FTP** and **SFTP**.	Yes
SFTP Public Key	Public key of the FTP server. This parameter is available only when **Dump Mode** is set to **SFTP**. You are advised to configure a public key. Otherwise, security risks may arise.	No

Step 4 Click **OK** to complete the settings.

----End

9.4 Alarm Management

9.4.1 Viewing and Manually Clearing an Alarm

Scenario

You can view and clear alarms on MRS Manager.

Generally, the system automatically clears an alarm when the fault is rectified. If the fault has been rectified and the alarm cannot be automatically cleared, you can manually clear the alarm.

You can view the latest 100,000 alarms (including uncleared, manually cleared, and automatically cleared alarms) on MRS Manager. If the number of cleared alarms exceeds 100,000 and is about to reach 110,000, the system automatically

dumps the earliest 10,000 cleared alarms to **${BIGDATA_HOME}/OMSV100R001C00x8664/workspace/data** on the active management node. A directory is automatically generated when alarms are dumped for the first time.

NOTE

Set an automatic refresh interval or click  for an immediate refresh.

The following refresh interval options are supported:

- Refresh every 30 seconds
- Refresh every 60 seconds
- Stop refreshing

Procedure

Step 1 On MRS Manager, click **Alarms** to view the alarm information in the alarm list.

- By default, the alarm list page displays the latest 10 alarms.
- By default, alarms are displayed in descending order by **Generated**. You can click **Alarm ID**, **Alarm Name**, **Severity**, **Generated**, **Location**, **Operation** to change the display mode.
- You can filter all alarms of the same severity in **Severity**, including cleared and uncleared alarms.
- You can click  ,  , , or  to filter out **Critical**, **Major**, **Minor**, or **Warning** alarms.

Step 2 Click **Advanced Search**. In the displayed alarm search area, set search criteria and click **Search** to view the information about specified alarms. Click **Reset** to clear the search criteria.

NOTE

You can set the **Start Time** and **End Time** to specify the time range. You can search for alarms generated within the time range.

Handle the alarm by referring to **Alarm Reference**. If the alarms in some scenarios are generated due to other cloud services that MRS depends on, you need to contact maintenance personnel of the corresponding cloud services.

Step 3 If the alarm needs to be manually cleared after errors are rectified, click **Clear Alarm**.

NOTE

If multiple alarms have been handled, you can select one or more alarms to be cleared and click **Clear Alarm** to clear the alarms in batches. A maximum of 300 alarms can be cleared in each batch.

----End

9.4.2 Configuring an Alarm Threshold

Scenario

You can configure an alarm threshold to learn the metric health status. After **Send Alarm** is selected, the system sends an alarm message when the monitored data reaches the alarm threshold. You can view the alarm information in **Alarms**.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In **Configuration**, click **Configure Alarm Threshold** under **Monitoring and Alarm**, select monitoring metrics as planned, and set their baselines.

Step 3 Click a metric, for example, **CPU Usage**, and click **Create Rule**.

Step 4 Set the monitoring metric rule parameters on the displayed configuration page.

Table 9-13 Monitoring metric rule parameters

Parameter	Description	Value
Rule Name	Specifies the rule name.	CPU_MAX (example)
Reference Date	Specifies the date on which the reference indicator history is generated.	2014/11/06 (example)
Threshold Type	Specifies the maximum or minimum value of a metric. If this parameter is set to **Max. Value**, the system generates an alarm when the actual value of the metric is greater than the threshold. If this parameter is set to **Min. Value**, the system generates an alarm when the actual value of the metric is smaller than the threshold.	• Max. value • Min. value
Alarm Severity	Alarm Severity	• Critical • Major • Minor • Warning
Time Range	Specifies the period in which the rule takes effect.	From 00:00 to 23:59 (example)
Threshold	Specifies the threshold of the rule monitoring metrics.	80 (example)

Parameter	Description	Value
Date	Specifies the type of date when the rule takes effect.	• Workday • Weekend • Other
Add Date	This parameter is valid only when **Date** is set to **Other**. You can select multiple dates.	11/30 (example)

Step 5 Click **OK**. A message is displayed in the upper right corner of the page, indicating that the template is saved successfully.

Send alarm is selected by default. MRS Manager checks whether the value of each monitored metric reaches the threshold. If the number of consecutive check times is equal to the value of **Trigger Count**, and the threshold is not reached in these checks, the system sends an alarm. The value can be customized. **Check Period (s)** indicates the interval at which MRS Manager checks monitoring metrics.

Step 6 Locate the row that contains the newly added rule, and click **Apply** in the **Operation** column. A message is displayed in the upper right corner, indicating that the rule *xx* is successfully added. Click **Cancel** in the **Operation** column. A message is displayed in the upper right corner, indicating that the rule *xx* is successfully canceled.

----End

9.4.3 Configuring Syslog Northbound Interface Parameters

Scenario

You can configure the northbound interface so that alarms generated on MRS Manager can be reported to your monitoring O&M system using Syslog.

NOTICE

If the Syslog protocol is not encrypted, data may be stolen.

Prerequisites

The ECS corresponding to the server must be in the same VPC as the Master node of the MRS cluster, and the Master node can access the IP address and specified port of the server.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In **Configuration**, click **Configure Syslog** under **Monitoring and Alarm**.

The **Syslog Service** is disabled by default. Click the switch to enable the Syslog service.

Step 3 Set the interconnection parameters listed in **Table 9-14**.

Table 9-14 Syslog parameters

Area	Parameter	Description
Syslog Protocol	Service IP Address	Specifies the IP address of the interconnection server.
	Server Port	Specifies the port number for interconnection.
	Protocol	Specifies the protocol type. The options are as follows: • **TCP** • **UDP**
	Severity	Specifies the severity of the reported message. The options are as follows: • **Informational** • **Emergency** • **Alert** • **Critical** • **Error** • **Warning** • **Notice** • **Debug**
	Facility	Specifies the module where the log is generated.
	Identifier	Specifies the product ID. The default value is **MRS Manager**.
Report Message	Report Format	Specifies the message format of the alarm report. For details, see help information on the web page.

Area	Parameter	Description
	Alarm Status	Specifies the type of the alarm to be reported. • **Fault**: indicates that the Syslog alarm message is reported when MRS Manager generates an alarm. • **Clear**: indicates that a Syslog alarm message is reported when an alarm on MRS Manager is cleared. • **Event**: indicates that the Syslog alarm message is reported when MRS Manager generates an event.
	Report Alarm Severity	Specifies the level of the alarm to be reported. The value can be **Suggestion**, **Minor**, **Major**, and **Critical**.
Uncleared Alarm Reporting	Periodic Uncleared Alarm Report	Specifies whether uncleared alarms are reported periodically. By default, the switch of **Periodic Uncleared Alarm Reporting** is disabled. You can click the switch to enable it.
	Report Interval (min)	Specifies the interval for periodically reporting uncleared alarms to the remote Syslog service. This parameter is valid only when **Periodic Uncleared Alarm Reporting** switch is enabled. The unit is minute. The default value is **15**. The value ranges from 5 minutes to one day (1,440 minutes).

Area	Parameter	Description
Heartbeat Settings	Heartbeat Report	Specifies whether to periodically report Syslog heartbeat messages. By default, the switch of **Periodic Uncleared Alarm Reporting** is disabled. You can click the switch to enable it.
	Heartbeat Period (min)	Specifies the interval for periodically reporting heartbeat messages. This parameter is valid only when **Heartbeat Report** switch is enabled. The unit is minute. The default value is **15**. The value ranges from 1 to 60.
	Heartbeat Packet	Specifies the content of the reported heartbeat message. This parameter is enabled when **Heartbeat Report** is enabled. The value can contain a maximum of 256 characters, including digits, letters, underscores (_), vertical bars (|), colons (:), spaces, commas (,), and periods (.).

NOTE

After the periodic heartbeat packet function is enabled, packets may be interrupted during automatic recovery of some cluster error tolerance (for example, active/standby management node switchover). In this case, wait for automatic recovery.

Step 4 Click **OK** to complete the settings.

----End

9.4.4 Configuring SNMP Northbound Interface Parameters

Scenario

You can configure the northbound interface so that alarms and monitoring metrics on MRS Manager can be integrated to the network management platform using SNMP.

Prerequisites

The ECS corresponding to the server must be in the same VPC as the Master node of the MRS cluster, and the Master node can access the IP address and specified port of the server.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In **Configuration**, click **Configure SNMP** under **Monitoring and Alarm**.

The **SNMP Service** is disabled by default. Click the switch to enable the SNMP service.

Step 3 Set the interconnection parameters listed in **Table 9-15**.

Table 9-15 Syslog parameters

Parameter	Description
Version	Specifies the version of the SNMP, which can be: • v2c: an earlier version with low security • v3: the latest version of SNMP with higher security than SNMPv2c The SNMP v3 version is recommended.
Local Port	Specifies the local port. The default value is **20000**. The value ranges from **1025** to **65535**.
Read Community Name	Specifies the read-only community name. This parameter is valid only when **Version** is set to **v2c**.
Write Community Name	Specifies the write community name. This parameter is valid only when **Version** is set to **v2c**.
Security Username	Specifies the SNMP security username. This parameter is valid only when **Version** is set to **v3**.
Authentication Protocol	Specifies the authentication protocol. You are advised to set this parameter to set this parameter to **SHA**. This parameter is valid only when **Version** is set to **v3**.
Authentication Password	Specifies the authentication key. This parameter is valid only when **Version** is set to **v3**.
Confirm Password	Used to confirm the authentication key. This parameter is valid only when **Version** is set to **v3**.
Encryption Protocol	Specifies the encryption protocol. You are advised to set this parameter to **AES256**. This parameter is valid only when **Version** is set to **v3**.
Encryption Password	Specifies the encryption key. This parameter is valid only when **Version** is set to **v3**.

Parameter	Description
Confirm Password	Used to confirm the encryption key. This parameter is valid only when **Version** is set to **v3**.

NOTE

- The **Authentication Password** and **Encryption Password** must contain 8 to 16 characters, including at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters. The two passwords must be different. The two passwords cannot be the same as the security username or the reverse of the security username.
- For security purposes, periodically change the authentication password and encryption password when the SNMP protocol is used.
- If SNMPv3 is used, a security user will be locked after five consecutive authentication failures within 5 minutes. The user will be automatically unlocked 5 minutes later.

Step 4 Click **Create Trap Target** in the **Trap Target** area. In the displayed dialog box, set the following parameters:

- **Target Symbol** specifies the trap target ID, which is the ID of the NMS or host that receives traps. The value consists of 1 to 255 characters, including letters or digits.
- **Target IP Address** specifies the IP address of the target trap. IP addresses of class A, B, and C can be used to communicate with the IP address of the management plane of the management node.
- **Target Port** specifies the port receiving traps. The port number must be consistent with the peer end and ranges from 0 to 65535.
- **Trap Community Name** is valid only when **Version** is set to **v2c**.

Click **OK**. The **Create Trap Target** dialog box is closed.

Step 5 Click **OK** to complete the settings.

----End

9.5 Alarm Reference (Applicable to MRS 2.x and Earlier Versions)

9.5.1 ALM-12001 Audit Log Dump Failure (For MRS 2.*x* or Earlier)

Description

Cluster audit logs need to be dumped on a third-party server due to the local historical data backup policy. Audit logs can be successfully dumped if the dump server meets the configuration conditions. This alarm is generated when the audit log dump fails because the disk space of the dump directory on the third-party server is insufficient or a user changes the username, password, or dump directory of the dump server.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12001	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system can only store a maximum of 50 dump files locally. If the fault persists on the dump server, the local audit log may be lost.

Possible Causes

- The network connection is abnormal.
- The username, password, or dump directory of the dump server does not meet the configuration conditions.
- The disk space of the dump directory is insufficient.

Procedure

Step 1 Check whether the username, password, and dump directory are correct.

1. Check on the dump configuration page of MRS Manager to see if they are correct.
   - If yes, go to **Step 3**.
   - If no, go to **Step 1.2**.
2. Change the username, password, or dump directory, and click **OK**.
3. Wait 2 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Reset the dump rule.

1. On MRS Manager, choose **System** > **Dump Audit Log**.
2. Reset dump rules, set the parameters properly, and click **OK**.

3. Wait 2 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

N/A

9.5.2 ALM-12002 HA Resource Abnormal (For MRS 2.*x* or Earlier)

Description

The high availability (HA) software periodically checks the WebService floating IP addresses and databases of Manager. This alarm is generated when the HA software detects that the WebService floating IP addresses or databases are abnormal.

This alarm is cleared when the HA software detects that the floating IP addresses or databases are normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12002	Major	Yes

Parameter

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
RESName	Specifies the resource for which the alarm is generated.

Impact on the System

If the WebService floating IP addresses of Manager are abnormal, users cannot log in to or use Manager. If databases of Manager are abnormal, all core services and related service processes, such as alarms and monitoring functions, are affected.

Possible Causes

- The floating IP address is abnormal.
- The database is abnormal.

Procedure

Step 1 Check the floating IP address status of the active management node.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host address and resource name of the alarm.
2. Log in to the active management node. Run the following commands to switch the user:

   sudo su - root

   su - omm
3. Go to the **${BIGDATA_HOME}/om-0.0.1/sbin/** directory, run the **status-oms.sh** script to check whether the floating IP address of the active Manager is normal. View the command output, locate the row where **ResName** is **floatip**, and check whether the following information is displayed.

   Example:

   ```
   10-10-10-160 floatip Normal Normal Single_active
   ```

   - If yes, go to **Step 2**.
   - If no, go to **Step 1.4**.
4. Contact the O&M personnel to check whether the floating IP NIC exists.
   - If yes, go to **Step 2**.
   - If no, go to **Step 1.5**.
5. Contact O&M personnel to rectify the NIC fault.

   Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Check the database status of the active and standby management nodes.

1. Log in to the active and standby management nodes, run the **sudo su - root** and **su - ommdba** commands to switch to user **ommdba**, and run the **gs_ctl query** command to check whether the following information is displayed in the command output.

   Command output of the active management node:

   ```
   Ha state:
   LOCAL_ROLE: Primary
   STATIC_CONNECTIONS: 1
   DB_STATE: Normal
   DETAIL_INFORMATION: user/password invalid
   ```



```
 Senders info:
No information
 Receiver info:
No information
```

Command output of the standby management node:

```
Ha state:
LOCAL_ROLE: Standby
STATIC_CONNECTIONS: 1
DB_STATE : Normal
DETAIL_INFORMATION: user/password invalid
 Senders info:
No information
 Receiver info:
No information
```

- If yes, go to **Step 2.3**.
- If no, go to **Step 2.2**.

2. Contact the O&M personnel to check whether a network fault occurs and rectify the fault.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.3 ALM-12004 OLdap Resource Abnormal (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the Ldap resource in Manager is abnormal.

This alarm is cleared when the Ldap resource in Manager recovers and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12004	Major	Yes

Parameter

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The Manager authentication services are unavailable and cannot provide security authentication and user management functions for web upper-layer services. Users may be unable to log in to Manager.

Possible Causes

The LdapServer process in Manager is abnormal.

Procedure

Step 1 Check whether the LdapServer process in Manager is normal.

1. Log in to the active management node.
2. Run **ps -ef | grep slapd** to check whether the LdapServer resource process in the **${BIGDATA_HOME}/om-0.0.1/** directory of the configuration file is running properly.

   You can determine that the resource is normal as follows:

   a. Run **sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh** and find that **ResHAStatus** of the OLdap process is **Normal**.

   b. Run **ps -ef | grep slapd** and find that the slapd process occupies port 21750.

   - If yes, go to **Step 2**.
   - If no, go to **Step 3**.

Step 2 Run **kill -2** *PID of the LdapServer process* and wait 20 seconds. The HA starts the OLdap process automatically. Check whether the status of the OLdap resource is normal.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.4 ALM-12005 OKerberos Resource Abnormal (For MRS 2.*x* or Earlier)

Description

The alarm module monitors the status of the Kerberos resource in Manager. This alarm is generated when the Kerberos resource is abnormal.

This alarm is cleared when the alarm handling is complete and the Kerberos resource status recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12005	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The authentication services are unavailable and cannot provide security authentication functions for web upper-layer services. Users may be unable to log in to MRS Manager.

Possible Causes

The OLdap resource on which OKerberos depends is abnormal.

Procedure

Step 1 Check whether the OLdap resource on which OKerberos depends is abnormal in Manager.

1. Log in to the active management node.

2. Run the following command to check whether the OLdap resource managed by HA is normal:

   sh ${BIGDATA_HOME}/OMSV100R001C00x8664/workspace0/ha/module/hacom/script/status_ha.sh

   The OLdap resource is normal when the OLdap resource is in the **Active_normal** state on the active node and in the **Standby_normal** state on the standby node.

   - If yes, go to **Step 3**.
   - If no, go to **Step 2**.

Step 2 Resolve the problem by following the instructions in **ALM-12004 OLdap Resource Abnormal (For MRS 2.x or Earlier)**. After the OLdap resource status recovers, check whether the OKerberos resource is normal.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.5 ALM-12006 Node Fault (For MRS 2.*x* or Earlier)

Description

Controller checks the NodeAgent status every 30 seconds. This alarm is generated when Controller fails to receive the status report of a NodeAgent for three consecutive times.

This alarm is cleared when Controller can properly receive the status report of the NodeAgent.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12006	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Services on the node are unavailable.

Possible Causes

The network is disconnected, or the hardware is faulty.

Procedure

Step 1 Check whether the network is disconnected or the hardware is faulty.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host address of the alarm.
2. Log in to the active management node.
3. Run the following command to check whether the faulty node is reachable:

   ping *IP address of the faulty host*

   a. If yes, go to **Step 2**.

   b. If no, go to **Step 1.4**.
4. Contact the O&M personnel to check whether the network is faulty.
   - If yes, go to **Step 2**.
   - If no, go to **Step 1.6**.
5. Rectify the network fault and check whether the alarm is cleared from the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 1.6**.
6. Contact the O&M personnel to check whether a hardware fault (for example, a CPU or memory fault) occurs on the node.
   - If yes, go to **Step 1.7**.
   - If no, go to **Step 2**.
7. Repair the faulty components and restart the node. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.6 ALM-12007 Process Fault (For MRS 2.*x* or Earlier)

Description

The process health check module checks the process status every 5 seconds. This alarm is generated when the process health check module detects that the process connection status is Bad for three consecutive times.

This alarm is cleared when the process can be connected.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12007	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The service provided by the process is unavailable.

Possible Causes

- The instance process is abnormal.
- The drive space is insufficient.

Procedure

Step 1 Check whether the instance process is abnormal.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host name and service name of the alarm.
2. On the **Alarms** page, check whether the alarm **ALM-12006 Node Fault (For MRS 2.x or Earlier)** is generated.

   If yes, go to **Step 1.3**.

   If no, go to **Step 1.4**.
3. Handle the alarm by following the instructions in **ALM-12006 Node Fault (For MRS 2.x or Earlier)**.
4. Check whether the installation directory user, user group, and permission of the alarm role are correct. The correct user, user group, and the permission are **omm**, **ficommon**, and **750**, respectively.
   - If yes, go to **Step 1.6**.
   - If no, go to **Step 1.5**.
5. Run the following commands to set the permission to **750** and **User:Group** to **omm:ficommon**:

   chmod 750 *<folder_name>*

   chown omm:ficommon *<folder_name>*
6. Wait 5 minutes and check whether the ALM-12007 Process Fault alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check whether the disk space is insufficient.

1. On the MRS cluster details page, click the alarm management tab and check whether ALM-12017 Insufficient Disk Capacity is generated in the alarm list.
   - If yes, go to **Step 2.2**.
   - If no, go to **Step 3**.
2. Handle the alarm by following the instructions in **ALM-12017 Insufficient Disk Capacity (For MRS 2.x or Earlier)**.
3. Wait 5 minutes and check whether the ALM-12017 Insufficient Disk Capacity alarm is cleared.

   If yes, go to **Step 2.4**.

   If no, go to **Step 3**.
4. Wait 5 minutes and check whether the alarm is cleared.

   If yes, no further action is required.

   If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.7 ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the active Manager does not receive any heartbeat signal from the standby Manager within 7 seconds.

This alarm is cleared when the active Manager receives heartbeat signals from the standby Manager.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12010	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local Manager HA Name	Specifies a local Manager HA.
Peer Manager HA Name	Specifies a peer Manager HA.

Impact on the System

When the active Manager process is abnormal, an active/standby failover cannot be performed, and services are affected.

Possible Causes

The link between the active and standby Manager servers is abnormal.

Procedure

Step 1 Check whether the network between the active and standby Manager servers is normal.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the standby Manager server.

2. Log in to the active management node.
3. Run the following command to check whether the standby Manager is reachable:

   ping *heartbeat IP address of the standby Manager*

   - If yes, go to **Step 2**.
   - If no, go to **Step 1.4**.
4. Contact the O&M personnel to check whether the network is faulty.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Rectify the network fault and check whether the alarm is cleared from the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Log in to all master nodes in the cluster and run the following commands to find all **sed***xxx* files and delete them:

find /srv/BigData/ -name "sed*"

find /opt -name "sed*"

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.8 ALM-12011 Data Synchronization Exception Between the Active and Standby Manager Nodes (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the standby Manager fails to synchronize files with the active Manager.

This alarm is cleared when the standby Manager synchronizes files with the active Manager.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12011	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local Manager HA Name	Specifies a local Manager HA.
Peer Manager HA Name	Specifies a peer Manager HA.

Impact on the System

Because the configuration files on the standby Manager are not updated, some configurations will be lost after an active/standby switchover. Manager and some components may not run properly.

Possible Causes

The link between the active and standby Manager nodes is interrupted.

Procedure

Step 1 Check whether the network between the active and standby Manager servers is normal.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the standby Manager server.
2. Log in to the active management node. Run the following command to check whether the standby Manager is reachable:

   ping *IP address of the standby Manager*

   - If yes, go to **Step 2**.
   - If no, go to **Step 1.3**.
3. Contact the O&M personnel to check whether the network is faulty.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2**.
4. Rectify the network fault and check whether the alarm is cleared from the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.9 ALM-12012 NTP Service Abnormal (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the NTP service on the current node fails to synchronize time with the NTP service on the active OMS node.

This alarm is cleared when the NTP service on the current node synchronizes time properly with the NTP service on the active OMS node.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12012	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The time on the node is inconsistent with that on other nodes in the cluster. Therefore, some MRS applications on the node may not run properly.

Possible Causes

- The NTP service on the current node cannot start properly.
- The current node fails to synchronize time with the NTP service on the active OMS node.

- The key value authenticated by the NTP service on the current node is inconsistent with that on the active OMS node.
- The time offset between the node and the NTP service on the active OMS node is large.

Procedure

Step 1 Check the NTP service on the current node.

1. Check whether the ntpd process is running on the node using the following method. Log in to the node for which the alarm is generated and run the **sudo su - root** command to switch to user **root**. Then run the following command to check whether the command output contains the ntpd process:

   ps -ef | grep ntpd | grep -v grep

   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.2**.
2. Run **service ntp start** to start the NTP service.
3. Wait 10 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check whether the current node can synchronize time properly with the NTP service on the active OMS node.

1. Check whether the node can synchronize time with the NTP service on the active OMS node based on additional information of the alarm.

   If yes, go to **Step 2.2**.

   If no, go to **Step 3**.
2. Check whether the synchronization with the NTP service on the active OMS node is faulty.

   Log in to the node for which the alarm is generated, run the **sudo su - root** command to switch to user **root**, and run the **ntpq -np** command.

   If an asterisk (*) exists before the IP address of the NTP service on the active OMS node in the command output, the synchronization is in normal state. The command output is as follows:

   ```
   remote refid st t when poll reach delay offset jitter
   ==============================================================================
   *10.10.10.162 .LOCL. 1 u 1 16 377 0.270 -1.562 0.014
   ```

   If there is no asterisk (*) before the IP address of the NTP service on the active OMS node, as shown in the following command output, and the value of **refid** is **.INIT.**, the synchronization is abnormal.

   ```
   remote refid st t when poll reach delay offset jitter
   ==============================================================================
   10.10.10.162 .INIT. 1 u 1 16 377 0.270 -1.562 0.014
   ```

   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
3. Rectify the fault, wait 10 minutes, and then check whether the alarm is cleared.

   An NTP synchronization failure is usually related to the system firewall. If the firewall can be disabled, disable it and then check whether the fault is

rectified. If the firewall cannot be disabled, check the firewall configuration policies and ensure that port **UDP 123** is enabled (you need to follow specific firewall configuration policies of each system).

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Check whether the key value authenticated by the NTP service on the current node is consistent with that on the active OMS node.

Run **cat** to check whether the authentication code whose key value index is 1 is the same as the value of the NTP service on the active OMS node.

- If yes, go to **Step 4.1**.
- If no, go to **Step 5**.

Step 4 Check whether the time offset between the node and the NTP service on the active OMS node is large.

1. Check whether the time offset is large in additional information of the alarm.
   - If yes, go to **Step 4.2**.
   - If no, go to **Step 5**.
2. On the **Hosts** page, select the host of the node, and choose **More** > **Stop All Roles** to stop all the services on the node.

   If the time on the alarm node is later than that on the NTP service of the active OMS node, adjust the time of the alarm node. After adjusting the time, choose **More** > **Start All Roles** to start the services on the node.

   If the time on the alarm node is earlier than that on the NTP service of the active OMS node, wait until the time offset is due and adjust the time of the alarm node. After adjusting the time, choose **More** > **Start All Roles** to start the services on the node.

   NOTE

   If you do not wait, data loss may occur.
3. Wait 10 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.10 ALM-12014 Device Partition Lost (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the system detects that a partition to which service directories are mounted is lost (because the device is removed or goes offline, or the partition is deleted). The system checks the partition status periodically.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12014	Major	• Yes: MRS 1.9.3.10 and later patch versions • No: MRS 2.*x* and earlier versions

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DirName	Specifies the directory for which the alarm is generated.
PartitionName	Specifies the device partition for which the alarm is generated.

Impact on the System

Service data fails to be written into the partition, and the service system runs abnormally.

Possible Causes

- The disk is removed.
- The disk is offline, or a bad sector exists on the disk.

Procedure

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the real-time alarm list, click the row that contains the alarm.

Step 3 In the **Alarm Details** area, obtain the values of **HostName**, **PartitionName**, and **DirName** from **Location**.

Step 4 Check whether the disk corresponding to **PartitionName** on **HostName** is inserted to the correct server slot.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Contact hardware engineers to remove the faulty disk.

Step 6 Use PuTTY to log in to the **HostName** node where an alarm is reported and check whether there is a line containing **DirName** in the **/etc/fstab** file.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 Run the **vi /etc/fstab** command to edit the file and delete the line containing **DirName**.

Step 8 Contact hardware engineers to insert a new disk. For details, see the hardware product document of the relevant model. If the faulty disk is in a RAID group, configure the RAID group. For details, see the configuration methods of the relevant RAID controller card.

Step 9 Wait 20 to 30 minutes (The disk size determines the waiting time), and run the **mount** command to check whether the disk has been mounted to the **DirName** directory.

- If yes, perform **Step 10** for MRS 1.9.3.10 or later. For other versions, clear the alarm. No further action is required.
- If no, perform **Step 11**.

Step 10 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, perform **Step 11**.

Step 11 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

MRS 1.9.3.10 and later patch versions: After the fault is rectified, the system automatically clears the alarm.

MRS 2.*x* and earlier versions: After the fault is rectified, the system does not automatically clear the alarm. You need to clear the alarm.

Reference

None

9.5.11 ALM-12015 Device Partition File System Read-Only (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the system detects that a partition to which service directories are mounted enters the read-only mode (due to a bad sector or a faulty file system). The system checks the partition status periodically.

This alarm is cleared when the system detects that the partition to which service directories are mounted exits from the read-only mode (because the file system is restored to read/write mode, the device is removed, or the device is formatted).

Attribute

Alarm ID	Alarm Severity	Auto Clear
12015	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DirName	Specifies the directory for which the alarm is generated.
PartitionName	Specifies the device partition for which the alarm is generated.

Impact on the System

Service data fails to be written into the partition, and the service system runs abnormally.

Possible Causes

The disk is faulty, for example, a bad sector exists.

Procedure

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the real-time alarm list, click the row that contains the alarm.

Step 3 In the **Alarm Details** area, obtain **HostName** and **PartitionName** from **Location**. **HostName** indicates the node for which the alarm is generated, and **PartitionName** indicates the partition of the faulty disk.

Step 4 Contact hardware engineers to check whether the disk is faulty. If the disk is faulty, remove it from the server.

Step 5 After the disk is removed, the system reports ALM-12014 Partition Lost. Handle the alarm by following the instructions in **ALM-12014 Device Partition Lost (For MRS 2.x or Earlier)**. After the handling, the alarm is automatically cleared.

----End

Reference

None

9.5.12 ALM-12016 CPU Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the CPU usage every 30 seconds and compares the check result with the default threshold. The CPU usage has a default threshold. This alarm is generated when the CPU usage exceeds the threshold for several times (configurable, 10 times by default) consecutively.

This alarm is cleared when the average CPU usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12016	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- The alarm threshold or alarm hit number is improperly configured.
- The CPU configuration cannot meet service requirements. The CPU usage reaches the upper limit.

Procedure

Step 1 Check whether the alarm threshold or alarm hit number is properly configured.

1. Log in to MRS Manager and change the alarm threshold and alarm hit number based on CPU usage.
2. Choose **System** > **Threshold Configuration** > **Device** > **Host** > **CPU** > **CPU Usage** > **CPU Usage** and change the alarm threshold based on the actual CPU usage.
3. Choose **System** > **Threshold Configuration** > **Device** > **Host** > **CPU** > **CPU Usage** > **CPU Usage** and change **hit number** based on the actual CPU usage.

   NOTE

   This option defines the alarm check phase. **Interval** indicates the alarm check period and **hit number** indicates the number of times when the CPU usage exceeds the threshold. An alarm is generated when the CPU usage exceeds the threshold for several times consecutively.
4. Wait 2 minutes and check whether the alarm is automatically cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Expand the system.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the node.
2. Log in to the node for which the alarm is generated.
3. Run **cat /proc/stat | awk 'NR==1'|awk '{for(i=2;i<=NF;i++)j+=$i;print "" 100 - ($5+$6) * 100 / j;}'** to check the system CPU usage.
4. If the CPU usage exceeds the threshold, expand the CPU capacity.
5. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.13 ALM-12017 Insufficient Disk Capacity (For MRS 2.*x* or Earlier)

Description

The system checks the host disk usage every 30 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold. This alarm is generated if the disk usage exceeds the threshold.

To change the threshold, choose **System** > **Threshold Configuration**.

This alarm is cleared when the host disk usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12017	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition for which the alarm is generated.

Parameter	Description
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Service processes become unavailable.

Possible Causes

The disk configuration cannot meet service requirements. The disk usage reaches the upper limit.

Procedure

Step 1 Log in to MRS Manager and check whether the threshold is appropriate.

1. The default threshold is 90%. You can change the threshold to meet service requirements.
   - If yes, go to **Step 2**.
   - If no, go to **Step 1.2**.
2. Choose **System** > **Threshold Configuration** and change the alarm threshold based on the actual disk usage.
3. Wait 2 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Check whether the disk is a system disk.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host name and disk partition information.
2. Log in to the node for which the alarm is generated.
3. Run the **df -h** command to check the system disk partition usage. Check whether the disk is mounted to any of the following directories by using the disk partition name obtained in **Step 2.1**: **/**, **/boot**, **/home**, **/opt**, **/tmp**, **/var**, **/var/log**, **/boot**, and **/srv/BigData**.
   - If yes, the disk is a system disk. Then go to **Step 3.1**.
   - If no, the disk is not a system disk. Then go to **Step 2.4**.
4. Run the **df -h** command to check the system disk partition usage. Determine the role of the disk based on the disk partition name obtained in **Step 2.1**.
5. Check whether the disk is used by HDFS or Yarn.
   - If yes, expand the disk capacity for the Core node. Then go to **Step 2.6**.
   - If no, go to **Step 4**.
6. Wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Check whether large files are written to the disk.

1. Run the **find / -xdev -size +500M -exec ls -l {} \;** command to view files larger than 500 MB on the node. Check whether such files are written to the disk.
   - If yes, go to **Step 3.2**.
   - If no, go to **Step 4**.
2. Handle the large files and check whether the alarm is cleared 2 minutes later.
   - If yes, no further action is required.
   - If no, go to **Step 4**.
3. Expand the disk capacity.
4. Wait 2 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.14 ALM-12018 Memory Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the memory usage every 30 seconds and compares the actual memory usage with the threshold. The memory usage has a default threshold. This alarm is generated when the detected memory usage exceeds the threshold.

This alarm is cleared when the host memory usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12018	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

Memory configuration cannot meet service requirements. The memory usage reaches the upper threshold.

Procedure

Step 1 Expand the system.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the host address of the alarm.
2. Log in to the node for which the alarm is generated.
3. Run **free -m | grep Mem\: | awk '{printf("%s,", ($3-$6-$7) * 100 / $2)}'** to check the system memory usage.
4. If the memory usage exceeds the threshold, expand the memory capacity.
5. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.15 ALM-12027 Host PID Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the PID usage every 30 seconds and compares the actual PID usage with the default threshold. This alarm is generated when the PID usage exceeds the threshold.

This alarm is cleared when the host PID usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12027	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

No PID is available for new processes and service processes are unavailable.

Possible Causes

Too many processes are running on the node. You need to increase the value of **pid_max**. The system is abnormal.

Procedure

Step 1 Increase the value of **pid_max**.

1. On the MRS cluster details page, click the alarm from the real-time alarm list. In the **Alarm Details** area, obtain the IP address of the host for which the alarm is generated.

2. Log in to the node for which the alarm is generated.
3. Run the **cat /proc/sys/kernel/pid_max** command to check the value of **pid_max**.
4. If the PID usage exceeds the threshold, open the **/etc/sysctl.conf** file and change the value of **kernel.pid_max** to twice the value of **pid_max** queried in **Step 1.3**. If **kernel.pid_max** does not exist, add it to the end of the file.

   For example, change the parameter value to **kernel.pid_max=65536** and run the following command to make the parameter take effect immediately:

   sysctl -p

   NOTE

   The maximum value of **kernel.pid_max** is as follows:
   - 32-bit OS: **32768**
   - 64-bit OS: **4194304** (22nd power of 2)
5. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Check whether the system environment is abnormal.

1. Contact the O&M personnel to check whether the operating system is abnormal.
   - If yes, rectify the operating system fault and go to **Step 2.2**.
   - If no, go to **Step 3**.
2. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.16 ALM-12028 Number of Processes in the D State on the Host Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system periodically checks the number of D state processes of user **omm** on the host every 30 seconds and compares the number with the threshold. The number of processes in the D state on the host has a default threshold. This alarm is generated when the number of processes in the D state exceeds the threshold.

This alarm is cleared when the number is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12028	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Excessive system resources are used and the service process responds slowly.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and a process is in the D state.

Procedure

Step 1 Check the process that is in the D state.

1. Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the address of the host.
2. Log in to the node for which the alarm is generated.
3. Run the following commands to switch the user:

   sudo su - root

   su - omm
4. Run the following command as user **omm** to view the PID of the process that is in the D state:

   ps -elf | grep -v "\[thread_checkio\]" | awk 'NR!=1 {print $2, $3, $4}' | grep omm | awk -F' ' '{print $1, $3}' | grep D | awk '{print $2}'
5. Check whether the command output is empty.

- If yes, the service process is running properly. Then go to **Step 1.7**.
- If no, go to **Step 1.6**.

6. Switch to user **root** and run the **reboot** command to restart the alarm host.

   Restarting the host brings certain risks. Ensure that the service process runs properly after the restart.

7. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.17 ALM-12031 User omm or Password Is About to Expire (For MRS 2.*x* or Earlier)

Description

The system starts at 00:00 every day to check whether user **omm** and the password are about to expire every eight hours. This alarm is generated if the user or password is about to expire in 15 days.

The alarm is cleared when the validity period of user **omm** is changed or the password is reset and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12031	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The node trust relationship is unavailable and Manager cannot manage the services.

Possible Causes

User **omm** or the password is about to expire.

Procedure

Step 1 Check whether user **omm** and the password in the system are valid.

1. Log in to the faulty node.
2. Run the following command to view the information about user **omm** and the password:

   chage -l omm
3. Check whether the user has expired based on the system message.
   a. View the value of **Password expires** to check whether the password is about to expire.
   b. View the value of **Account expires** to check whether the user is about to expire.

   NOTE

   If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2**.
4. Run the following command to modify the validity period configuration:
   - Run the following command to set a validity period for user **omm**:

     chage -E *'specified date'* **omm**
   - Run the following command to set the number of validity days for user **omm**:

     chage -M *'number of days'* **omm**
5. Check whether the alarm is cleared automatically in the next periodic check.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.18 ALM-12032 User ommdba or Password Is About to Expire (For MRS 2.*x* or Earlier)

Description

The system starts at 00:00 every day to check whether user **ommdba** and the password are about to expire every eight hours. This alarm is generated if the user or password is about to expire in 15 days.

The alarm is cleared when the validity period of user **ommdba** is changed or the password is reset and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12032	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The OMS database cannot be managed and data cannot be accessed.

Possible Causes

User **ommdba** or the password is about to expire.

Procedure

Step 1 Check whether user **ommdba** and the password in the system are valid.

1. Log in to the faulty node.
2. Run the following command to view the information about user **ommdba** and the password:

chage -l ommdba

3. Check whether the user has expired based on the system message.
   a. View the value of **Password expires** to check whether the password is about to expire.
   b. View the value of **Account expires** to check whether the user is about to expire.

   NOTE

   If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2**.
4. Run the following command to modify the validity period configuration:
   - Run the following command to set a validity period for user **ommdba**:

     chage -E *'specified date'* **ommdba**
   - Run the following command to set the number of validity days for user **ommdba**:

     chage -M *'number of days'* **ommdba**
5. Check whether the alarm is cleared automatically in the next periodic check.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.19 ALM-12033 Slow Disk Fault (For MRS 2.*x* or Earlier)

Description

For MRS 2.*x* or earlier:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - The system runs the **iostat** command every 3 seconds, and detects that the **svctm** value exceeds 1000 ms for 10 consecutive periods within 30 seconds.
  - The system runs the **iostat** command every 3 seconds, and detects that more than 60% of I/O exceeds 150 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:

  - The system runs the **iostat** command every 3 seconds, and detects that the **svctm** value exceeds 1000 ms for 10 consecutive periods within 30 seconds.
  - The system runs the **iostat** command every 3 seconds, and detects that more than 60% of I/O exceeds 20 ms within 300 seconds.

This alarm is automatically cleared when the preceding conditions have not been met for 15 minutes.

For MRS 1.9.3.10 or later:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 150 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 20 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when none of the preceding conditions are met for three consecutive detection periods (30 or 300 seconds).

NOTE

For details about how to obtain the related parameters, see **Related Information**.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12033	• Minor: MRS 1.9.3.10 and later patch versions • Major: MRS 2.*x* and earlier versions	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
Host Name	Specifies the host for which the alarm is generated.
DiskName	Specifies the disk for which the alarm is generated.

Impact on the System

Service performance deteriorates, service processing capabilities become poor, and services may be unavailable.

Possible Causes

The disk is aged or has bad sectors.

Procedure

Check the disk status.

Step 1 On the MRS cluster details page, click the alarm from the real-time alarm list. In the **Alarm Details** area, obtain information about the host for which the alarm is generated and information about the faulty disk.

Step 2 Check whether the node for which the alarm is generated is in a virtualization environment.

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Step 3 Check whether the storage performance provided by the virtualization environment meets the hardware requirements. Then, go to **Step 4**.

Step 4 Log in to the alarm node as user **root**, run the **df -h** command, and check whether the command output contains the value of the **DiskName** field.

- If yes, go to **Step 6**.
- If no, go to **Step 5**.

Step 5 Run the **lsblk** command to check whether the mapping between the value of **DiskName** and the disk has been created.

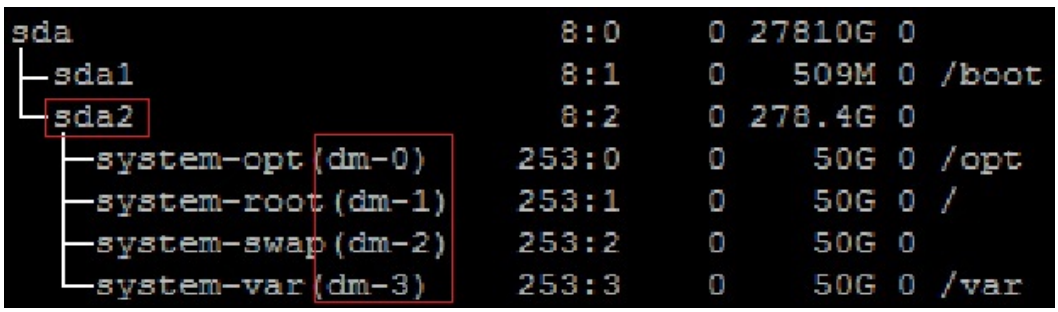

- If yes, go to **Step 6**.
- If no, go to **Step 21**.

Step 6 Log in to the alarm node as user **root**, run the **lsscsi | grep "/dev/sd[x]"** command to view the disk information, and check whether RAID has been set up.

NOTE

In the command, **/dev/sd[x]** indicates the disk name obtained in **Step 1**.

Example:

lsscsi | grep "/dev/sda"

In the command output, if **ATA**, **SATA**, or **SAS** is displayed in the third line, the disk has not been organized into a RAID group. If other information is displayed, RAID has been set up.

- If yes, go to **Step 11**.
- If no, go to **Step 7**.

Step 7 Run the **smartctl -i /dev/sd[x]** command to check whether the hardware supports the SMART tool.

Example:

smartctl -i /dev/sda

In the command output, if "SMART support is: Enabled" is displayed, the hardware supports SMART. If "Device does not support SMART" or other information is displayed, the hardware does not support SMART.

- If yes, go to **Step 8**.
- If no, go to **Step 16**.

Step 8 Run the **smartctl -H --all /dev/sd[x]** command to check basic SMART information and determine whether the disk is working properly.

Example:

smartctl -H --all /dev/sda

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or **Elements in grown defect list**. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to **Step 9**.
- If no, go to **Step 17**.

Step 9 Run the **smartctl -l error -H /dev/sd[x]** command to check the Glist of the disk and determine whether the disk is normal.

Example:

smartctl -l error -H /dev/sda

Check the **Command/Featrue_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other

errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can perform step 9 to trigger the disk SMART self-check.

- If yes, go to **Step 10**.
- If no, go to **Step 17**.

Step 10 Run the **smartctl -t long /dev/sd[x]** command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be completed is displayed. After the self-check is completed, repeat **Step 8** and **Step 9** to check whether the disk is working properly.

Example:

smartctl -t long /dev/sda

- If yes, go to **Step 16**.
- If no, go to **Step 17**.

Step 11 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]** command to check whether the hardware supports SMART.

NOTE

- In the command, **[sat|scsi]** indicates the disk type. Both types need to be used.
- **[DID]** indicates the slot information. Slots 0 to 15 need to be used.

For example, run the following commands in sequence:

smartctl -d sat+megaraid,0 -H --all /dev/sda

smartctl -d sat+megaraid,1 -H --all /dev/sda

smartctl -d sat+megaraid,2 -H --all /dev/sda

...

Try the command combinations of different disk types and slot information. If "SMART support is: Enabled" is displayed in the command output, the disk supports SMART. Record the parameters of the disk type and slot information when a command is successfully executed. If "SMART support is: Enabled" is not displayed in the command output, the disk does not support SMART.

- If yes, go to **Step 12**.
- If no, go to **Step 15**.

Step 12 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]** command recorded in **Step 11** to check basic SMART information and determine whether the disk is normal.

Example:

smartctl -d sat+megaraid,2 -H --all /dev/sda

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or

Elements in grown defect list. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to **Step 13**.
- If no, go to **Step 17**.

Step 13 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -l error -H /dev/sd[x]** command to check the Glist of the disk and determine whether the hard disk is working properly.

Example:

smartctl -d sat+megaraid,2 -l error -H /dev/sda

Check the **Command/Featrue_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can perform step 9 to trigger the disk SMART self-check.

- If yes, go to **Step 14**.
- If no, go to **Step 17**.

Step 14 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -t long /dev/sd[x]** command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be completed is displayed. After the self-check is completed, repeat **Step 12** and **Step 13** to check whether the disk is working properly.

Example:

smartctl -d sat+megaraid,2 -t long /dev/sda

- If yes, go to **Step 16**.
- If no, go to **Step 17**.

Step 15 If the configured RAID controller card does not support SMART, the disk does not support SMART. In this case, use the check tool provided by the corresponding RAID controller card vendor to rectify the fault. Then go to **Step 16**.

For example, LSI is a MegaCLI tool.

Step 16 On the alarm details page, click **Clear Alarm**. Check whether the alarm is reported on the same disk again.

If the alarm is reported for more than three times, replace the disk.

- If yes, go to **Step 17**.
- If no, no further action is required.

Replace the disk.

Step 17 On MRS Manager, choose **Alarms**.

Step 18 View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

Step 19 Replace a disk.

Step 20 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 21**.

Collect the fault information.

Step 21 On MRS Manager, choose **System** > **Export Log**.

Step 22 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

To obtain the related parameters, perform the following steps:

- For MRS 2.*x* or earlier versions:

  Perform the following operations to detect slow disk faults:

  On the Linux platform, run the **iostat -x -t 1** command to check whether the I/O is faulty. Specifically, check the **svctm** value in the red box in the figure below.

  svctm indicates the I/O service time of the disk.

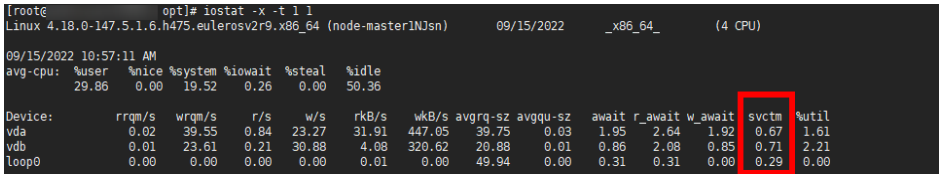

- For MRS 1.9.3.10 or later patch versions:

  The **svctm** value can be obtained through the following expression:

  svctm = (tot_ticks_new - tot_ticks_old) / (rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old)

  When the detection period is 30 seconds, if **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, then **svctm = 0**.

  When the detection period is 300 seconds and **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, if **tot_ticks_new - tot_ticks_old = 0**, then **svctm = 0**; otherwise, the value of **svctm** is infinite.

The parameters in the preceding expression can be obtained as follows:

Obtain the parameter values from the data collected via the **cat /proc/diskstats** command run by the system every 3 seconds. The following shows an example.

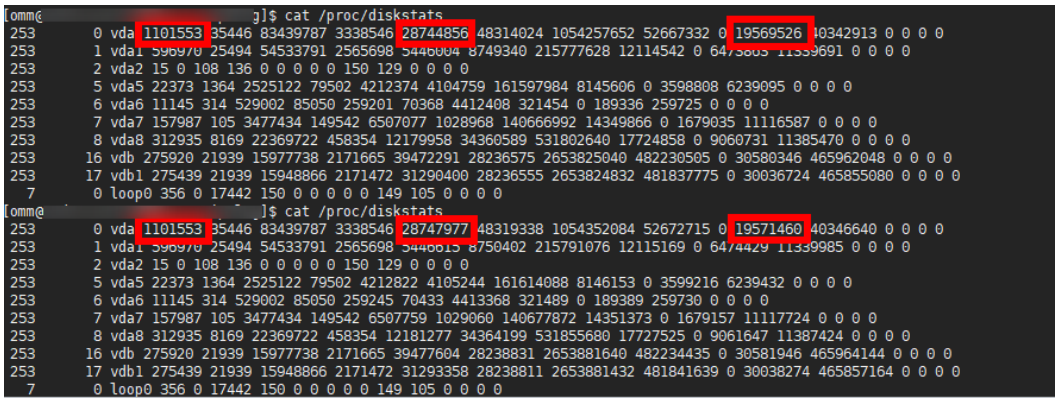

In the data collected for the first time, the number in the fourth column is the value of **rd_ios_old**, the number in the eighth column is the value of **wr_ios_old**, and the number in the thirteenth column is the value of **tot_ticks_old**.

In the data collected for the second time, the number in the fourth column is the value of **rd_ios_new**, the number in the eighth column is the value of **wr_ios_new**, and the number in the thirteenth column is the value of **tot_ticks_new**.

In this case, the value of **svctm** is as follows:

(19571460 - 19569526) / (1101553 + 28747977 - 1101553 - 28744856) = 0.6197

9.5.20 ALM-12034 Periodic Backup Failure (For MRS 2.*x* or Earlier)

Description

This alarm is generated when a periodic backup task fails to be executed. This alarm is cleared when the next backup task is executed successfully.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12034	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
TaskName	Specifies the task name.

Impact on the System

No backup package is available for a long time, so the system cannot be restored in case of exceptions.

Possible Causes

The alarm cause depends on the task details. Handle the alarm according to the logs and alarm details.

Procedure

Checking whether the disk space is insufficient

Step 1 On MRS Manager, choose **Alarms**.

Step 2 In the alarm list, click ⌄ of the alarm and obtain the task name from the **Location** area.

Step 3 Choose **System** > **Back Up Data**.

Step 4 Search for the backup task based on the task name and choose **More** > **View History** in the **Operation** column to view detailed information about the backup task.

Step 5 Choose **Details** > **View** and check whether message "Failed to backup xx due to insufficient disk space, move the data in the /srv/BigData/LocalBackup directory to other directories." exists.

- If yes, go to **Step 6**.
- If no, go to **Step 13**.

Step 6 Choose **Backup Path** > **View** to obtain the backup path.

Step 7 Log in to the node as user **root** and view the mounting details of the node.

df -h

Step 8 Check whether the available space of the node to which the backup path is mounted is less than 20 GB.

- If yes, go to **Step 9**.
- If no, go to **Step 13**.

Step 9 Check whether the backup package exists in the backup directory and whether the available space of the node to which the backup directory is mounted is less than 20 GB.

- If yes, go to **Step 10**.

- If no, go to **Step 13**.

Step 10 Ensure that the available space of the node to which the backup directory is mounted to be greater than 20 GB by moving backup packages out of the backup directory or deleting the backup packages.

Step 11 Start the backup task again and check whether the backup task is executed.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collecting fault information

Step 13 On MRS Manager, choose **System** > **Export Log**.

Step 14 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Reference

None

9.5.21 ALM-12035 Unknown Data Status After Recovery Task Failure (For MRS 2.*x* or Earlier)

Description

If a recovery task fails, the system attempts to automatically roll back. If the rollback fails, data may be lost. If this occurs, an alarm is reported. This alarm is cleared when the recovery task is successfully executed later.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12035	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TaskName	Specifies the task name.

Impact on the System

The data may be lost or the data status may be unknown, which may affect services.

Possible Causes

The possible cause of this alarm is that the component status does not meet the requirements before the restoration task is executed or an error occurs in a step during the restoration task. The error depends on the task details. You can obtain logs and task details to handle the alarm.

Procedure

Checking the component status

Step 1 Log in to MRS Manager and choose **Services**. On the page that is displayed, check whether the running status of the components meets the requirements. (OMS and DBService must be in the normal status, and other components must be stopped.)

- If yes, go to **Step 7**.
- If no, go to **Step 2**.

Step 2 Restore the component status as required and start the recovery task again.

Step 3 Log in to MRS Manager and choose **Alarms**. In the alarm list, click the row containing the alarm and obtain the task name from the **Location** area.

Step 4 Choose **System** > **Recovery Management**. Search for the restoration task based on the task name and view the task details.

Step 5 Start the restoration task and check whether the task is executed.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collecting fault information

Step 7 On MRS Manager, choose **System** > **Export Log**.

Step 8 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Reference

None

9.5.22 ALM-12037 NTP Server Abnormal (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the NTP server is abnormal.

This alarm is cleared when the NTP server recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12037	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the IP address of the NTP server for which the alarm is generated.

Impact on the System

The NTP server configured on the active OMS node is abnormal. In this case, the active OMS node cannot synchronize time with the NTP server and a time offset may be generated in the cluster.

Possible Causes

- The NTP server network is faulty.

- The NTP server authentication fails.
- The time cannot be obtained from the NTP server.
- The time obtained from the NTP server is not continuously updated.

Procedure

Step 1 Check the NTP server network.

1. On the MRS cluster details page, click the alarm from the real-time alarm list.
2. In the **Alarm Details** area, view the additional information to check whether the NTP server fails to be pinged.
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 2**.
3. Contact the O&M personnel to check the network configuration and ensure that the network between the NTP server and the active OMS node is in normal state. Then, check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Check whether the NTP server authentication fails.

1. Log in to the active management node.
2. Run **ntpq -np** to check whether the NTP server authentication fails. If **refid** of the NTP server is **.AUTH.**, the authentication fails.
   - If yes, go to **Step 5**.
   - If no, go to **Step 3**.

Step 3 Check whether the time can be obtained from the NTP server.

1. View the alarm additional information to check whether the time cannot be obtained from the NTP server.
   - If yes, go to **Step 3.2**.
   - If no, go to **Step 4**.
2. Contact the O&M personnel to rectify the NTP server fault. After the NTP server is in normal state, check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Check whether the time obtained from the NTP server fails to be updated.

1. View the alarm additional information to check whether the time obtained from the NTP server fails to be updated.
   - If yes, go to **Step 4.2**.
   - If no, go to **Step 5**.
2. Contact the provider of the NTP server to rectify the NTP server fault. After the NTP server is in normal state, check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.23 ALM-12038 Monitoring Indicator Dump Failure (For MRS 2.*x* or Earlier)

Description

This alarm is generated when dumping fails after monitoring indicator dumping is configured on MRS Manager.

This alarm is cleared when dumping is successful.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12038	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The upper-layer management system fails to obtain monitoring indicators from the MRS Manager system.

Possible Causes

- The server cannot be connected.
- The save path on the server cannot be accessed.
- The monitoring indicator file fails to be uploaded.

Procedure

Step 1 Contact the O&M personnel to check whether the network connection between the MRS Manager system and the server is normal.

- If yes, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 Contact the O&M personnel to restore the network and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Choose **System** > **Monitor Dumping Configuration** and check whether the FTP username, password, port, dump mode, and public key configured on the monitoring indicator dumping configuration page are consistent with those on the server.

- If yes, go to **Step 5**.
- If no, go to **Step 4**.

Step 4 Enter the correct configuration, click **OK**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Choose **System** > **Monitor Dumping Configuration** and check the configuration items, including the FTP username, save path, and dumping mode.

- If the FTP mode is used, go to **Step 6**.
- If the SFTP mode is used, go to **Step 7**.

Step 6 Log in to the server. In the default path, check whether the save path (relative path) has the read and write permission on the FTP username.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

Step 7 Log in to the server. In the default path, check whether the save path (absolute path) has the read and write permission on the FTP username.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

Step 8 Add the read and write permission and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Log in to the server and check whether the save path has sufficient disk space.

- If yes, go to **Step 11**.
- If no, go to **Step 10**.

Step 10 Delete unnecessary files or go to the monitoring indicator dumping configuration page to change the save path. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Step 11 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.24 ALM-12039 GaussDB Data Is Not Synchronized (For MRS 2.*x* or Earlier)

Description

The system checks the data synchronization status between the active and standby GaussDB nodes every 10 seconds. This alarm is generated when the synchronization status cannot be queried for six consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the data synchronization status is normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12039	Critical	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.
Local GaussDB HA IP	HA IP address of the local GaussDB.
Peer GaussDB HA IP	HA IP address of the peer GaussDB.
SYNC_PERSENT	Synchronization percentage.

Impact on the System

When data is not synchronized between the active and standby GaussDBs, the data may be lost or abnormal if the active instance becomes abnormal.

Possible Causes

- The network between the active and standby nodes is unstable.
- The standby GaussDB is abnormal.
- The disk space of the standby node is full.

Procedure

Step 1 Go to the MRS cluster details page. In the alarm list on the alarm management tab page, click the row that contains the alarm. In the alarm details, view the IP address of the standby GaussDB node.

Step 2 Log in to the active management node.

Step 3 Run the following command to check whether the standby GaussDB is reachable:

ping *heartbeat IP address of the standby GaussDB*

If yes, go to **Step 6**.

If no, go to **Step 4**.

Step 4 Contact the O&M personnel to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 6**.

Step 6 Log in to the standby GaussDB node.

Step 7 Run the following commands to switch the user:

sudo su - root

su - omm

Step 8 Go to the **${BIGDATA_HOME}/om-0.0.1/sbin/** directory.

Run the following command to check whether the resource status of the standby GaussDB is normal:

sh status-oms.sh

In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:

```
10_10_10_231 gaussDB Standby_normal Normal Active_standby
```

- If yes, go to **Step 9**.
- If no, go to **Step 15**.

Step 9 Log in to the standby GaussDB node.

Step 10 Run the following commands to switch the user:

sudo su - root

su - omm

Step 11 Run the **echo ${BIGDATA_DATA_HOME}/dbdata_om** command to obtain the GaussDB data directory.

Step 12 Run the **df -h** command to check the system disk partition usage.

Step 13 Check whether the disk where the GaussDB data directory is mounted is full.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Contact the O&M personnel to expand the disk capacity. After capacity expansion, wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Step 15 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.25 ALM-12040 Insufficient System Entropy (For MRS 2.*x* or Earlier)

Description

The system checks the entropy at 00:00:00 every day and performs five consecutive checks each time. First, the system checks whether the rng-tools tool is enabled and correctly configured. If not, the system checks the current entropy. This alarm is generated if the entropy is less than 500 in the five checks.

This alarm is cleared if the true random number mode is configured, random numbers are configured in pseudo-random number mode, or neither the true random number mode nor the pseudo-random number mode is configured but the entropy is greater than or equal to 500 in at least one check among the five checks.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12040	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Decryption failures occur and functions related to decryption are affected, for example, DBService installation.

Possible Causes

The rngd service is abnormal.

Procedure

Step 1 Go to the cluster details page and choose **Alarms**.

Step 2 View the alarm details to obtain the value of the **HostName** field in **Location**.

Step 3 Log in to the node for which the alarm is generated and run the **sudo su - root** command to switch to user **root**.

Step 4 Run the **/bin/rpm -qa | grep -w "rng-tools"** command. If the command is executed successfully, run the **ps -ef | grep -v "grep" | grep rngd | tr -d " " | grep "\-o/dev/random" | grep "\-r/dev/urandom"** command and view the command output.

- If the command is executed successfully, the rngd service is installed, correctly configured, and is running properly. Go to **Step 8**.
- If the command is not executed successfully, the rngd service is not running properly. Then go to **Step 5**.

Step 5 Run the following command to start the rngd service:

echo 'EXTRAOPTIONS="-r /dev/urandom -o /dev/random"' >> /etc/sysconfig/rngd

service rngd start

Step 6 Run the **service rngd status** command to check whether the rngd service is in the running state.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 Wait until 00:00:00 when the system checks the entropy again. Check whether the alarm is cleared automatically.

- If yes, no further action is required.
- If no, go to **Step 8**.

Step 8 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.26 ALM-12041 Permission of Key Files Is Abnormal (For MRS 2.*x* or Earlier)

Description

The system checks the permission, users, and user groups of key directories or files every hour. This alarm is generated if any of these is abnormal.

This alarm is cleared after the problem that causes abnormal permission, users, or user groups is solved.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12041	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PathName	Specifies the file path or file name.

Impact on the System

System functions are unavailable.

Possible Causes

The user has manually modified the file permission, user information, or user groups, or the system has experienced an unexpected power-off.

Procedure

Step 1 Check the file permission.

1. Go to the MRS cluster details page and choose **Alarms**.
2. In the details of the alarm, query the **HostName** (name of the alarmed host) and **PathName** (path or name of the involved file).
3. Log in to the alarmed node.
4. Run the **ll** *PathName* command to query the current user, permission, and user group of the file or path.
5. Go to the **${BIGDATA_HOME}/nodeagent/etc/agent/autocheck** directory and run the **vi keyfile** command. Search for the name of the involved file and query the correct permission of the file.
6. Compare the actual permission of the file with the permission obtained in **Step 1.5**. If they are different, change the actual permission, user information, and user group to the correct values.
7. Wait until the next system check is complete and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.27 ALM-12042 Key File Configurations Are Abnormal (For MRS 2.*x* or Earlier)

Description

The system checks key file configurations every hour. This alarm is generated if any key configuration is abnormal.

This alarm is cleared after the configuration becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12042	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PathName	Specifies the file path or file name.

Impact on the System

Functions related to the file are abnormal.

Possible Causes

The user has manually modified the file configurations or the system has experienced an unexpected power-off.

Procedure

Step 1 Check the file configurations.

1. Go to the MRS cluster details page and choose **Alarms**.
2. In the details of the alarm, query the **HostName** (name of the alarmed host) and **PathName** (path or name of the involved file).
3. Log in to the alarmed node.
4. Manually check and modify the file configurations according to the criteria in **Related Information**.
5. Wait until the next system check is complete and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Related Information

- **Checking /etc/fstab**

  Check whether partitions configured in **/etc/fstab** exist in **/proc/mounts** and whether swap partitions configured in **/etc/fstab** match those in **/proc/swaps**.

- **Checking /etc/hosts**

  Run the **cat /etc/hosts** command. If any of the following situations exists, the file configurations are abnormal.

  - The **/etc/hosts** file does not exist.
  - The host name is not configured in the file.
  - The IP address of the host is duplicate.
  - The IP address of the host does not exist in the **ipconfig** list.
  - An IP address in the file is used by multiple hosts.

9.5.28 ALM-12043 DNS Parsing Duration Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the DNS parsing duration every 30 seconds. This alarm is generated when the DNS parsing duration exceeds the threshold (the default threshold is 20,000 ms) for multiple times (the default value is **2**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Network Status** > **DNS Resolution Duration** > **DNS Resolution Duration**.

This alarm is cleared when **hit number** is **1** and the DNS resolution duration is less than or equal to the threshold. This alarm is cleared when **hit number** is not **1** and the DNS resolution duration is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12043	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- Kerberos-based secondary authentication is slow.
- The ZooKeeper service is abnormal.
- The node is faulty.

Possible Causes

- The node is configured with the DNS client.
- The node is equipped with the DNS server and the DNS server is started.

Procedure

Check whether the node is configured with the DNS client.

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 View the alarm details to obtain the value of the **HostName** field in **Location**.

Step 3 Use PuTTY to log in to the node for which the alarm is generated as user **root**.

Step 4 Run the **cat /etc/resolv.conf** command to check whether the DNS client is installed.

If information similar to the following is displayed, the DNS client is installed and started:

```
namesever 10.2.3.4
namesever 10.2.3.4
```

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Run the **vi /etc/resolv.conf** command to comment out the following content using the number signs (#) and save the file:

```
# namesever 10.2.3.4
# namesever 10.2.3.4
```

Step 6 Check whether this alarm is cleared after 5 minutes.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check whether the node is equipped with the DNS server and the DNS server is started.

Step 7 Run the **service named status** command to check whether the DNS service is installed on the node.

If information similar to the following is displayed, the DNS server is installed and started:

```
Checking for nameserver BIND
version: 9.6-ESV-R7-P4
CPUs found: 8
worker threads: 8
number of zones: 17
debug level: 0
xfers running: 0
xfers deferred: 0
soa queries in progress: 0
query logging is ON
recursive clients: 4/0/1000
tcp clients: 0/100
server is up and running
```

- If yes, go to **Step 8**.
- If no, go to **Step 10**.

Step 8 Run the **service named stop** command to stop the DNS server.

Step 9 Check whether this alarm is cleared after 5 minutes.

- If yes, no further action is required.
- If no, go to **Step 10**.

Step 10 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.29 ALM-12045 Read Packet Dropped Rate Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the read packet dropped rate every 30 seconds. This alarm is generated when the read packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Network Reading** > **Network Read Packet Rate Information** > **Read Packet Dropped Rate**.

This alarm is cleared when **hit number** is 1 and the read packet dropped rate is less than or equal to the threshold. This alarm is cleared when **hit number** is greater than 1 and the read packet dropped rate is less than or equal to 90% of the threshold.

The alarm detection is disabled by default. If you want to enable this function, check whether this function can be enabled based on Checking System Environments.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12045	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The service performance deteriorates or some services time out.

Risk warning: In SUSE kernel 3.0 or later or Red Hat 7.2, the system kernel modifies the mechanism for counting the number of dropped read packets. In this case, this alarm may be generated even if the network is running properly, but services are not affected. You are advised to check the system environment first.

Possible Causes

- An OS exception occurs.
- The NICs are bonded in active/standby mode.
- The alarm threshold is improperly configured.
- The network environment is abnormal.

Procedure

View the network packet dropped rate.

Step 1 Use PuTTY to log in to any non-alarm node in the cluster as user **omm** and run the **ping** *IP address of the node for which the alarm is generated* **-c 100** command to check whether packet drop occurs on the network.

```
# ping 10.10.10.12 -c 5
PING 10.10.10.12 (10.10.10.12) 56(84) bytes of data.
64 bytes from 10.10.10.11: icmp_seq=1 ttl=64 time=0.033 ms
64 bytes from 10.10.10.11: icmp_seq=2 ttl=64 time=0.034 ms
```



```
64 bytes from 10.10.10.11: icmp_seq=3 ttl=64 time=0.021 ms
64 bytes from 10.10.10.11: icmp_seq=4 ttl=64 time=0.033 ms
64 bytes from 10.10.10.11: icmp_seq=5 ttl=64 time=0.030 ms
--- 10.10.10.12 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4001ms   rtt min/avg/max/mdev =
0.021/0.030/0.034/0.006 ms
```

NOTE

- *IP address of the node for which the alarm is generated*: Query the IP address of the node for which the alarm is generated on the node management page of the MRS cluster details page based on the value of **HostName** in the alarm location information. Check both the IP addresses of the management plane and service plane.
- **-c**: number of check times. The default value is **100**.

- If yes, go to **Step 11**.
- If no, go to **Step 2**.

Check the system environment.

Step 2 Use PuTTY to log in to the active OMS node or the node for which the alarm is generated as user **omm**.

Step 3 Run the **cat /etc/*-release** command to check the OS type.

- If the OS is EulerOS, go to **Step 4**.

```
# cat  /etc/*-release                          EulerOS release 2.0 (SP2)
EulerOS release 2.0 (SP2)
```

- If the OS is SUSE, go to **Step 5**.

```
# cat /etc/*-release
SUSE Linux Enterprise Server 11 (x86_64)
VERSION = 11
PATCHLEVEL = 3
```

- Otherwise, go to **Step 11**.

Step 4 Run the **cat /etc/euleros-release** command to check whether the OS version is EulerOS 2.2.

```
# cat/etc/euleros-release
EulerOS release 2.0 (SP2)
```

- If yes, the alarm sending function cannot be enabled. Go to **Step 6**.
- If no, go to **Step 11**.

Step 5 Run the **cat /proc/version** command to check whether the SUSE kernel version is 3.0 or later.

```
# cat /proc/version
Linux version 3.0.101-63-default (geeko@buildhost) (gcc version 4.3.4 [gcc-4_3-branch revision 152973]
(SUSE Linux) ) #1 SMP Tue Jun 23 16:02:31 UTC 2015 (4b89d0c)
```

- If yes, the alarm sending function cannot be enabled. Go to **Step 6**.
- If no, go to **Step 11**.

Step 6 Log in to MRS Manager and choose **System** > **Configuration** > **Threshold Configuration**.

Step 7 In the navigation pane of the **Threshold Configuration** page, choose **Network Reading** > **Network Read Packet Rate Information** > **Read Packet Dropped Rate**. In the right pane, check whether **Send Alarm** is selected.

- If yes, the alarm sending function is enabled. Go to **Step 8**.
- If no, the alarm sending function is disabled. Go to **Step 10**.

Step 8 In the right pane, deselect **Send Alarm** to shield alarm "Network Read Packet Dropped Rate Exceeds the Threshold."

Step 9 Go to the MRS cluster details page and choose **Alarms**.

Step 10 Search for alarm 12045 and manually clear the alarms that are not automatically cleared. No further action is required.

NOTE

The ID of alarm Network Read Packet Dropped Rate Exceeds the Threshold is 12045.

Check whether the NICs are bonded in active/standby mode.

Step 11 Use PuTTY to log in to the node for which the alarm is generated as user **omm** and run the **ls -l /proc/net/bonding** command to check whether the **/proc/net/bonding** directory exists on the node.

- If yes, as shown in the following figure, the bond mode is configured for the node. Go to **Step 12**.

```
# ls -l /proc/net/bonding/
total 0
-r--r--r-- 1 root root 0 Oct 11 17:35 bond0
```

- If no, the bond mode is not configured for the node. Go to **Step 14**.

```
# ls -l /proc/net/bonding/
ls: cannot access /proc/net/bonding/: No such file or directory
```

Step 12 Run the **cat /proc/net/bonding/bond0** command to check whether the value of **Bonding Mode** in the configuration file is **fault-tolerance**.

NOTE

In the preceding command, **bond0** is the name of the bond configuration file. Use the file name obtained in **Step 11**.

```
# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)

Bonding Mode: fault-tolerance (active-backup)
Primary Slave: eth1 (primary_reselect always)
Currently Active Slave: eth1
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0

Slave Interface: eth0
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0

Slave Interface: eth1
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0
```

- If yes, the NICs are bonded in active/standby mode. Go to **Step 13**.
- If no, go to **Step 14**.

Step 13 Check whether the NIC specified by **NetworkCardName** in the alarm details is the standby NIC.

- If yes, the alarm of the standby NIC cannot be automatically cleared. Manually clear the alarm on the alarm management page. No further action is required.
- If no, go to **Step 14**.

  NOTE

  To determine the standby NIC, check the **/proc/net/bonding/bond0** configuration file. If the NIC name corresponding to **NetworkCardName** is **Slave Interface** but not **Currently Active Slave** (the current active NIC), the NIC is the standby one.

Check whether the threshold is set properly.

Step 14 Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.

- If yes, go to **Step 17**.
- If no, go to **Step 15**.

Step 15 Choose **System** > **Threshold Configuration** > **Device** > **Host** > **Network Reading** > **Network Read Packet Rate Information** > **Read Packet Dropped Rate** and change the alarm threshold based on the actual service usage.

Step 16 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 17**.

Check whether the network is normal.

Step 17 Contact the system administrator to check whether the network is normal.

- If yes, rectify the network fault and go to **Step 18**.
- If no, go to **Step 19**.

Step 18 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 19**.

Step 19 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.30 ALM-12046 Write Packet Dropped Rate Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the write packet dropped rate every 30 seconds. This alarm is generated when the write packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Network Writing** > **Network Write Packet Rate Information** > **Write Packet Dropped Rate**.

When the **hit number** is **1**, this alarm is cleared when the network write packet dropped rate is less than or equal to the threshold. When the **hit number** is greater than **1**, this alarm is cleared when the network write packet dropped rate is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12046	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The service performance deteriorates or some services time out.

Possible Causes

- The alarm threshold is improperly configured.

- The network environment is abnormal.

Procedure

Check whether the threshold is set properly.

Step 1 Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **System** > **Threshold Configuration** > **Device** > **Host** > **Network Write Information** > **Network Write Packet Rate** > **Write Packet Dropped Rate** and change the alarm threshold based on the actual service usage.

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network is normal.

Step 4 Contact the system administrator to check whether the network is normal.

- If yes, rectify the network fault and go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Step 6 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.31 ALM-12047 Read Packet Error Rate Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the read packet error rate every 30 seconds. This alarm is generated when the read packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Network Reading** > **Network Read Packet Rate Information** > **Read Packet Error Rate**.

If the **hit number** is **1**, this alarm is cleared when the read packet error rate is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the read packet error rate is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12047	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The communication is intermittently interrupted, and services time out.

Possible Causes

- The alarm threshold is improperly configured.
- The network environment is abnormal.

Procedure

Check whether the threshold is set properly.

Step 1 Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **System** > **Threshold Configuration** > **Device** > **Host** > **Network Reading** > **Network Read Packet Rate Information** > **Read Packet Error Rate** and change the alarm threshold based on the actual service usage.

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network is normal.

Step 4 Contact the system administrator to check whether the network is normal.

- If yes, rectify the network fault and go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Step 6 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.32 ALM-12048 Write Packet Error Rate Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the write packet error rate every 30 seconds. This alarm is generated when the write packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Network Writing** > **Network Write Packet Rate Information** > **Write Packet Error Rate**.

If **hit number** is **1**, this alarm is cleared when the write packet error rate is less than or equal to the threshold. If **hit number** is greater than **1**, this alarm is cleared when the write packet error rate is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12048	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The communication is intermittently interrupted, and services time out.

Possible Causes

- The alarm threshold is improperly configured.
- The network environment is abnormal.

Procedure

Check whether the threshold is set properly.

Step 1 Log in to MRS Manager and check whether the threshold (configurable, 0.5% by default) is appropriate.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **System** > **Threshold Configuration** > **Device** > **Host** > **Network Writing** > **Network Write Packet Rate Information** > **Write Packet Error Rate** and change the alarm threshold based on the actual service usage.

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network is normal.

Step 4 Contact the system administrator to check whether the network is normal.

- If yes, rectify the network fault and go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Step 6 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.33 ALM-12049 Read Throughput Rate Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the read throughput rate every 30 seconds. This alarm is generated when the read throughput rate exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Network Reading** > **Network Read Throughput Rate** > **Read Throughput Rate**.

If the **hit number** is **1**, this alarm is cleared when the read throughput rate is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the read throughput rate is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12049	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The service system runs abnormally or is unavailable.

Possible Causes

- The alarm threshold is improperly configured.
- The network port rate does not meet service requirements.

Procedure

Check whether the threshold is set properly.

Step 1 Log in to MRS Manager and check whether the threshold (configurable, 80% by default) is appropriate.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Choose **System** > **Threshold Configuration** > **Device** > **Host** > **Network Reading** > **Network Read Throughput Rate** > **Read Throughput Rate** to change the alarm threshold based on the actual service usage.

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network port rate meets the requirements.

Step 4 In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address and network port name of the host for which the alarm is generated.

Step 5 Use PuTTY to log in to the host for which the alarm is generated as user **root**.

Step 6 Run the **ethtool** *network port name* command to check the maximum network port rate **Speed**.

NOTE

In a VM environment, you may fail to obtain the network port rate by running commands. You are advised to contact the system administrator to check whether the network port rate meets the requirements.

Step 7 If the read throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.34 ALM-12050 Write Throughput Rate Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the write throughput rate every 30 seconds. This alarm is generated when the write throughput rate exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Network Writing** > **Network Write Throughput Rate** > **Write Throughput Rate**.

If the **hit number** is **1**, this alarm is cleared when the write throughput rate is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the write throughput rate is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12050	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The service system runs abnormally or is unavailable.

Possible Causes

- The alarm threshold is improperly configured.
- The network port rate does not meet service requirements.

Procedure

Check whether the threshold is set properly.

Step 1 Log in to MRS Manager and check whether the threshold (configurable, 80% by default) is appropriate.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **System** > **Threshold Configuration** > **Device** > **Host** > **Network Writing** > **Network Write Throughput Rate** > **Write Throughput Rate** to change the alarm threshold based on the actual service usage.

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network port rate meets the requirements.

Step 4 In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address and network port of the host for which the alarm is generated.

Step 5 Use PuTTY to log in to the host for which the alarm is generated as user **root**.

Step 6 Run the **ethtool** *network port name* command to check the maximum network port rate **Speed**.

NOTE

In a VM environment, you may fail to obtain the network port rate by running commands. You are advised to contact the system administrator to check whether the network port rate meets the requirements.

Step 7 If the write throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.35 ALM-12051 Disk Inode Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the disk inode usage every 30 seconds. This alarm is generated when the disk inode usage exceeds the threshold (the default threshold is 80%) for multiple times (the default value is 5).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Disk** > **Disk Inode Usage** > **Disk Inode Usage**.

If the **hit number** is **1**, this alarm is cleared when the disk inode usage is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the disk inode usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12051	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition for which the alarm is generated.

Parameter	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Data cannot be written to the file system.

Possible Causes

- There are too many small files on the disk.
- The system is abnormal.

Procedure

There are too many small files on the disk.

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address and disk partitions of the host for which the alarm is generated.

Step 3 Use PuTTY to log in to the host for which the alarm is generated as user **root**.

Step 4 Run the **df -i** *partition name* command to check the current inode usage of the disk.

Step 5 If the inode usage exceeds the threshold, manually check whether the small files in the partition can be deleted.

- If yes, delete the files and go to **Step 6**.
- If no, adjust the capacity. Then go to **Step 7**.

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check whether the system environment is normal.

Step 7 Contact the operating system maintenance personnel to check whether the system environment is abnormal.

- If yes, rectify the operating system fault and go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.36 ALM-12052 Usage of Temporary TCP Ports Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the usage of temporary TCP ports every 30 seconds. This alarm is generated when the usage of temporary TCP ports exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Host** > **Network Status** > **TCP Ephemeral Port Usage** > **TCP Ephemeral Port Usage**.

If the **hit number** is **1**, this alarm is cleared when the usage of temporary TCP ports is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the usage of temporary TCP ports is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12052	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Services on the host fail to establish connections with the external and services are interrupted.

Possible Causes

- The temporary ports do not meet service requirements.
- The system is abnormal.

Procedure

Expand the range of temporary ports.

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address of the host for which the alarm is generated.

Step 3 Use PuTTY to log in to the host for which the alarm is generated as user **omm**.

Step 4 Run the **cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 1** command to obtain the start port number. Run the **cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 2** command to obtain the end port number. Subtract the start port number from the end port number to obtain the total number of temporary ports. If the total number of temporary ports is less than 28,232, the random port range of the OS is too small. In this case, contact the system administrator to expand the port range.

Step 5 Run the **ss -ant 2>/dev/null | grep -v LISTEN | awk 'NR > 2 {print $4}'|cut -d ':' -f 2 | awk '$1 >"***start port number***" {print $1}' | sort -u | wc -l** command to calculate the number of used temporary ports.

Step 6 Calculate the usage of temporary ports using the following formula: Usage of temporary ports = (Number of used temporary ports/Total number of temporary ports) x 100. Check whether the usage exceeds the threshold.

- If yes, go to **Step 8**.
- If no, go to **Step 7**.

Step 7 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check whether the system environment is normal.

Step 8 Run the following command to import the temporary file and view the frequently used ports in the **port_result.txt** file:

netstat -tnp > $BIGDATA_HOME/tmp/port_result.txt

```
netstat -tnp

Active Internet connections (w/o servers)

Proto Recv Send LocalAddress ForeignAddress State PID/ProgramName tcp   0   0 10-120-85-154:45433
10-120-8:25009 CLOSE_WAIT 94237/java
tcp   0   0 10-120-85-154:45434  10-120-8:25009 CLOSE_WAIT 94237/java
tcp   0   0 10-120-85-154:45435  10-120-8:25009 CLOSE_WAIT 94237/java
...
```


Step 9 Run the following command to check the processes that occupy a large number of ports:

ps -ef |grep *PID*

NOTE

- *PID* indicates the process ID of the port queried in **Step 8**.
- Run the following command to collect information about all processes in the system and check the processes that occupy a large number of ports:

  ps -ef > $BIGDATA_HOME/tmp/ps_result.txt

Step 10 Contact the system administrator to clear the processes that occupy a large number of ports. Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Step 11 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.37 ALM-12053 File Handle Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the handler usage every 30 seconds. This alarm is generated when the handle usage exceeds the threshold (the default threshold is **80%**) for multiple times (the default value is **5**).

You can change the threshold by choosing **System** > **Threshold Configuration** > **Device** > **Host** > **Host Status** > **Host File Handle Usage** > **Host File Handle Usage**.

If the **hit number** is **1**, this alarm is cleared when the file handle usage is less than or equal to the threshold. If the **hit number** is greater than **1**, this alarm is cleared when the file handle usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12053	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The system applications fail to open files, access networks, and perform other I/O operations. The applications are running improperly.

Possible Causes

- The number of file handles does not meet service requirements.
- The system is abnormal.

Procedure

Increase the number of file handles.

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the real-time alarm list, click the alarm. In the **Alarm Details** area, obtain the IP address of the host for which the alarm is generated.

Step 3 Use PuTTY to log in to the host for which the alarm is generated as user **root**.

Step 4 Run the **ulimit -n** command to check the maximum number of handles set in the system.

Step 5 If the file handle usage exceeds the threshold, contact the system administrator to increase the number of system file handles.

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check whether the system environment is normal.

Step 7 Contact the system administrator to check whether the OS is abnormal.

- If yes, rectify the operating system fault and go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.38 ALM-12054 Invalid Certificate File (For MRS 2.*x* or Earlier)

Description

The system checks whether the certificate file is invalid (has expired or is not yet valid) on 23:00 every day. This alarm is generated when the certificate file is invalid.

This alarm is cleared when the status of the newly imported certificate is valid.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12054	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system reminds users that the certificate file is invalid. If the certificate file is invalid, some functions are restricted and cannot be used properly.

Possible Causes

No certificate (HA root certificate or HA user certificate) is imported to the system, the certificate fails to be imported, or the certificate file is invalid.

Procedure

Check the alarm cause.

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the real-time alarm list, click the row that contains the alarm.

In the **Alarm Details** area, view the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, use PuTTY to log in to the active OMS management node as user **omm** and go to **Step 3**.
- If **HA root Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to **Step 4**.
- If **HA server Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to **Step 5**.

Check the validity period of the certificate files in the system.

Step 3 Check whether the current system time is in the validity period of the CA certificate.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/cert/root/ca.crt** command to check the effective time and due time of the root certificate.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 4 Check whether the current system time is in the validity period of the HA root certificate.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/root-ca.crt** command to check the effective time and due time of the HA root certificate.

- If yes, go to **Step 8**.
- If no, go to **Step 7**.

Step 5 Check whether the current system time is in the validity period of the HA user certificate.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/server.crt** command to check the effective time and due time of the HA user certificate.

- If yes, go to **Step 8**.
- If no, go to **Step 7**.

The following is an example of the effective time and expiration time of a CA or HA certificate:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            97:d5:0e:84:af:ec:34:d8
        Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=CountryName, ST=State, L=Locality, O=Organization, OU=IT, CN=HADOOP.COM
        Validity
            Not Before: Dec 13 06:38:26 2016 GMT          // Effective time
            Not After : Dec 11 06:38:26 2026 GMT          // Expiration time
```

Import certificate files.

Step 6 Import a new CA certificate file.

Contact O&M personnel to apply for or generate a new CA certificate file and import it. Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to **Step 8**.
- If no, no further action is required.

Step 7 Import a new HA certificate file.

Apply for or generate a new HA certificate file and import it by referring to **Replacing the HA Certificate**. Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to **Step 8**.
- If no, no further action is required.

Step 8 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.39 ALM-12055 Certificate File Is About to Expire (For MRS 2.*x* or Earlier)

Description

The system checks the certificate file on 23:00 every day. This alarm is generated if the certificate file is about to expire with a validity period less than days set in the alarm threshold.

This alarm is generated if the status of the newly imported certificate is valid.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12055	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system reminds users that the license is about to expire. If the license expires, some functions are restricted and cannot be used properly.

Possible Causes

The remaining validity period of the CA certificate, HA root certificate, or HA user certificate is smaller than the alarm threshold.

Procedure

Check the alarm cause.

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the real-time alarm list, click the row that contains the alarm.

In the **Alarm Details** area, view the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, use PuTTY to log in to the active OMS management node as user **omm** and go to **Step 3**.
- If **HA root Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to **Step 4**.
- If **HA server Certificate** is displayed in the additional information, check **Location** to obtain the name of the host involved in this alarm. Then use PuTTY to log in to the host as user **omm** and go to **Step 5**.

Check the validity period of the certificate files in the system.

Step 3 Check whether the remaining validity period of the CA certificate is smaller than the alarm threshold.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/cert/root/ca.crt** command to check the effective time and due time of the root certificate.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 4 Check whether the remaining validity period of the HA root certificate is smaller than the alarm threshold.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/root-ca.crt** command to check the effective time and due time of the HA root certificate.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 5 Check whether the remaining validity period of the HA user certificate is smaller than the alarm threshold.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/server.crt** command to check the effective time and due time of the HA user certificate.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

  The following is an example of the effective time and expiration time of a CA or HA certificate:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            97:d5:0e:84:af:ec:34:d8
        Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=CountryName, ST=State, L=Locality, O=Organization, OU=IT, CN=HADOOP.COM
        Validity
            Not Before: Dec 13 06:38:26 2016 GMT            // Effective time
            Not After : Dec 11 06:38:26 2026 GMT            // Expiration time
```

Import certificate files.

Step 6 Import a new CA certificate file.

Contact O&M personnel to apply for or generate a new CA certificate file and import it. Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to **Step 8**.
- If no, no further action is required.

Step 7 Import a new HA certificate file.

Apply for or generate a new HA certificate file and import it by referring to **Replacing the HA Certificate**. Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to **Step 8**.
- If no, no further action is required.

Step 8 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.40 ALM-12180 Disk Card I/O (For MRS 2.*x* or Earlier)

Description

For MRS 2.*x* or earlier:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - The system collects data every 3 seconds, and detects that the **svctm** value exceeds 6s for 10 consecutive periods within 30 seconds.
  - The system collects data every 3 seconds, and detects that the **avgqu-sz** value is greater than 0, the IOPS or bandwidth is 0, and the **ioutil** value is greater than **99%** for 10 consecutive periods within 30 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - The system collects data every 3 seconds, and detects that the **svctm** value exceeds 2s for 10 consecutive periods within 30 seconds.
  - The system collects data every 3 seconds, and detects that the **avgqu-sz** value is greater than 0, the IOPS or bandwidth is 0, and the **ioutil** value is greater than **99%** for 10 consecutive periods within 30 seconds.

This alarm is automatically cleared when none of the conditions are met for 90 seconds.

For MRS 1.9.3.10 or later:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 6 seconds within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. Disk queue depth (**avgqu-sz**) > 0 and IOPS = 0, or bandwidth = 0 and **ioutil** > 99% in at least 10 collection periods within 30 seconds.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 1000 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 3 seconds within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. Disk queue depth (**avgqu-sz**) > 0 and IOPS = 0, or bandwidth = 0 and **ioutil** > 99% in at least 10 collection periods within 30 seconds.

- By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 500 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when none of the conditions are met for three consecutive detection periods (30 or 300 seconds).

NOTE

For details about how to obtain related parameters, see **Related Information**.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12180	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DiskName	Specifies the disk for which the alarm is generated.

Impact on the System

A continuously high I/O usage may adversely affect service operations and result in service loss.

Possible Causes

The disk is aged.

Procedure

Replace the disk.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**.

Step 2 View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

Step 3 Replace the faulty disk.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, perform **Step 5**.

Collect the fault information.

Step 5 On MRS Manager, choose **System** > **Export Log**.

Step 6 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

To obtain the related parameters, perform the following steps:

- Run the following command in the OS to collect data:

  iostat -x -t 1 1

  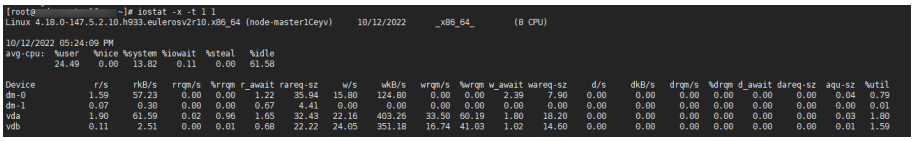

  The command parameters are as follows:

  avgqu-sz indicates the disk queue depth.

  The sum of **r/s** and **w/s** is the IOPS.

  The sum of **rkB/s** and **wkB/s** is the bandwidth.

  %util is the value of **ioutil**.

- The value of **svctm** is calculated as follows:

  svctm = (tot_ticks_new - tot_ticks_old) / (rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old)

  For MRS 2.*x* or earlier:

  If **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old** is **0**, then **svctm** is **0**.

  For MRS 1.9.3.10 or later:

  When the detection period is 30 seconds, if **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, then **svctm = 0**.

  When the detection period is 300 seconds and **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, if **tot_ticks_new - tot_ticks_old = 0**, then **svctm = 0**; otherwise, the value of **svctm** is infinite.

  The parameters in the preceding expression can be obtained as follows:

Obtain the parameter values from the data collected via the **cat /proc/diskstats** command run by the system every 3 seconds. The following shows an example.

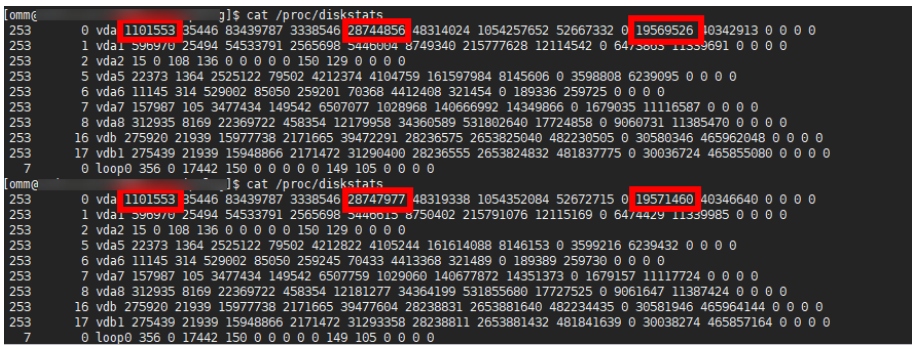

In the data collected for the first time, the number in the fourth column is the value of **rd_ios_old**, the number in the eighth column is the value of **wr_ios_old**, and the number in the thirteenth column is the value of **tot_ticks_old**.

In the data collected for the second time, the number in the fourth column is the value of **rd_ios_new**, the number in the eighth column is the value of **wr_ios_new**, and the number in the thirteenth column is the value of **tot_ticks_new**.

In this case, the value of **svctm** is as follows:

(19571460 - 19569526) / (1101553 + 28747977 - 1101553 - 28744856) = 0.6197

9.5.41 ALM-12357 Failed to Export Audit Logs to OBS (For MRS 2.*x* or Earlier)

Description

If the user has configured audit log export to the OBS on MRS Manager, the system regularly exports audit logs to the OBS. This alarm is reported if the system fails to access the OBS.

This alarm is cleared after the system exports audit logs to the OBS successfully.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12357	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The local system saves a maximum of seven compressed service audit log files. If this alarm persists, local service audit logs may be lost.

The local system saves a maximum of 50 management audit log files (each file contains 100,000 records). If this alarm persists, local management audit logs may be lost.

Possible Causes

- Connection to the OBS server fails.
- The specified OBS file system does not exist.
- The user AK/SK information is invalid.
- The local OBS configuration cannot be obtained.

Procedure

Step 1 Log in to the OBS server and check whether the OBS server can be properly accessed.

- If yes, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 Contact the maintenance personnel to repair OBS. Then check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 On MRS Manager, choose **System** > **Export Audit Log**. Check whether the AK/SK information, file system name, and path are correct.

- If yes, go to **Step 5**.
- If no, go to **Step 4**.

Step 4 Correct the information. Then check whether the alarm is cleared when the export task is executed again.

NOTE

To check alarm clearance quickly, you can set the start time of audit log collection to 10 or 30 minutes later than the current time. After checking the result, restore the original start time.

- If yes, no further action is required.

- If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.42 ALM-13000 ZooKeeper Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the ZooKeeper service status every 30 seconds. This alarm is generated when the ZooKeeper service is unavailable.

This alarm is cleared when the ZooKeeper service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
13000	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

ZooKeeper fails to provide coordination services for upper-layer components and the components depending on ZooKeeper may not run properly.

Possible Causes

- The ZooKeeper instance is abnormal.

- The disk capacity is insufficient.
- The network is faulty.
- The DNS is installed on the ZooKeeper node.

Procedure

Check the ZooKeeper service instance status.

Step 1 On the MRS cluster details page, choose **Components** > **ZooKeeper** > **quorumpeer**.

Step 2 Check whether the ZooKeeper instances are normal.

- If yes, go to **Step 6**.
- If no, go to **Step 3**.

Step 3 Select instances whose status is not good and choose **More** > **Restart Instance**.

Step 4 Check whether the instance status is good after restart.

- If yes, go to **Step 5**.
- If no, go to **Step 19**.

Step 5 On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check disk status.

Step 6 On the MRS cluster details page, choose **Components** > **ZooKeeper** > **quorumpeer**, and check the host information of each node housing the ZooKeeper instance.

Step 7 On the MRS cluster details page, click the **Nodes** tab and expand a node group.

Step 8 In the **Disk Usage** column, check whether the disk space of each node housing ZooKeeper instances is insufficient (disk usage exceeds 80%).

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

Step 9 Expand the disk capacity. For details, see **ALM-12017 Insufficient Disk Capacity (For MRS 2.x or Earlier)**.

Step 10 On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Check network communication status.

Step 11 On the Linux node housing the ZooKeeper instance, run the **ping** command to check whether the host names of other nodes housing the ZooKeeper instances can be pinged successfully.

- If yes, go to **Step 15**.
- If no, go to **Step 12**.

Step 12 Modify the IP addresses in **/etc/hosts** and add the mapping between host names and IP addresses.

Step 13 Run the **ping** command again to check whether the host names of other nodes housing the ZooKeeper instances can be pinged successfully.

- If yes, go to **Step 14**.
- If no, go to **Step 19**.

Step 14 On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Check the DNS.

Step 15 Check whether the DNS is installed on the node housing the ZooKeeper instance. On the Linux node housing the ZooKeeper instance, run the **cat /etc/resolv.conf** command to check whether the file is empty.

- If yes, go to **Step 16**.
- If no, go to **Step 19**.

Step 16 Run the **service named status** command to check whether the DNS is started.

- If yes, go to **Step 17**.
- If no, go to **Step 19**.

Step 17 Run the **service named stop** command to stop the DNS service. If "Shutting down name server BIND waiting for named to shut down (28s)" is displayed, the DNS service is stopped successfully. Comment out the content (if any) in **/etc/resolv.conf**.

Step 18 On the **Alarms** tab page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 19**.

Step 19 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.43 ALM-13001 Available ZooKeeper Connections Are Insufficient (For MRS 2.*x* or Earlier)

Description

The system checks ZooKeeper connections every 30 seconds. This alarm is generated when the system detects that the number of used ZooKeeper instance connections exceeds the threshold (80% of the maximum connections).

This alarm is cleared when the number of used ZooKeeper instance connections is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
13001	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Available ZooKeeper connections are insufficient. When the connection usage reaches 100%, external connections cannot be handled.

Possible Causes

The number of connections to the ZooKeeper node exceeds the threshold. Connection leakage occurs on some connection processes, or the maximum number of connections does not meet the requirement of the actual scenario.

Procedure

Step 1 Check the connection status.

1. On the MRS cluster details page, choose **Alarms** > **ALM-13001 Available ZooKeeper Connections Are Insufficient** > **Location**. Check the IP address of the node for which the alarm is generated.
2. Obtain the PID of the ZooKeeper process. Log in to the node for which this alarm is generated and run the **pgrep -f proc_zookeeper** command.
3. Check whether the PID can be successfully obtained.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2**.

4. Obtain all the IP addresses connected to the ZooKeeper instance and the number of connections and check 10 IP addresses with top connections. Run the **lsof -i|grep $pid | awk '{print $9}' | cut -d : -f 2 | cut -d \> -f 2 | awk '{a[$1]++} END {for(i in a){print i,a[i] | "sort -r -g -k 2"}}' | head -10** command based on the obtained PID value. (**$pid** is the PID obtained in the preceding step.)
5. Check whether the node IP addresses and the number of connections are successfully obtained.
   - If yes, go to **Step 1.6**.
   - If no, go to **Step 2**.
6. Obtain the ID of the port connected to the process. Run the **lsof -i|grep $pid | awk '{print $9}'|cut -d \> -f 2 |grep $IP| cut -d : -f 2** command based on the obtained PID and IP address. (**$pid** and **$IP** are the PID and IP address obtained in the preceding step.)
7. Check whether the port ID is successfully obtained.
   - If yes, go to **Step 1.8**.
   - If no, go to **Step 2**.
8. Obtain the ID of the connected process. Log in to each IP address and run the following command based on the obtained port ID: **lsof -i|grep $port**. (**$port** is the port ID obtained in the preceding step.)
9. Check whether the process ID is successfully obtained.
   - If yes, go to **Step 1.10**.
   - If no, go to **Step 2**.
10. Check whether connection leakage occurs on the process based on the obtained process ID.
    - If yes, go to **Step 1.11**.
    - If no, go to **Step 1.12**.
11. Close the process where connection leakage occurs and check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to **Step 1.12**.
12. On the MRS cluster details page, choose **Components** > **ZooKeeper** > **Service Configuration**. Set **Type** to **All**, choose **quorumpeer** > **Performance**, and change the value of **maxCnxns** to **20000** or more.
13. Check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.44 ALM-13002 ZooKeeper Memory Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the ZooKeeper service status every 30 seconds. The alarm is generated when the memory usage of a ZooKeeper instance exceeds the threshold (80% of the maximum memory).

The alarm is cleared when the memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
13002	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

If the available ZooKeeper memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The memory usage of the ZooKeeper instance is overused or the memory is inappropriately allocated.

Procedure

Step 1 Check the memory usage.

1. On the MRS cluster details page, choose **Alarms** > **ALM-13002 ZooKeeper Memory Usage Exceeds the Threshold** > **Location**. Check the IP address of the instance for which the alarm is generated.

2. On the MRS cluster details page, choose **Components** > **ZooKeeper** > **Instances** > **quorumpeer** (IP address of the instance for which the alarm is generated) > **Customize** > **ZooKeeper Heap And Direct Buffer Resource**. Check the heap memory usage.
3. Check whether the used heap memory of ZooKeeper reaches 80% of the maximum heap memory specified for ZooKeeper.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 1.6**.
4. On MRS Manager, choose **Services** > **ZooKeeper** > **Configuration** > **All** > **quorumpeer** > **System**. Increase the value of **-Xmx** in **GC_OPTS** as required.
5. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 1.6**.
6. On the MRS cluster details page, choose **Components** > **ZooKeeper** > **Instances** > **quorumpeer** (IP address of the instance for which the alarm is generated) > **Customize** > **ZooKeeper Heap And Direct Buffer Resource**. Check the direct buffer memory usage.
7. Check whether the used direct buffer memory of ZooKeeper reaches 80% of the maximum direct buffer memory specified for ZooKeeper.
   - If yes, go to **Step 1.8**.
   - If no, go to **Step 2**.
8. On the MRS cluster details page, choose **Components** > **ZooKeeper** > **Service Configuration**. Set **Type** to **All** and choose **quorumpeer** > **System**.

   Increase the value of **-XX:MaxDirectMemorySize** in **GC_OPTS** as required.
9. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.45 ALM-14000 HDFS Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the service status of NameService every 30 seconds. This alarm is generated when the system considers that the HDFS service is unavailable because all the NameService services are abnormal.

This alarm is cleared when at least one NameService service is normal and the system considers that the HDFS service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14000	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

HDFS fails to provide services for HDFS service-based upper-layer components, such as HBase and MapReduce. As a result, users cannot read or write files.

Possible Causes

- ZooKeeper is abnormal.
- All NameService services are abnormal.

Procedure

Step 1 Check the ZooKeeper status.

1. Go to the MRS cluster details page. On the **Components** tab page, check whether the health status of the ZooKeeper service is **Good**.
   - If yes, go to **Step 1.2**.
   - If no, go to **Step 2.1**.
2. Rectify the health status of the ZooKeeper service. For details, see **ALM-13000 ZooKeeper Service Unavailable (For MRS 2.x or Earlier)**. Then check whether the health status of the ZooKeeper service is **Good**.
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 3**.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.

- If no, go to **Step 2.1**.

Step 2 Handle the NameService service exception alarm.

1. Go to the MRS cluster details page. On the **Alarms** page, check whether all NameService services have abnormal alarms.
   - If yes, go to **Step 2.2**.
   - If no, go to **Step 3**.
2. Handle the abnormal NameService services following the instructions in **ALM-14010 NameService Is Abnormal (For MRS 2.x or Earlier)** and check whether each NameService service exception alarm is cleared.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.46 ALM-14001 HDFS Disk Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the disk usage of the HDFS cluster every 30 seconds and compares the actual disk usage with the threshold. The HDFS cluster disk usage indicator has a default threshold. This alarm is generated when the HDFS disk usage exceeds the threshold.

This alarm is cleared when the disk usage of the HDFS cluster is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14001	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NSName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

The performance of writing data to HDFS is affected.

Possible Causes

The disk space configured for the HDFS cluster is insufficient.

Procedure

Step 1 Check the disk capacity and delete unnecessary files.

1. On the MRS cluster details page, choose **Components** > **HDFS**. The **Service Status** page is displayed.
2. In the **Charts** area, view the value of the monitoring indicator **Percentage of HDFS Capacity** to check whether the HDFS disk usage exceeds the threshold (80% by default).
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 3**.
3. Use the client on the cluster node and run the **hdfs dfsadmin -report** command to check whether the value of **DFS Used%** is less than 100% minus the threshold.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 3**.
4. Use the client on the cluster node and run the **hdfs dfs -rm -r** *file or directory path* command to delete unnecessary files.
5. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Expand the system.

1. Expand the disk capacity.
2. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.47 ALM-14002 DataNode Disk Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the DataNode disk usage every 30 seconds and compares the actual disk usage with the threshold. The **Percentage of DataNode Capacity** indicator has a default threshold. This alarm is generated when the value of the **Percentage of DataNode Capacity** indicator exceeds the threshold.

This alarm is cleared when the value of the **Percentage of DataNode Capacity** indicator is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14002	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Insufficient disk space will impact read/write to HDFS.

Possible Causes

- The disk space configured for the HDFS cluster is insufficient.
- Data skew occurs among DataNodes.

Procedure

Step 1 Check the cluster disk capacity.

1. Go to the MRS cluster details page. On the **Alarms** page, check whether the ALM-14001 HDFS Disk Usage Exceeds the Threshold alarm exists.
   - If yes, go to **Step 1.2**.
   - If no, go to **Step 2.1**.
2. Handle the alarm by following the instructions in ALM-14001 HDFS Disk Usage Exceeds the Threshold and check whether the alarm is cleared.
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 3**.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check the balance status of DataNodes.

1. Use the client on the cluster node, run the **hdfs dfsadmin -report** command to view the value of **DFS Used%** on the DataNode for which the alarm is generated, and compare the value with those on other DataNodes. Check whether the difference between the values is larger than 10.
   - If yes, go to **Step 2.2**.
   - If no, go to **Step 3**.
2. If data skew occurs, use the client on the cluster node and run the **hdfs balancer -threshold 10** command.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.48 ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the number of lost blocks every 30 seconds and compares the number of lost blocks with the threshold. The lost blocks indicator has a default threshold. This alarm is generated when the number of lost blocks exceeds the threshold.

This alarm is cleared when the number of lost blocks is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14003	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NSName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Data stored in HDFS is lost. HDFS may enter the safe mode and cannot provide write services. Lost block data cannot be restored.

Possible Causes

- The DataNode instance is abnormal.
- Data is deleted.

Procedure

Step 1 Check the DataNode instance.

1. On the MRS cluster details page, choose **Components** > **HDFS** > **Instances**.
2. Check whether the status of all DataNode instances is **Good**.
   - If yes, go to **Step 3**.
   - If no, go to **Step 1.3**.
3. Restart the DataNode instance and check whether the restart is successful.
   - If yes, go to **Step 2.2**.
   - If no, go to **Step 2.1**.

Step 2 Delete the damaged file.

1. Use the client on the cluster node. Run the **hdfs fsck / -delete** command to delete the lost file. Then rewrite the file and recover the data.
2. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.49 ALM-14004 Number of Damaged HDFS Blocks Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the number of damaged blocks every 30 seconds and compares the number of damaged blocks with the threshold. The damaged blocks indicator has a default threshold. This alarm is generated when the number of damaged blocks exceeds the threshold.

This alarm is cleared when the number of damaged blocks is less than or equal to the threshold. You are advised to run the **hdfs fsck /** command to check whether any file is completely damaged.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14004	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NSName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Data is damaged and HDFS fails to read files.

Possible Causes

- The DataNode instance is abnormal.
- Data verification information is damaged.

Procedure

Step 1 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.50 ALM-14006 Number of HDFS Files Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system periodically checks the number of HDFS files every 30 seconds and compares the number of HDFS files with the threshold. This alarm is generated when the system detects that the number of HDFS files exceeds the threshold.

This alarm is cleared when the number of HDFS files is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14006	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NSName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Disk storage space is insufficient, which may result in data import failure. The performance of the HDFS system is affected.

Possible Causes

The number of HDFS files exceeds the threshold.

Procedure

Step 1 Check whether unnecessary files exist in the system.

1. Use the client on the cluster node and run the **hdfs dfs -ls** *file or directory path* command to check whether the file or directory can be deleted.
   - If yes, go to **Step 1.2**.
   - If no, go to **Step 2.1**.
2. Run the **hdfs dfs -rm -r** *file or directory path* command. Delete unnecessary files, wait 5 minutes, and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check the number of files in the system.

1. On MRS Manager, choose **System** > **Threshold Configuration**.
2. In the navigation tree on the left, choose **Services** > **HDFS** > **HDFS File** > **Total Number of Files**.
3. In the right pane, modify the threshold in the rule based on the number of current HDFS files.

   To check the number of HDFS files, choose **Services** > **HDFS**, click **Customize** in the **Real-Time Statistics** area on the right, and select the **HDFS File** monitoring item.
4. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.51 ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the HDFS NameNode memory usage every 30 seconds and compares the actual memory usage with the threshold. The HDFS NameNode memory usage has a default threshold. This alarm is generated when the HDFS NameNode memory usage exceeds the threshold.

This alarm is cleared when the HDFS NameNode memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14007	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

If the memory usage of the HDFS NameNode is too high, data read/write performance of HDFS will be affected.

Possible Causes

The HDFS NameNode memory is insufficient.

Procedure

Step 1 Delete unnecessary files.

1. Use the client on the cluster node and run the **hdfs dfs -rm -r** *file or directory path* command to delete unnecessary files.
2. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.52 ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the HDFS DataNode memory usage every 30 seconds and compares the actual memory usage with the threshold. The HDFS DataNode memory usage has a default threshold. This alarm is generated when the HDFS DataNode memory usage exceeds the threshold.

This alarm is cleared when the HDFS DataNode memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14007	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

The HDFS DataNode memory usage is too high, which affects the data read/write performance of the HDFS.

Possible Causes

The HDFS DataNode memory is insufficient.

Procedure

Step 1 Delete unnecessary files.

1. Use the client on the cluster node and run the **hdfs dfs -rm -r** *file or directory path* command to delete unnecessary files.
2. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.53 ALM-14009 Number of Faulty DataNodes Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system periodically checks the number of faulty DataNodes in the HDFS cluster every 30 seconds, and compares the number with the threshold. The number of faulty DataNodes has a default threshold. This alarm is generated when the number of faulty DataNodes in the HDFS cluster exceeds the threshold.

This alarm is cleared when the number of faulty DataNodes in the HDFS cluster is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14009	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Faulty DataNodes cannot provide HDFS services.

Possible Causes

- DataNodes are faulty or overloaded.
- The network between the NameNode and the DataNode is disconnected or busy.
- NameNodes are overloaded.

Procedure

Step 1 Check whether DataNodes are faulty.

1. Use the client on the cluster node to run the **hdfs dfsadmin -report** command to check whether DataNodes are faulty.
   - If yes, go to **Step 1.2**.
   - If no, go to **Step 2.1**.
2. On the MRS cluster details page, choose **Components** > **HDFS** > **Instances** to check whether the DataNode is stopped.
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 2.1**.
3. Select the DataNode instance, and choose **More** > **Restart Instance** to restart it. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check the status of the network between the NameNode and the DataNode.

1. Log in to the service IP address of the node where the faulty DataNode is located, and run the **ping** *IP address of the NameNode* command to check whether the network between the DataNode and the NameNode is abnormal.
   - If yes, go to **Step 2.2**.
   - If no, go to **Step 3.1**.
2. Rectify the network fault. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.

- If no, go to **Step 3.1**.

Step 3 Check whether the DataNode is overloaded.

1. On the MRS cluster details page, click **Alarms** and check whether the alarm ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold exists.
   - If yes, go to **Step 3.2**.
   - If no, go to **Step 4.1**.
2. Follow procedures in **ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)** to handle the alarm and check whether the alarm is cleared.
   - If yes, go to **Step 3.3**.
   - If no, go to **Step 4.1**.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4.1**.

Step 4 Check whether the NameNode is overloaded.

1. On the MRS cluster details page, click **Alarms** and check whether the alarm ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold exists.
   - If yes, go to **Step 4.2**.
   - If no, go to **Step 5**.
2. Follow procedures in **ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold (For MRS 2.x or Earlier)** to handle the alarm and check whether the alarm is cleared.
   - If yes, go to **Step 4.3**.
   - If no, go to **Step 5**.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.54 ALM-14010 NameService Is Abnormal (For MRS 2.*x* or Earlier)

Description

The system checks the NameService service status every 180 seconds. This alarm is generated when the NameService service is unavailable.

This alarm is cleared when the NameService service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14010	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NSName	Specifies the NameService service for which the alarm is generated.

Impact on the System

HDFS fails to provide services for upper-layer components based on the NameService service, such as HBase and MapReduce. As a result, users cannot read or write files.

Possible Causes

- The JournalNode is faulty.
- The DataNode is faulty.
- The disk capacity is insufficient.
- The NameNode enters safe mode.

Procedure

Step 1 Check the status of the JournalNode instance.

1. On the MRS Manager home page, click **Components**.
2. Click **HDFS**.
3. Click **Instance**.
4. Check whether the **Health Status** of the JournalNode is **Good**.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.5**.
5. Select the faulty JournalNode, and choose **More** > **Restart Instance**. Check whether the JournalNode successfully restarts.

- If yes, go to **Step 1.6**.
- If no, go to **Step 5**.

6. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check the status of the DataNode instance.

1. On the MRS cluster details page, click **Components**.
2. Click **HDFS**.
3. In **Operation and Health Summary**, check whether the **Health Status** of all DataNodes is **Good**.
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 2.4**.
4. Click **Instances**. On the DataNode management page, select the faulty DataNode, and choose **More** > **Restart Instance**. Check whether the DataNode successfully restarts.
   - If yes, go to **Step 2.5**.
   - If no, go to **Step 3.1**.
5. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4.1**.

Step 3 Check the disk status.

1. On the MRS cluster details page, click the **Nodes** tab and expand a node group.
2. In the **Disk Usage** column, check whether disk space is insufficient.
   - If yes, go to **Step 3.3**.
   - If no, go to **Step 4.1**.
3. Expand the disk capacity.
4. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4.1**.

Step 4 Check whether NameNode is in the safe mode.

1. Use the client on the cluster node, and run the **hdfs dfsadmin -safemode get** command to check whether **Safe mode is ON** is displayed.

   Information behind **Safe mode is ON** is alarm information and is displayed based actual conditions.
   - If yes, go to **Step 4.2**.
   - If no, go to **Step 5**.
2. Use the client on the cluster node and run the **hdfs dfsadmin -safemode leave** command.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.

- If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.55 ALM-14011 HDFS DataNode Data Directory Is Not Configured Properly (For MRS 2.*x* or Earlier)

Description

The DataNode parameter **dfs.datanode.data.dir** specifies the DataNode data directory. This alarm is generated in any of the following scenarios:

- A configured data directory cannot be created.
- A data directory uses the same disk as other critical directories in the system.
- Multiple directories use the same disk.

This alarm is cleared when the DataNode data directory is configured properly and this DataNode is restarted.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14011	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the DataNode data directory is mounted on critical directories such as the root directory, the disk space of the root directory will be used up after running for a long time. This causes a system fault.

If the DataNode data directory is not configured properly, HDFS performance will deteriorate.

Possible Causes

- The DataNode data directory fails to be created.
- The DataNode data directory uses the same disk as critical directories, such as **/** or **/boot**.
- Multiple directories in the DataNode data directory use the same disk.

Procedure

Step 1 Check the alarm cause and information about the DataNode for which the alarm is generated.

1. On the MRS cluster details page, click **Alarms**. In the alarm list, click the alarm.
2. In the **Alarm Details** area, view **Alarm Cause** to obtain the cause of the alarm. In **HostName** of **Location**, obtain the host name of the DataNode for which the alarm is generated.

Step 2 Delete directories that do not comply with the disk plan from the DataNode data directory.

1. Choose **Components** > **HDFS** > **Instances**. In the instance list, click the DataNode instance on the node for which the alarm is generated.
2. Click **Instance Configuration** and view the value of the DataNode parameter **dfs.datanode.data.dir**.
3. Check whether all DataNode data directories are consistent with the disk plan.
   - If yes, go to **Step 2.4**.
   - If no, go to **Step 2.7**.
4. Modify the DataNode parameter **dfs.datanode.data.dir** and delete the incorrect directories.
5. Choose **Components** > **HDFS** > **Instances** to restart the DataNode instance.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.7**.
7. Log in to the DataNode for which the alarm is generated.
   - If the alarm cause is "The DataNode data directory fails to be created", go to **Step 3.1**.
   - If the alarm cause is "The DataNode data directory uses the same disk as critical directories, such **/** or **/boot**", go to **Step 4.1**.
   - If the alarm cause is "Multiple directories in the DataNode data directory use the same disk", go to **Step 5.1**.

Step 3 Check whether the DataNode data directory fails to be created.

1. Run the following commands to switch the user:

   sudo su - root

   su - omm

2. Run the **ls** command to check whether the directories exist in the DataNode data directory.
   - If yes, go to **Step 7**.
   - If no, go to **Step 3.3**.
3. Run the **mkdir** *data directory* command to create a directory and check whether the directory is successfully created.
   - If yes, go to **Step 6.1**.
   - If no, go to **Step 3.4**.
4. Click **Alarms** to check whether alarm ALM-12017 Insufficient Disk Capacity exists.
   - If yes, go to **Step 3.5**.
   - If no, go to **Step 3.6**.
5. Adjust the disk capacity and check whether alarm ALM-12017 Insufficient Disk Capacity is cleared. For details, see **ALM-12017 Insufficient Disk Capacity (For MRS 2.x or Earlier)**.
   - If yes, go to **ALM-12017 Insufficient Disk Capacity (For MRS 2.x or Earlier)**.
   - If no, go to **Step 7**.
6. Check whether user **omm** has the **rwx** or **x** permission of all the upper-layer directories of the directory. (For example, for **/tmp/abc/**, user **omm** has the **x** permission for directory **tmp** and the **rwx** permission for directory **abc**.)
   - If yes, go to **Step 6.1**.
   - If no, go to **Step 3.7**.
7. Run the **chmod u+rwx** *path* or **chmod u+x** *path* command as the **root** user to add the **rwx** or **x** permission to the paths. Then, go to **Step 3.3**.

Step 4 Check whether the DataNode data directory uses the same disk as other critical directories in the system.

1. Run the **df** command to obtain the disk mounting information of each directory in the DataNode data directory.
2. Check whether the directories mounted to the disk are critical directories, such as **/** or **/boot**.
   - If yes, go to **Step 4.3**.
   - If no, go to **Step 6.1**.
3. Change the value of the DataNode parameter **dfs.datanode.data.dir** and delete the directories that use the same disk as critical directories.
4. Go to **Step 6.1**.

Step 5 Check whether multiple directories in the DataNode data directory use the same disk.

1. Run the **df** command to obtain the disk mounting information of each directory in the DataNode data directory. Record the mounted directory in the command output.
2. Modify the DataNode node parameter **dfs.datanode.data.dir** to reserve one of the directories mounted on the same disk directory.
3. Go to **Step 6.1**.

Step 6 Restart the DataNode and check whether the alarm is cleared.

1. Choose **Components** > **HDFS** > **Instances** to restart the DataNode instance.
2. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 7**.

Step 7 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.56 ALM-14012 HDFS Journalnode Data Is Not Synchronized (For MRS 2.*x* or Earlier)

Description

On the active NameNode, the system checks data synchronization on all JournalNodes in the cluster every 5 minutes. This alarm is generated when data on a JournalNode is not synchronized with that on other JournalNodes.

This alarm is cleared in 5 minutes after data on JournalNodes is synchronized.

Attribute

Alarm ID	Alarm Severity	Auto Clear
14012	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
IP	Specifies the service IP address of the JournalNode instance for which the alarm is generated.

Impact on the System

When a JournalNode is working incorrectly, data on the node is not synchronized with that on other JournalNodes. If data on more than half of JournalNodes is not synchronized, the NameNode cannot work correctly, making the HDFS service unavailable.

Possible Causes

- The JournalNode instance has not been started or has been stopped.
- The JournalNode instance is working incorrectly.
- The network of the JournalNode is unreachable.

Procedure

Step 1 Check whether the JournalNode instance has been started.

1. On the MRS cluster details page, click **Alarms**. In the alarm list, click the alarm.
2. In the **Alarm Details** area, check **Location** and obtain the IP address of the JournalNode for which the alarm is generated.
3. Choose **Components** > **HDFS** > **Instances**. In the instance list, click the JournalNode for which the alarm is generated and check whether **Operating Status** of the node is **Started**.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.4**.
4. Select the JournalNode instance and choose **More** > **Start Instance** to start it.
5. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 2 Check whether the JournalNode instance is working correctly.

1. Check whether **Health Status** of the JournalNode instance is **Good**.
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 2.2**.
2. Select the JournalNode instance and choose **More** > **Restart Instance** to restart it.
3. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 3 Check whether the network of the JournalNode is reachable.

1. On the MRS cluster details page, choose **Components** > **HDFS** > **Instances** to check the service IP address of the active NameNode.
2. Log in to the active NameNode.
3. Run the **ping** command to check whether a timeout occurs or the network between the active NameNode and the JournalNode is unreachable.

   ping *service IP address of the JournalNode*

   - If yes, go to **Step 3.4**.
   - If no, go to **Step 4**.
4. Contact O&M personnel to rectify the network fault. Wait 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.57 ALM-16000 Percentage of Sessions Connected to the HiveServer to the Maximum Number Allowed Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the percentage of sessions connected to the HiveServer to the maximum number allowed every 30 seconds. This indicator can be viewed on the Hive service monitoring page. This alarm is generated when the the percentage of sessions connected to the HiveServer to the maximum number allowed exceeds the specified threshold (90% by default).

This alarm can be automatically cleared when the percentage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16000	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

If a connection alarm is generated, too many sessions are connected to the HiveServer and new connections cannot be created.

Possible Causes

Too many clients are connected to the HiveServer.

Procedure

Step 1 Increase the maximum number of connections to Hive.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Hive** > **Service Configuration** and switch **Basic** to **All**.
3. Increase the value of the **hive.server.session.control.maxconnections** configuration item. Suppose the value of the configuration item is A, the threshold is B, and sessions connected to the HiveServer is C. Adjust the value of the configuration item according to A x B > C. Sessions connected to the HiveServer can be viewed on the Hive monitoring page.
4. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.58 ALM-16001 Hive Warehouse Space Usage Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the Hive warehouse space usage every 30 seconds. The indicator **Percentage of HDFS Space Used by Hive to the Available Space** can be viewed on the Hive service monitoring page. This alarm is generated when the Hive warehouse space usage exceeds the specified threshold (85% by default).

This alarm is cleared when the Hive warehouse space usage is less than or equal to the threshold. You can reduce the warehouse space usage by expanding the warehouse capacity or releasing the used space.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16001	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

The system fails to write data, which causes data loss.

Possible Causes

- The upper limit of the HDFS capacity available for Hive is too small.
- The system disk space is insufficient.
- Some data nodes break down.

Procedure

Step 1 Expand the system configuration.

1. Analyze the cluster HDFS capacity usage and increase the upper limit of the HDFS capacity available for Hive.

   Go to the MRS cluster details page, choose **Components** > **Hive** > **Service Configuration**, set **Type** to **All**, search for **hive.metastore.warehouse.size.percent**, and increase the value of this parameter. Suppose that the value of the configuration item is A, total HDFS storage space is B, the threshold is C, and HDFS space used by Hive is D. Adjust the value of the configuration item according to A x B x C > D. The total HDFS storage space can be viewed on the HDFS monitoring page, and HDFS space used by Hive can be viewed on the Hive monitoring page.

2. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Expand the system.

1. Add nodes.
2. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3.1**.

Step 3 Check whether the data node is normal.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether ALM-12006 Node Fault, ALM-12007 Process Fault, or ALM-14002 DataNode Disk Usage Exceeds the Threshold exists.
   - If yes, go to **Step 3.3**.
   - If no, go to **Step 4**.
3. Clear the alarm by following the steps provided in ALM-12006 Node Fault, ALM-12007 Process Fault, or ALM-14002 DataNode Disk Usage Exceeds the Threshold.
4. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.59 ALM-16002 Hive SQL Execution Success Rate Is Lower Than the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the percentage of the HiveQL statements that are executed successfully every 30 seconds. Percentage of HiveQL statements that are executed successfully = Number of HiveQL statements that are executed successfully by Hive in a specified period/Total number of HiveQL statements that are executed by Hive. This indicator can be viewed on the Hive service monitoring page. This alarm is generated when the percentage of the HiveQL statements that are executed successfully exceeds the specified threshold (90% by default). The name of the host for which the alarm is generated can be obtained from the location information of the alarm. The host IP address is the IP address of the HiveServer node.

This alarm is cleared when the percentage of the HiveQL statements that are executed successfully in a test period is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16002	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold for triggering the alarm.

Impact on the System

The system configuration and performance cannot meet service processing requirements.

Possible Causes

- A syntax error occurs in HiveQL commands.

- The HBase service is abnormal when a Hive on HBase task is being performed.
- Basic services that are depended on are abnormal, such as HDFS, Yarn, and ZooKeeper.

Procedure

Step 1 Check whether the HiveQL commands comply with syntax.

1. Use the Hive client to log in to the HiveServer node for which the alarm is generated. Query the HiveQL syntax standard provided by Apache, and check whether the HiveQL commands are correct.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.2**.

   NOTE

   To view the user who runs an incorrect statement, download HiveServerAudit logs of the HiveServer node for which this alarm is generated. Set **Start time** and **End time** to 10 minutes before and after the alarm generation time respectively. Open the log file and search for the **Result=FAIL** keyword to filter the log information about the incorrect statement, and then view the user who runs the incorrect statement according to **UserName** in the log information.
2. Enter correct HiveQL statements, and check whether the command can be properly executed.
   - If yes, go to **Step 4.5**.
   - If no, go to **Step 2.1**.

Step 2 Check whether the HBase service is abnormal.

1. Check whether a Hive on HBase task is performed.
   - If yes, go to **Step 2.2**.
   - If no, go to **Step 3.1**.
2. Check whether the HBase service is normal in the service list.
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 2.3**.
3. Check the alarms displayed on the alarm page and clear them according to **Alarm Help**.
4. Enter correct HiveQL statements, and check whether the command can be properly executed.
   - If yes, go to **Step 4.5**.
   - If no, go to **Step 3.1**.

Step 3 Check whether the Spark service is abnormal.

1. Check whether the Spark service is normal in the service list.
   - If yes, go to **Step 4.1**.
   - If no, go to **Step 3.2**.
2. Check the alarms displayed on the alarm page and clear them according to **Alarm Help**.

3. Enter correct HiveQL statements, and check whether the command can be properly executed.
   - If yes, go to **Step 4.5**.
   - If no, go to **Step 4.1**.

Step 4 Check whether HDFS, Yarn, and ZooKeeper are normal.

1. Go to the MRS cluster details page and click **Components**.
2. In the service list, check whether the services, such as HDFS, Yarn, and ZooKeeper are normal.
   - If yes, go to **Step 4.5**.
   - If no, go to **Step 4.3**.
3. Check the alarms displayed on the alarm page and clear them according to **Alarm Help**.
4. Enter correct HiveQL statements, and check whether the command can be properly executed.
   - If yes, go to **Step 4.5**.
   - If no, go to **Step 5**.
5. Wait one minute and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.60 ALM-16004 Hive Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Hive service status every 30 seconds. This alarm is generated when the Hive service is unavailable.

This alarm is cleared when the Hive service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16004	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system cannot provide data loading, query, and extraction services.

Possible Causes

- Basic services, such as ZooKeeper, HDFS, Yarn, and DBService work incorrectly, or the Hive process is faulty.
  - ZooKeeper is abnormal.
  - HDFS is abnormal.
  - Yarn is abnormal.
  - DBService is abnormal.
  - The Hive service process is faulty. If the alarm is caused by a Hive process fault, the alarm report has a delay of about 5 minutes.
- The network communication between the Hive service and basic services is interrupted.

Procedure

Step 1 Check the HiveServer/MetaStore process status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Hive** > **Instances**. In the Hive instance list, check whether the status of all HiveSserver/MetaStore instances is **Unknown**.
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 2**.
3. Above the Hive instance list, choose **More** > **Restart Instance** to restart the HiveServer/MetaStore process.
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Check the ZooKeeper status.

1. Go to the cluster details page and choose **Alarms**.

2. On MRS Manager, check whether the ALM-12007 Process Fault alarm is reported.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
3. In the **Alarm Details** area of ALM-12007 Process Fault, check whether **ServiceName** is **ZooKeeper**.
   - If yes, go to **Step 2.4**.
   - If no, go to **Step 3**.
4. Rectify the fault by following steps provided in ALM-12007 Process Fault.
5. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Check the HDFS status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, check whether the alarm ALM-14000 HDFS Service Unavailable exists.
   - If yes, go to **Step 3.3**.
   - If no, go to **Step 4**.
3. Rectify the fault by following the steps provided in ALM-14000 HDFS Service Unavailable.
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Check the Yarn status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list on MRS Manager, check whether the alarm ALM-18000 Yarn Service Unavailable is generated.
   - If yes, go to **Step 4.3**.
   - If no, go to **Step 4**.
3. Rectify the fault by following the steps provided in ALM-18000 Yarn Service Unavailable.
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 5 Check the DBService status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list on MRS Manager, check whether ALM-27001 DBService Unavailable is generated.
   - If yes, go to **Step 5.3**.

- If no, go to **Step 6**.

3. Rectify the fault by following the handling procedure in **ALM-27001 DBService Unavailable (For MRS 2.x or Earlier)**.
4. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 6**.

Step 6 Check the network connection between Hive and ZooKeeper, HDFS, Yarn, and DBService.

1. Go to the MRS cluster details page and click **Components**.
2. Click **Hive**.
3. Click **Instances**.

   The HiveServer instance list is displayed.
4. Click **Host Name** in the row of **HiveServer**.

   The HiveServer host status page is displayed.
5. Record the IP address under **Summary**.
6. Use the IP address obtained in **Step 6.5** to log in to the host where HiveServer is located.
7. Run the **ping** command to check whether the network connection between the host that runs HiveServer and the hosts that run the ZooKeeper, HDFS, Yarn, and DBService services is normal. Methods of obtaining IP addresses of the hosts that run ZooKeeper, HDFS, Yarn, and DBService services as well as the HiveServer IP address are the same.
   - If yes, go to **Step 7**.
   - If no, go to **Step 6.8**.
8. Contact the O&M personnel to restore the network.
9. In the alarm list, check whether ALM-16004 Hive Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 7**.

Step 7 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.61 ALM-16005 Number of Failed Hive SQL Executions in the Last Period Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks whether the number of Hive SQL statements that fail to be executed has exceeded the threshold in the last 10-minute period. This alarm is generated when the number of failed Hive SQL statement executions in the last 10 minutes is greater than the threshold. In the next 10 minutes, if the number of failed Hive SQL statement executions is less than the threshold, the alarm is automatically cleared.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16005	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

None

Possible Causes

The Hive SQL syntax is incorrect. As a result, the Hive SQL statements fail to be executed.

Procedure

Check the Hive SQL statements that fail to be executed, correct the syntax, and execute the SQL statements again.

Reference

None

9.5.62 ALM-18000 Yarn Service Unavailable (For MRS 2.*x* or Earlier)

Description

The alarm module checks the Yarn service status every 30 seconds. This alarm is generated when the Yarn service is unavailable.

This alarm is cleared when the Yarn service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18000	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The cluster cannot provide the Yarn service. Users cannot run new applications. Submitted applications cannot be run.

Possible Causes

- ZooKeeper is abnormal.
- HDFS is abnormal.
- There is no active ResourceManager node in the Yarn cluster.
- All NodeManager nodes in the Yarn cluster are abnormal.

Procedure

Step 1 Check the ZooKeeper status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, check whether the alarm ALM-13000 ZooKeeper Service Unavailable exists.

- If yes, go to **Step 1.3**.
- If no, go to **Step 2.2**.

3. Rectify the fault by following the handling procedure in **ALM-13000 ZooKeeper Service Unavailable (For MRS 2.x or Earlier)**. Then, check whether this alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.2**.

Step 2 Check the HDFS status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, check whether an HDFS alarm is generated.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3.2**.
3. Click **Alarms**, and handle HDFS alarms according to **Alarm Help**. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3.2**.

Step 3 Check the ResorceManager status in the Yarn cluster.

1. Go to the MRS cluster details page and click **Components**.
2. Click **Yarn**.
3. In **Yarn Summary**, check whether there is an active ResourceManager node in the Yarn cluster.
   - If yes, go to **Step 4.2**.
   - If no, go to **Step 5**.

Step 4 Check the NodeManager node status in the Yarn cluster.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Yarn** > **Instances**.
3. Check **Health Status** of NodeManager, and check whether there are unhealthy nodes.
   - If yes, go to **Step 4.4**.
   - If no, go to **Step 5**.
4. Rectify the fault by following the procedure provided in **ALM-18002 NodeManager Heartbeat Lost (For MRS 2.x or Earlier)** or **ALM-18003 NodeManager Unhealthy (For MRS 2.x or Earlier)**. Then, check whether this alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.63 ALM-18002 NodeManager Heartbeat Lost (For MRS 2.*x* or Earlier)

Description

The system checks the number of lost NodeManager nodes every 30 seconds, and compares the number of lost nodes with the threshold. The **Lost Nodes** indicator has a default threshold. This alarm is generated when the value of the **Lost Nodes** indicator exceeds the threshold.

This alarm is cleared when the value of **Lost Nodes** is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18002	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

- The lost NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

Possible Causes

- NodeManager is forcibly deleted without decommission.
- All NodeManager instances are stopped or the NodeManager process is faulty.

- The host where the NodeManager node resides is faulty.
- The network between the NodeManager and ResourceManager is disconnected or busy.

Procedure

Step 1 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.64 ALM-18003 NodeManager Unhealthy (For MRS 2.*x* or Earlier)

Description

The system checks the number of abnormal NodeManager nodes every 30 seconds, and compares the number of abnormal nodes with the threshold. The **Unhealthy Nodes** indicator has a default threshold. This alarm is generated when the value of the **Unhealthy Nodes** indicator exceeds the threshold.

This alarm is cleared when the value of **Unhealthy Nodes** is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18003	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

- The faulty NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

Possible Causes

- The disk space of the host where the NodeManager node resides is insufficient.
- User **omm** does not have the permission to access a local directory on the NodeManager node.

Procedure

Step 1 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.65 ALM-18004 NodeManager Disk Usability Ratio Is Lower Than the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the available disk space of each NodeManager node every 30 seconds and compares the disk availability rate with the threshold. A default threshold range is provided for the **NodeManager Disk Usability Ratio**. This alarm is generated when the system detects that the actual **NodeManager Disk Usability Ratio** is lower than the threshold.

This alarm is automatically cleared when the value of **NodeManager Disk Usability Ratio** is greater than the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18004	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold for triggering the alarm.

Impact on the System

- The NodeManager node whose disk availability rate is lower than the threshold may fail to provide the Yarn service.
- The number of containers decreases, so the cluster performance may deteriorate.

Possible Causes

- The disk space of the host where the NodeManager node resides is insufficient.
- User **omm** does not have the permission to access a local directory on the NodeManager node.

Procedure

Step 1 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.66 ALM-18006 MapReduce Job Execution Timeout (For MRS 2.*x* or Earlier)

Description

The alarm module checks the MapReduce job execution every 30 seconds. This alarm is generated when the execution of a submitted MapReduce job times out.

This alarm must be manually cleared.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18006	Major	No

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Execution of the submitted MapReduce job times out, so no execution result can be obtained. Execute the job again after rectifying the fault.

Possible Causes

It takes a long time to execute a MapReduce job. However, the specified time is less than the required execution time.

Procedure

Step 1 Check whether time is improperly set.

Set **-Dapplication.timeout.interval** to a larger value, or do not set the parameter. Check whether the MapReduce job can be executed.

- If yes, go to **Step 2.5**.
- If no, go to **Step 2.2**.

Step 2 Check the Yarn status.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list on MRS Manager, check whether the alarm ALM-18000 Yarn Service Unavailable is generated.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.

3. Rectify the fault by following the handling procedure in **ALM-18000 Yarn Service Unavailable (For MRS 2.x or Earlier)**.
4. Run the MapReduce job command again to check whether the MapReduce job can be executed.
   - If yes, go to **Step 2.5**.
   - If no, go to **Step 4**.
5. In the alarm list, click  in the **Operation** column of the alarm to manually clear the alarm. No further action is required.

Step 3 Adjust the timeout threshold.

On MRS Manager, choose **System** > **Threshold Configuration** > **Services** > **Yarn** > **Timed out Applications**, and increase the maximum number of timeout tasks allowed by the current threshold rule. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.67 ALM-18008 Heap Memory Usage of Yarn ResourceManager Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the heap memory usage of Yarn ResourceManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Yarn ResourceManager exceeds the threshold (80% of the maximum memory by default).

To change the threshold, choose **System** > **Threshold Configuration** > **Service** > **Yarn**. The alarm is cleared when the heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18008	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

When the heap memory usage of Yarn ResourceManager is overhigh, the performance of Yarn task submission and operation is affected. What is more, a memory overflow occurs so that the Yarn service is unavailable.

Possible Causes

The heap memory of the Yarn ResourceManager instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Step 1 Check the heap memory usage.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **18008** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Yarn** > **Instances** > **ResourceManager** (IP address of the instance for which the alarm is generated) > **Customize** > **Percentage of Used Heap Memory of the ResourceManager**. Check the heap memory usage.
4. Check whether the heap memory usage of ResourceManager has reached the threshold (80% of the maximum memory).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Yarn** > **Service Configuration**. Set **Type** to **All** and choose **ResourceManager** > **System**. Change the values of **-Xmx** and **-Xms** in the **GC_OPTS** parameter based on the site requirements to ensure that the value of **-Xms** is less than that of **-Xmx**. Click **Save Configuration** and select **Restart Role Instance**. Click **OK**.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.

- If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.68 ALM-18009 Heap Memory Usage of MapReduce JobHistoryServer Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the heap memory usage of MapReduce JobHistoryServer every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of MapReduce JobHistoryServer exceeds the threshold (80% of the maximum memory by default).

To change the threshold, choose **System** > **Threshold Configuration** > **Service** > **MapReduce**. The alarm is cleared when the heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18009	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

When the heap memory usage of MapReduce JobHistoryServer is overhigh, the performance of MapReduce log archiving is affected. What is more, a memory overflow occurs so that the Yarn service is unavailable.

Possible Causes

The heap memory of the MapReduce JobHistoryServer instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Step 1 Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **18009** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **MapReduce** > **Instance** > **JobHistoryServer** (IP address of the instance for which the alarm is generated) > **Customize** > **JobHistoryServer Heap Memory Usage Statistics**. Check the heap memory usage.
4. Check whether the heap memory usage of JobHistoryServer has reached the threshold (80% of the maximum heap memory).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **MapReduce** > **Service Configuration**. Set **Type** to **All** and choose **JobHistoryServer** > **System**. Increase the value of **-Xmx** in the **GC_OPTS** parameter as required, click **Save Configuration**, and select **Restart the affected services or instances.** Click **OK**.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.69 ALM-18010 Number of Pending Yarn Tasks Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the number of pending Yarn tasks every 30 seconds and compares the number of tasks with the threshold. This alarm is generated when the number of pending tasks exceeds the threshold.

You can change the threshold by choosing **System** > **Configure Alarm Threshold** > **Service** > **Yarn** > **Queue Root Pending Applications** > **Queue Root Pending Applications** on MRS Manager.

This alarm is cleared when the number of pending tasks is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18010	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Tasks may be stacked and cannot be processed in a timely manner.

Possible Causes

The computing capability of the cluster is lower than the task submission rate. As a result, the task cannot be processed in a timely manner after being submitted.

Procedure

Step 1 Check the usage of memory and vCores on the Yarn page.

Check whether the values of **Memory Used|Memory Total** and **VCores Used|VCores Total** on the native Yarn page reach or approach the maximum values.

- If yes, go to **Step 2**.
- If no, go to **Step 5**.

Step 2 Check the number of submitted tasks.

Check whether the running tasks are submitted at a normal frequency.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Scale out the cluster.

The scale-out is based on the site requirements. For details, see .

Step 4 After the scale-out is completed, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.70 ALM-18011 Memory of Pending Yarn Tasks Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the memory of pending Yarn tasks every 30 seconds and compares the memory with the threshold. This alarm is generated when the memory of pending tasks exceeds the threshold.

You can change the threshold by choosing **System** > **Configure Alarm Threshold** > **Service** > **Yarn** > **Queue Root Pending Memory** > **Queue Root Pending Memory** on MRS Manager.

This alarm is cleared when the memory of pending tasks is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18011	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Tasks may be stacked and cannot be processed in a timely manner.

Possible Causes

The computing capability of the cluster is lower than the task submission rate. As a result, the task cannot be processed in a timely manner after being submitted.

Procedure

Step 1 Check the usage of memory and vCores on the Yarn page.

Check whether the values of **Memory Used|Memory Total** and **VCores Used|VCores Total** on the native Yarn page reach or approach the maximum values.

- If yes, go to **Step 2**.
- If no, go to **Step 5**.

Step 2 Check the number of submitted tasks.

Check whether the running tasks are submitted at a normal frequency.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Scale out the cluster.

The scale-out is based on the site requirements. For details, see .

Step 4 After the scale-out is completed, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.71 ALM-18012 Number of Terminated Yarn Tasks in the Last Period Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the number of terminated Yarn tasks every 10 minutes. This alarm is generated when the number of terminated Yarn tasks in the last 10 minutes is greater than the threshold. This alarm is automatically cleared when the number of terminated Yarn tasks is less than the threshold in the next 10 minutes.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18012	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

None

Possible Causes

A user manually stops a running Yarn task.

Procedure

Check the task termination operator in the Yarn logs and audit logs, and determine the cause of the task termination.

Reference

None

9.5.72 ALM-18013 Number of Failed Yarn Tasks in the Last Period Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the number of failed Yarn tasks every 10 minutes. This alarm is generated when the number of failed Yarn tasks in the last 10 minutes is greater than the threshold. This alarm is automatically cleared when the number of failed Yarn tasks is less than the threshold in the next 10 minutes.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18013	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

None

Possible Causes

The submitted Yarn job program is incorrect. For example, the parameter for Spark to submit a job is incorrect.

Procedure

Check the log of the failed job, locate the failure cause, modify the job, and submit the job again.

Reference

None

9.5.73 ALM-19000 HBase Service Unavailable (For MRS 2.*x* or Earlier)

Description

The alarm module checks the HBase service status every 30 seconds. This alarm is generated when the HBase service is unavailable.

This alarm is cleared when the HBase service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
19000	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Operations cannot be performed, such as reading or writing data and creating tables.

Possible Causes

- ZooKeeper is abnormal.
- HDFS is abnormal.
- HBase is abnormal.
- The network is abnormal.

Procedure

Step 1 Check the ZooKeeper status.

1. Go to the MRS cluster details page and click **Components**.
2. In the service list, check whether the health status of ZooKeeper is **Good**.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.3**.
3. In the alarm list, check whether the alarm ALM-13000 ZooKeeper Service Unavailable exists.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2.1**.
4. Rectify the fault by following the steps provided in ALM-13000 ZooKeeper Service Unavailable.
5. Wait several minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check the HDFS status.

1. On MRS Manager, check whether the ALM-14000 HDFS Service Unavailable alarm is reported.
   - If yes, go to **Step 2.2**.
   - If no, go to **Step 3**.
2. Rectify the fault by following the steps provided in ALM-14000 HDFS Service Unavailable.
3. Wait several minutes and check whether the alarm is cleared.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.74 ALM-19006 HBase Replication Sync Failed (For MRS 2.*x* or Earlier)

Description

This alarm is generated when disaster recovery (DR) data fails to be synchronized to a standby cluster.

This alarm is cleared when DR data synchronization succeeds.

Attribute

Alarm ID	Alarm Severity	Auto Clear
19006	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

HBase data in a cluster fails to be synchronized to the standby cluster, causing data inconsistency between active and standby clusters.

Possible Causes

- The HBase service on the standby cluster is abnormal.
- The network is abnormal.

Procedure

Step 1 Observe whether the system automatically clears the alarm.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, click the alarm to obtain alarm generation time from **Generated Time** in **Alarm Details**. Check whether the alarm has existed for over 5 minutes.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.3**.
3. Wait 5 minutes and check whether the alarm is automatically cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check the HBase service status of the standby cluster.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, click the alarm and obtain **HostName** from **Location** in **Alarm Details**.
3. Log in to the node where the HBase client of the active cluster is located. Run the following commands to switch the user:

sudo su - root

su - omm

4. Run the **status 'replication', 'source'** command to check the synchronization status of the faulty node.

   The synchronization status of a node is as follows.

   ```
   10-10-10-153:
    SOURCE: PeerID=abc, SizeOfLogQueue=0, ShippedBatches=2, ShippedOps=2, ShippedBytes=320, LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0, TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=0, TimeStampsOfLastShippedOp=Mon Jul 18 09:53:28 CST 2016, Replication Lag=0, FailedReplicationAttempts=0
    SOURCE: PeerID=abc1, SizeOfLogQueue=0, ShippedBatches=1, ShippedOps=1, ShippedBytes=160, LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0, TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=16788, TimeStampsOfLastShippedOp=Sat Jul 16 13:19:00 CST 2016, Replication Lag=16788, FailedReplicationAttempts=5
   ```

5. Obtain **PeerID** corresponding to a record whose **FailedReplicationAttempts** value is greater than 0.

   In the preceding step, data on the faulty node **10-10-10-153** fails to be synchronized to a standby cluster whose **PeerID** is **abc1**.

6. Run the **list_peers** command to find the cluster and the HBase instance corresponding to **PeerID**.

   ```
   PEER_ID CLUSTER_KEY STATE TABLE_CFS
    abc1 10.10.10.110,10.10.10.119,10.10.10.133:24002:/hbase2 ENABLED
    abc 10.10.10.110,10.10.10.119,10.10.10.133:24002:/hbase ENABLED
   ```

   In the preceding information, **/hbase2** indicates that data is synchronized to the HBase2 instance of the standby cluster.

7. In the service list of the standby cluster, check whether the health status of the HBase instance obtained in **Step 2.6** is **Good**.
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 2.8**.
8. In the alarm list, check whether the alarm ALM-19000 HBase Service Unavailable exists.
   - If yes, go to **Step 2.9**.
   - If no, go to **Step 3.1**.
9. Rectify the fault by following the steps provided in ALM-19000 HBase Service Unavailable.
10. Wait several minutes and check whether the alarm is cleared.
    - If yes, no further action is required.
    - If no, go to **Step 3.1**.

Step 3 Check the network connection between RegionServers on active and standby clusters.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, click the alarm and obtain **HostName** from **Location** in **Alarm Details**.
3. Log in to the faulty RegionServer node.
4. Run the **ping** command to check whether the network connection between the faulty RegionServer node and the host where RegionServer of the standby cluster resides is normal.

- If yes, go to **Step 4**.
- If no, go to **Step 3.5**.

5. Contact the O&M personnel to restore the network.
6. After the network recovers, check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.75 ALM-19007 HBase Merge Queue Exceeds the Threshold (for 2.x and Earlier Versions)

Description

The system checks the HBase compaction queue size every 30 seconds. This alarm is generated when the compaction queue size exceeds the alarm threshold (**100** by default) for three consecutive times. This alarm is cleared when the compaction queue size is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
19007	Minor	Yes

Parameters

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Host Name	Specifies the host for which the alarm is generated.

Impact on the System

The cluster performance may deteriorate, affecting data read and write.

Possible Causes

- The number of HBase RegionServers is too small.
- There are too many regions on a RegionServer of HBase.
- The HBase RegionServer heap size is small.
- Resources are insufficient.
- Related parameters are not configured properly.

Procedure

Step 1 Check whether related HBase parameters are properly configured.

1. Log in to the MRS cluster details page, choose **Components** > **HBase** > **Service Configuration**, switch **Basic Configuration** to **All Configurations**, and search for **hbase.hstore.compaction.min** and **hbase.hstore.compaction.max**, and increase the values of **hbase.regionserver.thread.compaction.small** and **hbase.regionserver.thread.compaction.throttle**.

   NOTE

   If you did not synchronize IAM users, perform synchronization first. (In the **Dashboard** tab, click **Synchronize** next to **IAM User Sync**.)

2. Save the configuration, and restart the HBase service during off-peak hours or perform a rolling restart to make the configuration take effect.
3. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

None

9.5.76 ALM-20002 Hue Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Hue service status every 60 seconds. This alarm is generated if the Hue service is unavailable.

This alarm is cleared when the Hue service is normal.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
20002	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system cannot provide data loading, query, and extraction services.

Possible Causes

- The KrbServer service on which Hue depends is abnormal.
- The DBService service on which Hue depends is abnormal.
- The network connection to DBService is abnormal.

Procedure

Check whether the KrbServer service is normal.

Step 1 Go to the MRS cluster details page and click **Components**.

Step 2 In the service list, check whether **Health Status** of **KrbServer** is **Good**.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Click **Restart** in the **Operation** column of the KrbServer service to restart the service.

Step 4 Wait for several minutes. Check whether ALM-20002 Hue Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether DBService is normal.

Step 5 Go to the MRS cluster details page and click **Components**.

Step 6 In the service list, check whether **Health Status** of **DBService** is **Good**.

- If yes, go to **Step 9**.
- If no, go to **Step 7**.

Step 7 Click **Restart** in the **Operation** column of the DBService service to restart the service.

NOTE

To restart the service, you need to enter the password of the MRS Manager administrator and select **Start or restart related services** .

Step 8 Wait for several minutes. Check whether ALM-20002 Hue Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether the network connected to DBService is normal.

Step 9 Choose **Components** > **Hue** > **Instance** and record the IP address of the active Hue node.

Step 10 Use PuTTY to log in to the active Hue.

Step 11 Run the **ping** command to check whether the network connection between the host where the active Hue is located and the host where DBService is located is normal. (The method of obtaining the DBService service IP address is the same as that of obtaining the active Hue IP address.)

- If yes, go to **Step 17**.
- If no, go to **Step 12**.

Step 12 Contact the network administrator to repair the network.

Step 13 Wait for several minutes. Check whether ALM-20002 Hue Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to **Step 17**.

Collect fault information.

Step 14 On MRS Manager, choose **System** > **Export Log**.

Step 15 Select the following nodes from the **Services** drop-down list and click **OK**.

- Hue
- Controller

Step 16 Set **Start Time** and **End Time** for log collection to 10 minutes before and after the alarm is generated, select an export type, and click **OK** to collect the corresponding fault log information.

Restart Hue.

Step 17 Choose **Components** > **Hue**.

Step 18 Choose **More** > **Restart Service** and click **OK**.

Step 19 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 20**.

Step 20 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.77 ALM-23001 Loader Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Loader service availability every 60 seconds. This alarm is generated if the Loader service is unavailable and is cleared after the Loader service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
23001	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Data loading, import, and conversion are unavailable.

Possible Causes

- The services that Loader depends on are abnormal.
  - ZooKeeper is abnormal.
  - HDFS is abnormal.
  - DBService is abnormal.
  - Yarn is abnormal.
  - MapReduce is abnormal.
- The network is faulty. Loader cannot communicate with its dependent services.
- Loader is running improperly.

Procedure

Step 1 Check the ZooKeeper status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **ZooKeeper** and check whether the health status of ZooKeeper is normal.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 1.3**.
3. Choose **More > Restart Service** to restart ZooKeeper. After ZooKeeper starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 1.4**.
4. On MRS Manager, check whether the ALM-12007 Process Fault alarm is reported.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2.1**.
5. In **Alarm Details** of the "ALM-12007 Process Fault" alarm, check whether **ServiceName** is **ZooKeeper**.
   - If yes, go to **Step 1.6**.
   - If no, go to **Step 2.1**.

6. Clear the alarm according to the handling suggestions of "ALM-12007 Process Fault".
7. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check the HDFS status.

1. Go to the MRS cluster details page and choose **Alarms**.
2. On MRS Manager, check whether the "ALM-14000 HDFS Service Unavailable alarm" is reported.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3.1**.
3. Clear the alarm according to the handling suggestions of "ALM-14000 HDFS Service Unavailable".
4. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3.1**.

Step 3 Check the DBService status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **DBService** to check whether the health status of DBService is normal.
   - If yes, go to **Step 4.1**.
   - If no, go to **Step 3.3**.
3. Choose **More > Restart Service** to restart DBService. After DBService starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4.1**.

Step 4 Check the MapReduce status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **MapReduce** and check whether the health status of MapReduce is normal.
   - If yes, go to **Step 5.1**.
   - If no, go to **Step 4.3**.
3. Choose **More > Restart Service** to restart MapReduce. After MapReduce starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5.1**.

Step 5 Check the Yarn status.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **Yarn** and check whether the health status of Yarn is normal.
   - If yes, go to **Step 5.4**.

- If no, go to **Step 5.3**.

3. Choose **More > Restart Service** to restart Yarn. After Yarn starts, check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5.4**.
4. On MRS Manager, check whether the "ALM-18000 Yarn Service Unavailable" alarm is reported.
   - If yes, go to **Step 5.5**.
   - If no, go to **Step 6.1**.
5. Clear the alarm according to the handling suggestions of "ALM-18000 Yarn Service Unavailable".
6. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 6.1**.

Step 6 Check the network connections between Loader and its dependent components.

1. Go to the MRS cluster details page and click **Components**.
2. Click **Loader**.
3. Click **Instance**. The Sqoop instance list is displayed.
4. Record the management IP addresses of all Sqoop instances.
5. Log in to the hosts using the IP addresses obtained in **Step 6.4**. Run the following commands to switch the user:

   sudo su - root

   su - omm
6. Run the **ping** command to check whether the network connection between the hosts where the Sqoop instances reside and the dependent components is normal. (The dependent components include ZooKeeper, DBService, HDFS, MapReduce, and Yarn. The method to obtain the IP addresses of the dependent components is the same as that used to obtain the IP addresses of the Sqoop instances.)
   - If yes, go to **Step 7**.
   - If no, go to **Step 6.7**.
7. Contact the network administrator to repair the network.
8. Check whether the "ALM-23001 Loader Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 7**.

Step 7 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

9.5.78 ALM-24000 Flume Service Unavailable (For MRS 2.*x* or Earlier)

Description

The alarm module checks the Flume service status every 180 seconds. This alarm is generated if the Flume service is abnormal.

This alarm is cleared after the Flume service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
24000	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Flume cannot work and data transmission is interrupted.

Possible Causes

- HDFS is unavailable.
- LdapServer is unavailable.

Procedure

Step 1 Check the HDFS status.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Check whether the ALM-14000 HDFS Service Unavailable alarm is generated.
   - If yes, clear the alarm according to the handling suggestions of "ALM-14000 HDFS Service Unavailable".
   - If no, go to **Step 2**.

Step 2 Check the LdapServer status.

Check whether the ALM-25000 LdapServer Service Unavailable alarm is generated.

- If yes, clear the alarm according to the handling suggestions of "ALM-25000 LdapServer Service Unavailable".
- If no, go to **Step 3.2**.

Step 3 Check whether the HDFS and LdapServer services are stopped.

1. Go to the MRS cluster details page and click **Components**.
2. In the service list on MRS Manager, check whether the HDFS and LdapServer services are stopped.
   - If yes, start the HDFS and LdapServer services and go to **Step 3.3**.
   - If no, go to **Step 4**.
3. Check whether the "ALM-24000 Flume Service Unavailable" alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.79 ALM-24001 Flume Agent Is Abnormal (For MRS 2.*x* or Earlier)

Description

This alarm is generated if the Flume agent monitoring module detects that the Flume agent process is abnormal.

This alarm is cleared after the Flume agent process recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
24001	Minor	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Functions of the alarmed Flume agent instance are abnormal. Data transmission tasks of the instance are suspended. In real-time data transmission, data will be lost.

Possible Causes

- The **JAVA_HOME** directory does not exist or the Java permission is incorrect.
- The permission of the Flume agent directory is incorrect.

Procedure

Step 1 Check the Flume agent's configuration file.

1. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:

   sudo su - root

2. Run the **cd** *Flume installation directory*/**fusioninsight-flume-1.6.0/conf/** command to go to Flume's configuration directory.
3. Run the **cat ENV_VARS** command. Check whether the **JAVA_HOME** directory exists and whether the Flume agent user has execute permission of Java.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.4**.
4. Specify the correct JAVA_HOME directory and grant the Flume agent user with the execute permission of Java. Then go to **Step 2.4**.

Step 2 Check the permission of the Flume agent directory.

1. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:

   sudo su - root

2. Run the following command to access the installation directory of the Flume agent:

   cd *Flume agent installation directory*

3. Run the **ls -al * -R** command. Check whether the owner of all files is the Flume agent user.
   - If yes, go to **Step 3**.
   - If no, run the **chown** command and change the owner of the files to the Flume agent user. Then go to **Step 2.4**.

4. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.80 ALM-24003 Flume Client Connection Interrupted (For MRS 2.*x* or Earlier)

Description

The alarm module monitors the port connection status on the Flume server. This alarm is generated if the Flume server fails to receive a connection message from the Flume client in 3 consecutive minutes.

This alarm is cleared after the Flume server receives a connection message from the Flume client.

Attribute

Alarm ID	Alarm Severity	Auto Clear
24003	Major	Yes

Parameters

Parameter	Description
ClientIP	Specifies the IP address of the Flume client.
ServerIP	Specifies the IP address of the Flume server.
ServerPort	Specifies the port on the Flume server.

Impact on the System

The communication between the Flume client and server fails. The Flume client cannot send data to the Flume server.

Possible Causes

- The network between the Flume client and server is faulty.
- The Flume client's process is abnormal.
- The Flume client is incorrectly configured.

Procedure

Step 1 Check the network between the Flume client and server.

1. Log in to the host where the alarmed Flume client resides. Run the following command to switch to user **root**:

   sudo su - root

2. Run the **ping** *Flume server IP address* command to check whether the network between the Flume client and server is normal.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 4**.

Step 2 Check whether the Flume client's process is normal.

1. Log in to the host where the alarmed Flume client resides. Run the following command to switch to user **root**:

   sudo su - root

2. Run the **ps -ef|grep flume |grep client** command to check whether the Flume client process exists.
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 4**.

Step 3 Check the Flume client configuration.

1. Log in to the host where the alarmed Flume client resides. Run the following command to switch to user **root**:

   sudo su - root

2. Run the **cd** *Flume installation directory*/**fusioninsight-flume-1.6.0/conf/** command to go to Flume's configuration directory.
3. Run the **cat properties.properties** command to query the current configuration file of the Flume client.
4. Check whether the **properties.properties** file is correctly configured according to the configuration description of the Flume agent.
   - If yes, go to **Step 3.5**.
   - If no, go to **Step 4**.
5. Modify the **properties.properties** configuration file.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.81 ALM-24004 Flume Fails to Read Data (For MRS 2.*x* or Earlier)

Description

The alarm module monitors the Flume source status. This alarm is generated if the duration that Flume source fails to read data exceeds the threshold.

Users can modify the threshold as required.

This alarm is cleared if the source reads data successfully.

Attribute

Alarm ID	Alarm Severity	Auto Clear
24004	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
ComponentType	Specifies the component type for which the alarm is generated.
ComponentName	Specifies the component name for which the alarm is generated.

Impact on the System

Data collection is stopped.

Possible Causes

- The Flume source is faulty.

- The network is faulty.

Procedure

Step 1 Check whether the Flume source is normal.

1. Check whether the Flume source is the spoolDir type.
   - If yes, go to **Step 1.2**.
   - If no, go to **Step 1.3**.
2. Query the **spoolDir** directory and check whether all files have been sent.
   - If yes, no further action is required.
   - If no, go to **Step 1.5**.
3. Check whether the Flume source is the Kafka type.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 1.5**.
4. Log in to the Kafka client and run the following commands to check whether all topic data configured for the Kafka source has been consumed.

   cd /opt/client/Kafka/kafka/bin

   ./kafka-consumer-groups.sh --bootstrap-server *Kafka cluster IP address*:**21007 --new-consumer --describe --group example-group1 --command-config**

   ../config/consumer.properties
   - If yes, no further action is required.
   - If no, go to **Step 1.5**.
5. Go to the cluster details page and click **Components**.
6. Choose **Flume** > **Instances**.
7. Click the Flume instance of the faulty node and check whether the value of the **Source Speed Metrics** is 0.
   - If yes, go to **Step 2.1**.
   - If no, no further action is required.

Step 2 Check the status of the network between the Flume source and faulty node.

1. Check whether the Flume source is the avro type.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
2. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:

   sudo su - root
3. Run the **ping** *Flume source IP address* command to check whether the Flume source can be pinged.
   - If yes, go to **Step 3**.
   - If no, go to **Step 2.4**.
4. Contact the network administrator to repair the network.
5. Wait for a while and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.82 ALM-24005 Data Transmission by Flume Is Abnormal (For MRS 2.*x* or Earlier)

Description

The alarm module monitors the capacity of Flume channels. This alarm is generated if the duration that a channel is full or the number of times that a source fails to send data to the channel exceeds the threshold.

Users can set the threshold as required by modifying the **channelfullcount** parameter.

This alarm is cleared after the Flume channel space is released.

Attribute

Alarm ID	Alarm Severity	Auto Clear
24005	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
ComponentType	Specifies the component type for which the alarm is generated.
ComponentName	Specifies the component name for which the alarm is generated.

Impact on the System

If the usage of the Flume channel continues to grow, the data transmission time increases. When the usage reaches 100%, the Flume agent process is suspended.

Possible Causes

- The Flume sink is faulty.
- The network is faulty.

Procedure

Step 1 Check whether the Flume sink is normal.

1. Check whether the Flume sink is the HDFS type.
   - If yes, go to **Step 1.2**.
   - If no, go to **Step 1.3**.
2. On MRS Manager, check whether the ALM-14000 HDFS Service Unavailable alarm is reported and whether the HDFS service is stopped.
   - If the alarm is reported, clear it according to the handling suggestions of ALM-14000 HDFS Service Unavailable; if the HDFS service is stopped, start it. Then go to **Step 1.7**.
   - If no, go to **Step 1.7**.
3. Check whether the Flume sink is the HBase type.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 1.7**.
4. On MRS Manager, check whether the ALM-19000 HBase Service Unavailable alarm is reported and whether the HBase service is stopped.
   - If the alarm is reported, clear it according to the handling suggestions of "ALM-19000 HBase Service Unavailable"; if the HBase service is stopped, start it. Then go to **Step 1.7**.
   - If no, go to **Step 1.7**.
5. Check whether the Flume sink is the Kafka type.
   - If yes, go to **Step 1.6**.
   - If no, go to **Step 1.7**.
6. On MRS Manager, check whether the ALM-38000 Kafka Service Unavailable alarm is reported and whether the Kafka service is stopped.
   - If the alarm is reported, clear it according to the handling suggestions of "ALM-38000 Kafka Service Unavailable"; if the Kafka service is stopped, start it. Then go to **Step 1.7**.
   - If no, go to **Step 1.7**.
7. Go to the MRS cluster details page and click **Components**.
8. Choose **Flume** > **Instances**.
9. Click the Flume instance of the faulty node and check whether the value of the **Sink Speed Metrics** is 0.
   - If yes, go to **Step 2.1**.

- If no, no further action is required.

Step 2 Check the status of the network between the Flume sink and faulty node.

1. Check whether the Flume sink is the Avro type.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
2. Log in to the host where the faulty node resides. Run the following command to switch to user **root**:

   sudo su - root
3. Run the **ping** *Flume sink IP address* command to check whether the Flume sink can be pinged.
   - If yes, go to **Step 3**.
   - If no, go to **Step 2.4**.
4. Contact the network administrator to repair the network.
5. Wait for a while and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.83 ALM-25000 LdapServer Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the LdapServer service status every 30 seconds. This alarm is generated when the active and standby LdapServer services are abnormal.

This alarm is cleared when either of the LdapServer services restores.

Attribute

Alarm ID	Alarm Severity	Auto Clear
25000	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

When this alarm is generated, no operation can be performed for the KrbServer users and LdapServer users in the cluster. For example, users, user groups, or roles cannot be added, deleted, or modified, and user passwords cannot be changed on MRS Manager. The authentication for existing users in the cluster is not affected.

Possible Causes

- The node where the LdapServer service locates is faulty.
- The LdapServer process is abnormal.

Procedure

Step 1 Check whether the nodes where the two SlapdServer instances of the LdapServer service locate are faulty.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **LdapServer** > **Instances**. Go to the LdapServer instance page to obtain the host name of the node where the two SlapdServer instances reside.
3. On the **Alarms** page of MRS Manager, check whether the alarm ALM-12006 Node Fault is generated.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2.1**.
4. Check whether the host name in the alarm information is the same as the actual host name in **Step 1.2**.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2.1**.
5. Rectify the fault by following steps provided in ALM-12006 Node Fault.
6. In the alarm list, check whether the alarm ALM-25000 LdapServer Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 2 Check whether the LdapServer process is in normal state.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether ALM-12007 Process Fault is generated.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
3. Check whether the service name and host name in the alarm are consistent with the LdapServer service and host names.
   - If yes, go to **Step 2.4**.
   - If no, go to **Step 3**.
4. Rectify the fault by following steps provided in ALM-12007 Process Fault.
5. In the alarm list, check whether the alarm ALM-25000 LdapServer Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.84 ALM-25004 Abnormal LdapServer Data Synchronization (For MRS 2.*x* or Earlier)

Description

This alarm is generated when LdapServer data on Manager is inconsistent. This alarm is cleared when the data becomes consistent.

This alarm is generated when LdapServer data in the cluster is inconsistent with LdapServer data on Manager. This alarm is cleared when the data becomes consistent.

Attribute

Alarm ID	Alarm Severity	Auto Clear
25004	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

LdapServer data inconsistency occurs because LdapServer data on Manager or in the cluster is damaged. The LdapServer process with damaged data cannot provide services externally, and the authentication functions of Manager and the cluster are affected.

Possible Causes

- The network of the node where the LdapServer process locates is faulty.
- The LdapServer process is abnormal.
- The OS restart damages data on LdapServer.

Procedure

Step 1 Check whether the network where the LdapServer nodes reside is faulty.

1. Go to the cluster details page and choose **Alarms**.
2. Record the IP address of **HostName** in **Location** of the alarm as **IP1** (if multiple alarms exist, record the IP addresses as **IP1**, **IP2**, and **IP3** respectively).
3. Contact O&M personnel and use PuTTY to log in to the node corresponding to **IP1**. Run the **ping** command on the node to check whether the IP address of the management plane of the active OMS node can be pinged.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2.1**.
4. Contact O&M personnel to recover the network and check whether the alarm **ALM-25004 Abnormal LdapServer Data Synchronization** is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check whether the LdapServer process is in normal state.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether ALM-12004 OLdap Resource Is Abnormal is generated for LdapServer.
   - If yes, go to **Step 2.3**.

- If no, go to **Step 2.5**.

3. Rectify the fault by following steps provided in **ALM-12004 OLdap Resource Is Abnormal**.
4. Check whether the alarm ALM-25004 Abnormal LdapServer Data Synchronization is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.5**.
5. On the **Alarms** page of MRS Manager, check whether the alarm ALM-12007 Process Fault of LdapServer is generated.
   - If yes, go to **Step 2.6**.
   - If no, go to **Step 3.1**.
6. Rectify the fault by following steps provided in ALM-12007 Process Fault.
7. Check whether the alarm ALM-25004 Abnormal LdapServer Data Synchronization is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3.1**.

Step 3 Check whether the OS restart damages data on LdapServer.

1. Go to the cluster details page and choose **Alarms**.
2. Record the IP address of **HostName** in **Location** of the alarm as **IP1** (if multiple alarms exist, record the IP addresses as **IP1**, **IP2**, and **IP3** respectively). Choose **Services** > **LdapServer** > **Service Configuration** and record the LdapServer port number as **PORT**. (If the IP address in the alarm location information is the IP address of the standby OMS node, the default port number is 21750.)
3. Log in to node **IP1** as user **omm** and run the **ldapsearch -H ldaps://IP1:PORT -x -LLL -b dc=hadoop,dc=com** command (if the IP address is the IP address of the standby OMS node, run the **ldapsearch -H ldaps://IP1:PORT -x -LLL -b dc=hadoop,dc=com** command before running this command). Check whether error information is displayed in the command output.
   - If yes, go to **Step 3.4**.
   - If no, go to **Step 4**.
4. Recover the LdapServer and OMS nodes using backup data before the alarm is generated. For details, see section "Recovering Manager Data" in the *Administrator Guide*.

   NOTE

   Use the OMS data and LdapServer data backed up at the same time to restore data. Otherwise, the service and operation may fail. To recover data when services run properly, you are advised to manually back up the latest management data and then recover the data. Otherwise, Manager data produced between the backup point in time and the recovery point in time will be lost.
5. Check whether the alarm ALM-25004 Abnormal LdapServer Data Synchronization is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.85 ALM-25500 KrbServer Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the KrbServer service status every 30 seconds. This alarm is generated when the KrbServer service is abnormal.

This alarm is cleared when the KrbServer service is in normal state.

Attribute

Alarm ID	Alarm Severity	Auto Clear
25500	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

When this alarm is generated, no operation can be performed for the KrbServer component in the cluster. The authentication of KrbServer in other components will be affected. The health status of components that depend on KrbServer in the cluster is **Bad**.

Possible Causes

- The node where the KrbServer service locates is faulty.

- The OLdap service is unavailable.

Procedure

Step 1 Check whether the node where the KrbServer service locates is faulty.

1. Go to the MRS cluster details page and click **Components**.
2. Choose **KrbServer** > **Instances**. Go to the KrbServer instance page and view the host name of the node where the KrbServer service is deployed.
3. On the **Alarms** page of MRS Manager, check whether the alarm ALM-12006 Node Fault is generated.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2.1**.
4. Check whether the host name in the alarm information is the same as the actual host name in **Step 1.2**.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2.1**.
5. Rectify the fault by following steps provided in ALM-12006 Node Fault.
6. In the alarm list, check whether the alarm ALM-25500 KrbServer Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 2 Check whether the OLdap service is unavailable.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether ALM-12004 OLdap Resource Is Abnormal is generated.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 3**.
3. Rectify the fault by following steps provided in ALM-12004 OLdap Resource Is Abnormal.
4. In the alarm list, check whether the alarm ALM-25500 KrbServer Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 3**.

Step 3 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.86 ALM-26051 Storm Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Storm service availability every 30 seconds. This alarm is generated if the Storm service becomes unavailable after all Nimbus nodes in a cluster become abnormal.

This alarm is cleared after the Storm service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
26051	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The cluster cannot provide the Storm service.
- Users cannot run new Storm tasks.

Possible Causes

- The Kerberos component is faulty.
- ZooKeeper is faulty or suspended.
- The active and standby Nimbus nodes in the Storm cluster are abnormal.

Procedure

Step 1 Check the Kerberos component status. For clusters without Kerberos authentication, skip this step and go to **Step 2**.

1. Go to the MRS cluster details page and click **Components**.
2. Check whether the health status of the Kerberos service is **Good**.

- If yes, go to **Step 2.1**.
- If no, go to **Step 1.3**.

3. Rectify the fault by following instructions in ALM-25500 KrbServer Service Unavailable.
4. Perform **Step 1.2** again.

Step 2 Check the ZooKeeper component status.

1. Check whether the health status of the ZooKeeper service is **Good**.
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 2.2**.
2. If the ZooKeeper service is stopped, start it. For other problems, follow the instructions in ALM-13000 ZooKeeper Service Unavailable.
3. Perform **Step 2.1** again.

Step 3 Check the status of the active and standby Nimbus nodes.

1. Choose **Components** > **Storm** > **Nimbus**.
2. In **Role**, check whether only one active Nimbus node exists.
   - If yes, go to **Step 4**.
   - If no, go to **Step 3.3**.
3. Select the two Nimbus instances and choose **More** > **Restart Instance**. Check whether the restart is successful.
   - If yes, go to **Step 3.4**.
   - If no, go to **Step 4**.
4. Log in to MRS Manager again and choose **Components** > **Storm** > **Nimbus**. Check whether the health status of Nimbus is **Good**.
   - If yes, go to **Step 3.5**.
   - If no, go to **Step 4**.
5. Wait 30 seconds and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.87 ALM-26052 Number of Available Supervisors in Storm Is Lower Than the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the number of supervisors every 60 seconds and compares it with the threshold. This alarm is generated if the number of supervisors is lower than the threshold.

To modify the threshold, users can choose **System** > **Threshold Configuration** on MRS Manager.

This alarm is cleared if the number of supervisors is greater than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
26052	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

- Existing tasks in the cluster cannot be executed.
- The cluster can receive new Storm tasks but cannot execute them.

Possible Causes

Supervisors are abnormal in the cluster.

Procedure

Step 1 Check the supervisor status.

1. Go to the cluster details page and click **Components**.
2. Choose **Storm** > **Supervisor**.
3. In **Role**, check whether the cluster has supervisor instances that are in the **Faulty** or **Recovering** state.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2**.
4. Select the supervisor instances that are in the **Faulty** or **Recovering** state and choose **More** > **Restart Instance**.
   - If yes, go to **Step 1.5**.
   - If the restart fails, go to **Step 2**.
5. Wait 30 seconds and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.88 ALM-26053 Slot Usage of Storm Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the slot usage of Storm every 60 seconds and compares it with the threshold. This alarm is generated if the slot usage exceeds the threshold.

To modify the threshold, users can choose **System** > **Threshold Configuration** on MRS Manager.

This alarm is cleared if the slot usage is lower than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
26053	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Users cannot run new Storm tasks.

Possible Causes

- Supervisors are abnormal in the cluster.
- Supervisors are normal but have poor processing capability.

Procedure

Step 1 Check the supervisor status.

1. Go to the cluster details page and click **Components**.
2. Choose **Storm** > **Supervisor**.
3. In **Role**, check whether the cluster has supervisor instances that are in the **Faulty** or **Recovering** state.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2.1** or **Step 3.1**.
4. Select the supervisor instances that are in the **Faulty** or **Recovering** state and choose **More** > **Restart Instance**.
   - If yes, go to **Step 1.5**.
   - If the restart fails, go to **Step 4**.
5. Wait a moment and then check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2.1** or **Step 3.1**.

Step 2 Increase the number of slots for the supervisors.

1. Go to the cluster details page and click **Components**.
2. Choose **Storm** > **Supervisor** > **Service Configuration**, and set **Type** to **All**.
3. Increase the value of **supervisor.slots.ports** to increase the number of slots for each supervisor. Then restart the instances.

4. Wait a moment and then check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 3 Expand the capacity of the supervisors.

1. Add nodes.
2. Wait a moment and then check whether the alarm is cleared.
   - If yes, no further action is required.
   - If the restart fails, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.89 ALM-26054 Heap Memory Usage of Storm Nimbus Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the heap memory usage of Storm Nimbus every 30 seconds and compares it with the threshold. This alarm is generated if the heap memory usage exceeds the threshold (80% by default).

To modify the threshold, users can choose **System** > **Threshold Configuration** > **Service** > **Storm** on MRS Manager.

This alarm is cleared if the heap memory usage is lower than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
26054	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Frequent memory garbage collection or memory overflow may occur, affecting submission of Storm services.

Possible Causes

The heap memory usage is high or the heap memory is improperly allocated.

Procedure

Step 1 Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Choose **ALM-26054 Heap Memory Usage of Storm Nimbus Exceeds the Threshold > Location**. Query the **HostName** of the alarmed instance.
3. Choose **Components > Storm > Instances > Nimbus (corresponding to the HostName of the alarmed instance) > Customize > Heap Memory Usage of Nimbus**.
4. Check whether the heap memory usage of Nimbus has reached the threshold (80%).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Adjust the heap memory.

   Choose **Components > Storm > Service Configuration**, and set **Type** to **All**. Choose **Nimbus** > **System**. Increase the value of **-Xmx** in **NIMBUS_GC_OPTS**. Click **Save Configuration**. Select **Restart the affected services or instances** and click **OK**.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.90 ALM-27001 DBService Unavailable (For MRS 2.*x* or Earlier)

Description

The alarm module checks the DBService status every 30 seconds. This alarm is generated when the system detects that DBService is unavailable.

This alarm is cleared when DBService recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
27001	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The database service is unavailable and cannot provide data import and query functions for upper-layer services, which results in service exceptions.

Possible Causes

- The floating IP address does not exist.
- There is no active DBServer instance.
- The active and standby DBServer processes are abnormal.

Procedure

Step 1 Check whether the floating IP address exists in the cluster environment.

1. Go to the MRS cluster details page and click **Components**.

2. Choose **DBService** > **Instances**.
3. Check whether the active instance exists.
   - If yes, go to **Step 1.4**.
   - If no, go to **Step 2.1**.
4. Select the active DBServer instance and record the IP address.
5. Log in to the host with the preceding IP address and run the **ifconfig** command to check whether the DBService floating IP address exists on the node.
   - If yes, go to **Step 1.6**.
   - If no, go to **Step 2.1**.
6. Run the **ping** *floating IP address* command to check whether the DBService floating IP address can be pinged.
   - If yes, go to **Step 1.7**.
   - If no, go to **Step 2.1**.
7. Log in to the host where the DBService floating IP address is located and run the **ifconfig** *interface* **down** command to delete the floating IP address.
8. Choose **Components** > **DBService** > **More** > **Restart Service** to restart DBService and check whether DBService is started successfully.
   - If yes, go to **Step 1.9**.
   - If no, go to **Step 2.1**.
9. Wait about 2 minutes and check whether the alarm is cleared in the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 13**.

Step 2 Check the status of the active DBServer instance.

1. Select the DBServer instance whose role status is abnormal and record the IP address.
2. On the **Alarms** page, check whether ALM-12007 Process Fault occurs in the DBServer instance on the host that corresponds to the IP address.
   - If yes, go to **Step 2.3**.
   - If no, go to **Step 4**.
3. Rectify the fault by following steps provided in ALM-12007 Process Fault.
4. Wait about 5 minutes and check whether the alarm is cleared in the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 3 Check the status of the active and standby DBServers.

1. Log in to the host where the DBService floating IP address is located, run the **sudo su - root** and **su - omm** commands to switch to user **omm**, and run the **cd ${BIGDATA_HOME}/FusionInsight/dbservice/** command to go to the DBService installation directory.
2. Run the **sh sbin/status-dbserver.sh** command to view the status of the active and standby HA processes of DBService. Determine whether the status can be viewed successfully.

- If yes, go to **Step 3.3**.
- If no, go to **Step 4**.

3. Check whether the active and standby HA processes are abnormal.
   - If yes, go to **Step 3.4**.
   - If no, go to **Step 4**.
4. Choose **Components** > **DBService** > **More** > **Restart Service** to restart DBService and check whether DBService is started successfully.
   - If yes, go to **Step 3.5**.
   - If no, go to **Step 4**.
5. Wait about 2 minutes and check whether the alarm is cleared in the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.91 ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the active or standby DBService node does not receive heartbeat messages from the peer node.

This alarm is cleared when the heartbeat recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
27003	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local DBService HA Name	Specifies a local DBService HA.
Peer DBService HA Name	Specifies a peer DBService HA.

Impact on the System

During the DBService heartbeat interruption, only one node can provide the service. If this node is faulty, no standby node is available for failover and the service is unavailable.

Possible Causes

The link between the active and standby DBService nodes is abnormal.

Procedure

Step 1 Check whether the network between the active and standby DBService servers is in normal state.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, locate the row that contains the alarm and view the IP address of the standby DBService server in the alarm details.
3. Log in to the active DBService server.
4. Run the **ping** *heartbeat IP address of the standby DBService* command to check whether the standby DBService server is reachable.
   - If yes, go to **Step 2**.
   - If no, go to **Step 1.5**.
5. Contact the network administrator to check whether the network is faulty.
   - If yes, go to **Step 1.6**.
   - If no, go to **Step 2**.
6. Rectify the network fault and check whether the alarm is cleared from the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.92 ALM-27004 Data Inconsistency Between Active and Standby DBServices (For MRS 2.*x* or Earlier)

Description

The system checks the data synchronization status between the active and standby DBServices every 10 seconds. This alarm is generated when the synchronization status cannot be queried for six consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the synchronization is in normal state.

Attribute

Alarm ID	Alarm Severity	Auto Clear
27004	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local DBService HA Name	Specifies a local DBService HA.
Peer DBService HA Name	Specifies a peer DBService HA.
SYNC_PERSENT	Synchronization percentage.

Impact on the System

When data is not synchronized between the active and standby DBServices, the data may be lost or abnormal if the active instance becomes abnormal.

Possible Causes

- The network between the active and standby nodes is unstable.
- The standby DBService is abnormal.

- The disk space of the standby node is full.

Procedure

Step 1 Check whether the network between the active and standby nodes is in normal state.

1. Go to the cluster details page and choose **Alarms**.
2. In the alarm list, locate the row that contains the alarm and view the IP address of the standby DBService node in the alarm details.
3. Log in to the active DBService node.
4. Run the **ping** *heartbeat IP address of the standby DBService* command to check whether the standby DBService node is reachable.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.5**.
5. Contact the O&M personnel to check whether the network is faulty.
   - If yes, go to **Step 1.6**.
   - If no, go to **Step 2.1**.
6. Rectify the network fault and check whether the alarm is cleared from the alarm list.
   - If yes, no further action is required.
   - If no, go to **Step 2.1**.

Step 2 Check whether the standby DBService is in normal state.

1. Log in to the standby DBService node.
2. Run the following commands to switch the user:

   sudo su - root

   su - omm
3. Go to the **${DBSERVER_HOME}/sbin** directory and run the **./status-dbserver.sh** command to check whether the GaussDB resource status of the standby DBService is in normal state. In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:

   Example:

   ```
   10_10_10_231 gaussDB Standby_normal Normal Active_standby
   ```
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 4**.

Step 3 Check whether the disk space of the standby node is insufficient.

1. Log in to the standby DBService node.
2. Run the following commands to switch the user:

   sudo su - root

   su - omm
3. Go to the **${DBSERVER_HOME}** directory, and run the following commands to obtain the DBService data directory:

   cd ${DBSERVER_HOME}

   source .dbservice_profile

echo ${DBSERVICE_DATA_DIR}

4. Run the **df -h** command to check the system disk partition usage.
5. Check whether the DBService data directory space is full.
   - If yes, go to **Step 3.6**.
   - If no, go to **Step 4**.
6. Perform upgrade and expand capacity.
7. After capacity expansion, wait 2 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.93 ALM-28001 Spark Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Spark service status every 30 seconds. This alarm is generated when the Spark service is unavailable.

This alarm is cleared when the Spark service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
28001	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The Spark tasks submitted by users fail to be executed.

Possible Causes

- The KrbServer service is abnormal.
- The LdapServer service is abnormal.
- The ZooKeeper service is abnormal.
- The HDFS service is abnormal.
- The Yarn service is abnormal.
- The corresponding Hive service is abnormal.

Procedure

Step 1 Check whether service unavailability alarms exist in services that Spark depends on.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Check whether the following alarms exist in the alarm list:
   a. ALM-25500 KrbServer Service Unavailable
   b. ALM-25000 LdapServer Service Unavailable
   c. ALM-13000 ZooKeeper Service Unavailable
   d. ALM-14000 HDFS Service Unavailable
   e. ALM-18000 Yarn Service Unavailable
   f. ALM-16004 Hive Service Unavailable
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 2**.
3. Handle the alarms based on the troubleshooting methods provided in the alarm help.

   After the alarm is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.94 ALM-38000 Kafka Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Kafka service availability every 30 seconds. This alarm is generated if the Kafka service becomes unavailable.

This alarm is cleared after the Kafka service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
38000	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The cluster cannot provide the Kafka service and users cannot run new Kafka tasks.

Possible Causes

- The KrbServer component is faulty.
- The ZooKeeper component is faulty or fails to respond.
- The Broker node in the Kafka cluster is abnormal.

Procedure

Step 1 Check the KrbServer component status. For clusters without Kerberos authentication, skip this step and go to **Step 2**.

1. Go to the MRS cluster details page and click **Components**.
2. Check whether the health status of the KrbServer service is **Good**.
   - If yes, go to **Step 2.1**.
   - If no, go to **Step 1.3**.
3. Rectify the fault by following instructions in ALM-25500 KrbServer Service Unavailable.
4. Perform **Step 1.2** again.

Step 2 Check the ZooKeeper component status.

1. Check whether the health status of the ZooKeeper service is **Good**.
   - If yes, go to **Step 3.1**.
   - If no, go to **Step 2.2**.
2. If the ZooKeeper service is stopped, start it. For other problems, follow the instructions in ALM-13000 ZooKeeper Service Unavailable.
3. Perform **Step 2.1** again.

Step 3 Check the Broker status.

1. Choose **Components** > **Kafka** > **Broker**.
2. In **Role**, check whether all instances are normal.
   - If yes, go to **Step 3.4**.
   - If no, go to **Step 3.3**.
3. Select all instances of Broker and choose **More** > **Restart Instance**.
   - If the restart is successful, go to **Step 3.4**.
   - If the restart fails, go to **Step 4**.
4. Choose **Components > Kafka**. Check whether the health status of Kafka is **Good**.
   - If yes, go to **Step 3.5**.
   - If no, go to **Step 4**.
5. Wait 30 seconds and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 4**.

Step 4 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.95 ALM-38001 Insufficient Kafka Disk Capacity (For MRS 2.*x* or Earlier)

Description

The system checks the Kafka disk usage every 60 seconds and compares it with the threshold. This alarm is generated if the disk usage exceeds the threshold.

To modify the threshold, users can choose **System** > **Threshold Configuration** on MRS Manager.

This alarm is cleared if the Kafka disk usage is lower than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
38001	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition where the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Kafka fails to write data to the disks.

Possible Causes

- The Kafka disk configurations (such as disk count and disk size) are insufficient for the data volume.
- The data retention period is long and historical data occupies large space.
- Services are improperly planned. As a result, data is unevenly distributed and some disks are full.

Procedure

Step 1 Go to the MRS cluster details page and choose **Alarms**.

Step 2 In the alarm list, click the alarm and view the **HostName** and **PartitionName** of the alarm in **Location** of **Alarm Details**.

Step 3 On the **Hosts** page, click the host name obtained in **Step 2**.

Step 4 Check whether the **Disk** area contains the **PartitionName** of the alarm.

- If yes, go to **Step 5**.
- If no, manually clear the alarm and no further action is required.

Step 5 In the **Disk** area, check whether the usage of the alarmed partition has reached 100%.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 In **Instance**, choose **Broker > Instance Configuration**. On the **Instance Configuration** page that is displayed, set **Type** to **All** and query the data directory parameter **log.dirs**.

Step 7 Choose **Components > Kafka > Instances**. On the **Kafka Instance** page that is displayed, stop the Broker instance corresponding to **Step 2**. Then log in to the alarmed node and manually delete the data directory in **Step 6**. After all subsequent operations are complete, start the Broker instance.

Step 8 Choose **Components > Kafka > Service Configuration**. The **Kafka Configuration** page is displayed.

Step 9 Check whether **disk.adapter.enable** is **true**.

- If yes, go to **Step 11**.
- If no, change the value to **true** and go to **Step 10**.

Step 10 Check whether the **adapter.topic.min.retention.hours** parameter, indicating the minimum data retention period, is properly configured.

- If yes, go to **Step 12**.
- If no, set it to a proper value and go to **Step 12**.

NOTE

If the retention period cannot be adjusted for certain topics, the topics can be added to **disk.adapter.topic.blacklist**.

Step 11 Wait 10 minutes and check whether the disk usage is reduced.

- If yes, wait until the alarm is cleared.
- If no, go to **Step 12**.

Step 12 Go to the **Kafka Topic Monitor** page and query the data retention period configured for Kafka. Determine whether the retention period needs to be shortened based on service requirements and data volume.

- If yes, go to **Step 13**.
- If no, go to **Step 14**.

Step 13 Find the topics with great data volumes based on the disk partition obtained in **Step 2**. Log in to the Kafka client and manually shorten the data retention period for these topics using the following command:

kafka-topics.sh --zookeeper *ZooKeeper address:24002/kafka* **--alter --topic** *Topic name* **--config retention.ms=***Retention period*

Step 14 Check whether partitions are properly configured for topics. For example, if the number of partitions for a topic with a large data volume is smaller than the number of disks, data may be unevenly distributed to the disks and the usage of some disks will reach the upper limit.

NOTE

To identify topics with great data volumes, log in to the relevant nodes that are obtained in **Step 2**, go to the data directory (the directory before **log.dirs** in **Step 6** is modified), and check the disk space occupied by the partitions of the topics.

- If the partitions are improperly configured, go to **Step 15**.
- If the partitions are properly configured, go to **Step 16**.

Step 15 On the Kafka client, add partitions to the topics.

kafka-topics.sh --zookeeper *ZooKeeper address:24002/kafka* **--alter --topic** *Topic name* **--partitions=***Number of new partitions*

NOTE

It is advised to set the number of new partitions to a multiple of the number of Kafka disks.

This operation may not quickly clear the alarm. Data will be gradually balanced among the disks.

Step 16 Check whether the cluster capacity needs to be expanded.

- If yes, add nodes to the cluster and go to **Step 17**.
- If no, go to **Step 17**.

Step 17 Wait a moment and then check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Step 18 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.96 ALM-38002 Heap Memory Usage of Kafka Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the heap memory usage of Kafka every 30 seconds. This alarm is generated if the heap memory usage of Kafka exceeds the threshold (80%).

This alarm is cleared if the heap memory usage is lower than the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
38002	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

Memory overflow may occur, causing service crashes.

Possible Causes

The heap memory usage is high or the heap memory is improperly allocated.

Procedure

Step 1 Check the heap memory usage.

1. Go to the MRS cluster details page and choose **Alarms**.
2. Choose **ALM-38002 Kafka Heap Memory Usage Exceeds the Threshold** > **Location**. Query the IP address of the alarmed instance.

3. Choose **Components > Kafka > Instance > Broker (corresponding to the IP address of the alarmed instance) > Customize > Kafka Heap Memory Resource Percentage** to check the heap memory usage.
4. Check whether the heap memory usage of Kafka has reached the threshold (80%).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Kafka** > **Service Configuration** > **All** > **Broker** > **Environment Variables**. Increase the value of **KAFKA_HEAP_OPTS** as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Related Information

N/A

9.5.97 ALM-43001 Spark Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Spark service status every 60 seconds. This alarm is generated when the Spark service is unavailable.

This alarm is cleared when the Spark service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43001	Critical	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The Spark tasks submitted by users fail to be executed.

Possible Causes

- The KrbServer service is abnormal.
- The LdapServer service is abnormal.
- ZooKeeper is abnormal.
- The HDFS service is abnormal.
- The Yarn service is abnormal.
- The corresponding Hive service is abnormal.

Procedure

Step 1 Check whether service unavailability alarms exist in services that Spark depends on.

1. Go to the cluster details page and choose **Alarms**.
2. Check whether the following alarms exist in the alarm list:
   a. ALM-25500 KrbServer Service Unavailable
   b. ALM-25000 LdapServer Service Unavailable
   c. ALM-13000 ZooKeeper Service Unavailable
   d. ALM-14000 HDFS Service Unavailable
   e. ALM-18000 Yarn Service Unavailable
   f. ALM-16004 Hive Service Unavailable
   - If yes, go to **Step 1.3**.
   - If no, go to **Step 2**.
3. Handle the alarm according to the alarm help.
   After the alarm is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.

2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.98 ALM-43006 Heap Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the JobHistory process status every 30 seconds. The alarm is generated when the heap memory usage of the JobHistory process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43006	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JobHistory process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the JobHistory process is overused or the heap memory is inappropriately allocated.

Procedure

Step 1 Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43006** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JobHistory** (IP address of the instance for which the alarm is generated) > **Customize** > **Heap Memory Statistics of the JobHistory Process**. Click **OK** to view the heap memory usage.
4. Check whether the used heap memory of JobHistory reaches 90% of the maximum heap memory specified for JobHistory.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JobHistory** > **Default**. Increase the value of **SPARK_DAEMON_MEMORY** as required.
6. Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK**.
7. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.99 ALM-43007 Non-Heap Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the JobHistory process status every 30 seconds. The alarm is generated when the non-heap memory usage of the JobHistory process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43007	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JobHistory process non-heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The non-heap memory of the JobHistory process is overused or the non-heap memory is inappropriately allocated.

Procedure

Step 1 Check non-heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43007** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JobHistory** (IP address of the instance for which the alarm is generated) > **Customize** > **Non-Heap Memory Statistics of the JobHistory Process**. Click **OK** to view the non-heap memory usage.
4. Check whether the non-heap memory usage of JobHistory has reached the threshold (90% of the maximum memory).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JobHistory** > **Default**. Increase the value of **-XX:MaxMetaspaceSize** in **SPARK_DAEMON_JAVA_OPTS** as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.100 ALM-43008 Direct Memory Usage of the JobHistory Process Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the JobHistory process status every 30 seconds. The alarm is generated when the direct memory usage of the JobHistory process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43008	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JobHistory process direct memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the JobHistory process is overused or the direct memory is inappropriately allocated.

Procedure

Step 1 Check the direct memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43008** and view the IP address and role name of the instance in **Location**.

3. Choose **Components** > **Spark** > **Instance** > **JobHistory** (IP address of the instance for which the alarm is generated) > **Customize** > **Direct Memory Statistics of the JobHistory Process**. Click **OK** to view the direct memory usage.
4. Check whether the direct memory usage of the JobHistory process has reached the threshold (90% of the maximum direct memory).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JobHistory** > **Default**. Increase the value of **-XX:MaxDirectMemorySize** in **SPARK_DAEMON_JAVA_OPTS** as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.101 ALM-43009 JobHistory GC Time Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the GC time of the JobHistory process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (12 seconds) for three consecutive times. You can change the threshold by choosing **System** > **Threshold Configuration** > **Service** > **Spark** > **JobHistory GC Time** > **Total JobHistory GC Time**. This alarm is cleared when the JobHistory GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43009	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the GC time exceeds the threshold, JobHistory may run in low performance.

Possible Causes

The heap memory of the JobHistory process is overused or inappropriately allocated, causing frequent GC.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43009** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JobHistory** (IP address of the instance for which the alarm is generated) > **Customize** > **GC Time of the JobHistory Process**. Click **OK** to view the GC time.
4. Check whether the GC time of the JobHistory process is longer than 12 seconds.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JobHistory** > **Default**. Increase the value of the **SPARK_DAEMON_MEMORY** parameter as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.102 ALM-43010 Heap Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the JDBCServer process status every 30 seconds. The alarm is generated when the heap memory usage of the JDBCServer process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43010	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JDBCServer process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the JDBCServer process is overused or the heap memory is inappropriately allocated.

Procedure

Step 1 Check the heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43010** and view the IP address and role name of the instance in **Location**.

3. Choose **Components** > **Spark** > **Instance** > **JDBCServer** (IP address of the instance for which the alarm is generated) > **Customize** > **Heap Memory Statistics of the JDBCServer Process**. Click **OK** to view the heap memory usage.
4. Check whether the heap memory usage of JDBCServer has reached the threshold (90% of the maximum heap memory).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JDBCServer** > **Tuning**. Increase the value of the **SPARK_DRIVER_MEMORY** parameter as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.103 ALM-43011 Non-Heap Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the JDBCServer process status every 30 seconds. The alarm is generated when the non-heap memory usage of the JDBCServer process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43011	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JDBCServer process non-heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The non-heap memory of the JDBCServer process is overused or the non-heap memory is inappropriately allocated.

Procedure

Step 1 Check non-heap memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43011** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JDBCServer** (IP address of the instance for which the alarm is generated) > **Customize** > **Non-heap Memory Statistics of the JDBCServer Process**. Click **OK** to view the non-heap memory usage.
4. Check whether the non-heap memory usage of JDBCServer has reached the threshold (90% of the maximum non-heap memory).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JDBCServer** > **Tuning**. Increase the value of **-XX:MaxMetaspaceSize** in **spark.driver.extraJavaOptions** as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.104 ALM-43012 Direct Memory Usage of the JDBCServer Process Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the JDBCServer process status every 30 seconds. The alarm is generated when the direct memory usage of the JDBCServer process exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43012	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the available JDBCServer process direct memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the JDBCServer process is overused or the direct memory is inappropriately allocated.

Procedure

Step 1 Check the direct memory usage.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43012** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JDBCServer** (IP address of the instance for which the alarm is generated) > **Customize** > **Direct Memory Statistics of the JDBCServer Process**. Click **OK** to view the direct memory usage.

4. Check whether the direct memory usage of the JDBCServer process has reached the threshold (90% of the maximum direct memory).
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JDBCServer** > **Tuning**. Increase the value of **-XX:MaxDirectMemorySize** in **spark.driver.extraJavaOptions** as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.105 ALM-43013 JDBCServer GC Time Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system checks the GC time of the JDBCServer process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (12 seconds) for three consecutive times. You can change the threshold by choosing **System** > **Threshold Configuration** > **Service** > **Spark** > **JDBCServer GC Time** > **Total JDBCServer GC Time**. This alarm is cleared when the JDBCServer GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
43013	Major	Yes

Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the GC time exceeds the threshold, JDBCServer may run in low performance.

Possible Causes

The heap memory of the JDBCServer process is overused or inappropriately allocated, causing frequent GC.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **43013** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Spark** > **Instance** > **JDBCServer** (IP address of the instance for which the alarm is generated) > **Customize** > **GC Time of the JDBCServer Process**. Click **OK** to view the GC time.
4. Check whether the GC time of the JDBCServer process is longer than 12 seconds.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Spark** > **Service Configuration**. Set **Type** to **All** and choose **JDBCServer** > **Tuning**. Increase the value of the **SPARK_DRIVER_MEMORY** parameter as required.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.106 ALM-44004 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold (For MRS 2.*x* or Earlier)

Description

This alarm is generated when the system detects that the number of queuing tasks in a resource group exceeds the threshold. The system queries the number of queuing tasks in a resource group through the JMX interface. You can choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Presto** > **resource-groups** to configure a resource group. You can choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Coordinator** > **Customize** > **resourceGroupAlarm** to configure the threshold of each resource group.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44004	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the number of queuing tasks in a resource group exceeds the threshold, a large number of tasks may be in the queuing state. The Presto task time exceeds the expected value. When the number of queuing tasks in a resource group exceeds the maximum number (**maxQueued**) of queuing tasks in the resource group, new tasks cannot be executed.

Possible Causes

The resource group configuration is improper or too many tasks in the resource group are submitted.

Procedure

Step 1 Choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Presto** > **resource-groups** to adjust the resource group configuration.

Step 2 You can choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Coordinator** > **Customize** > **resourceGroupAlarm** to modify the threshold of each resource group.

Step 3 Collect fault information.

1. Log in to the cluster node based on the host name in the fault information and query the number of queuing tasks based on **Resource Group** in the additional information on the Presto client.
2. Log in to the cluster node based on the host name in the fault information, view the **/var/log/Bigdata/nodeagent/monitorlog/monitor.log** file, and search for resource group information to view the monitoring collection information of the resource group.
3. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.107 ALM-44005 Presto Coordinator Process GC Time Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system collects GC time of the Presto Coordinator process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System** > **Configure Alarm Threshold** > **Service** > **Presto** > **Coordinator** > **Presto Process Garbage Collection Time** > **Garbage Collection Time of the Coordinator Process** on MRS Manager. This alarm is cleared when the Coordinator process GC time is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44005	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the GC time of the Coordinator process is too long, the Coordinator process running performance will be affected and the Coordinator process will even be unavailable.

Possible Causes

The heap memory of the Coordinator process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **44005** and view the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Presto** > **Instances** > **Coordinator** (business IP address of the instance for which the alarm is generated) > **Customize** > **Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Coordinator process is longer than 5 seconds.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Presto** > **Service Configuration**, and switch **Basic** to **All**. Choose **Presto** > **Coordinator**. Increase the value of **-Xmx** (maximum heap memory) in the **JAVA_OPTS** parameter based on the site requirements.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.108 ALM-44006 Presto Worker Process GC Time Exceeds the Threshold (For MRS 2.*x* or Earlier)

Description

The system collects GC time of the Presto Worker process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System** >

Configure Alarm Threshold > **Service** > **Presto** > **Worker** > **Presto Garbage Collection Time** > **Garbage Collection Time of the Worker Process** on MRS Manager. This alarm is cleared when the Worker process GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44006	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the GC time of the Worker process is too long, the Worker process running performance will be affected and the Worker process will even be unavailable.

Possible Causes

The heap memory of the Worker process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.
2. Select the alarm whose **Alarm ID** is **44006**. Then check the IP address and role name of the instance in **Location**.
3. Choose **Components** > **Presto** > **Instances** > **Worker** (business IP address of the instance for which the alarm is generated) > **Customize** > **Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Worker process is longer than 5 seconds.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Presto** > **Service Configuration**, and switch **Basic** to **All**, and choose **Presto** > **Worker** Increase the value of **-Xmx** (maximum heap memory) in the **JAVA_OPTS** parameter based on the site requirements.

6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M engineers and send the collected logs.

----End

Reference

None

9.5.109 ALM-45325 Presto Service Unavailable (For MRS 2.*x* or Earlier)

Description

The system checks the Presto service status every 60 seconds. This alarm is generated when the system detects that Presto is unavailable.

This alarm is cleared when the Presto service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45325	Critical	Yes

Parameters

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Presto cannot run SQL queries.

Possible Causes

- The Presto coordinator or worker process is faulty.
- The network communication between Presto coordinator and worker instances is interrupted.

Procedure

Step 1 Check the status of the coordinator and worker processes.

1. Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Presto**. On the page that is displayed, click the **Instance** tab. In the Presto instance list, check whether the status of all coordinator or worker instances is **Unknown**.
   - If yes, go to **2**.
   - If no, go to **1**.
2. In the upper part of the Presto instance list, choose **More** > **Restart Service** to restart the coordinator and worker processes.
3. In the alarm list, check whether ALM-45325 Presto Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **1** in **Step 2**.

Step 2 Collect fault information.

1. On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.
2. Select **Presto** for **Service**.
3. Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.
4. Contact the O&M engineers and send the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

9.6 Object Management

9.6.1 Managing Objects

MRS contains different types of basic objects as described in **Table 9-16**.

Table 9-16 MRS basic object overview

Object	Description	Example
Service	Function set that can complete specific business.	KrbServer service and LdapServer service
Service instance	Specific instance of a service, usually called service.	KrbServer service
Service role	Function entity that forms a complete service, usually called role.	KrbServer is composed of the KerberosAdmin role and KerberosServer role.
Role instance	Specific instance of a service role running on a host.	KerberosAdmin that is running on Host2 and KerberosServer that is running on Host3
Host	An ECS running Linux OS.	Host1 to Host5
Rack	Physical entity that contains multiple hosts connecting to the same switch.	Rack1 contains Host1 to Host5.
Cluster	Logical entity that consists of multiple hosts and provides various services.	Cluster names **Cluster1** consists of five hosts (Host1 to Host5) and provides services such as KrbServer and LdapServer.

9.6.2 Viewing Configurations

On MRS Manager, users can view the configurations of services (including roles) and role instances.

Procedure

- Query service configurations.
  a. On MRS Manager page, click **Services**.
  b. Select the target service from the service list.
  c. Click **Service Configuration**.
  d. Set **Type** to **All**. All configuration parameters of the service are displayed in the navigation tree. The root nodes from top down in the navigation tree represent the service names and role names.
  e. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

     The parameters under the service nodes and role nodes are service configuration parameters and role configuration parameters respectively.
  f. In the **Non-default** parameter, select **Non-default**. The parameters whose values are not default values will be displayed.

- Query role instance configurations.
  a. On MRS Manager page, click **Services**.
  b. Select the target service from the service list.
  c. Click the **Instances** tab.
  d. Click the target role instance from the role instance list.
  e. Click **Instance Configuration**.
  f. Set **Type** to **All**. The navigation tree of all configuration parameters of the role instance is displayed.
  g. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.
  h. In the **Non-default** parameter, select **Non-default**. The parameters whose values are not default values will be displayed.

9.6.3 Managing Services

You can perform the following operations on MRS Manager:

- Start the service in the **Stopped**, **Stop Failed**, or **Start Failed** state to use the service.
- Stop the services or stop abnormal services.
- Restart abnormal services or configure expired services to restore or enable the services.

Procedure

Step 1 On MRS Manager page, click **Services**.

Step 2 Locate the row that contains the target service, **Start**, **Stop**, or **Restart** to start, stop, or restart the service.

Services are interrelated. If a service is started, stopped, and restarted, services dependent on it will be affected.

The services will be affected in the following ways:

- If a service is to be started, the lower-layer services dependent on it must be started first.
- If a service is stopped, the upper-layer services dependent on it are unavailable.
- If a service is restarted, the running upper-layer services dependent on it must be restarted.

----End

9.6.4 Configuring Service Parameters

On MRS Manager, you can view and modify the default service configurations based on site requirements and export or import the configurations.

Impact on the System

- You need to download and update the client configuration files after configuring HBase, HDFS, Hive, Spark, Yarn, and MapReduce service properties.
- The parameters of DBService cannot be modified when only one DBService role instance exists in the cluster.

Procedure

- Modify a service.
  a. Click **Services**.
  b. Select the target service from the service list.
  c. Click **Service Configuration**.
  d. Set **Type** to **All**. All configuration parameters of the service are displayed in the navigation tree. The root nodes from top down in the navigation tree represent the service names and role names.
  e. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

     To cancel the change to a parameter value, click  .

     NOTE

     You can also use host groups to change role instance configurations in batches. Select a role name from the **Role** drop-down list and choose **< Select Host >** in the **Host** drop-down list. Enter a name in the **Host Group Name** text box, select the hosts to be modified from the **Host** list, add them to the **Selected hosts** area, and click **OK**. The added host group can be selected from **Host** and is only valid on the current page. The page cannot be saved after being refreshed.
  f. Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK** to restart the services.

     After **Operation successful.** is displayed, click **Finish**. The service is started successfully.

     NOTE

     To update the queue configuration of the Yarn service without restarting service, choose **More** > **Refresh Queue** to update the queue for the configuration to take effect.
- Export service configuration parameters.
  a. Click **Services**.
  b. Select a service.
  c. Click **Service Configuration**.
  d. Click **Export Service Configuration**. Select a path for saving the configuration files.
- Import service configuration parameters.
  a. Click **Services**.
  b. Select a service.

c. Click **Service Configuration**.

d. Click **Import Service Configuration**.

e. Select the target configuration file.

f. Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK**.

After **Operation successful.** is displayed, click **Finish**. The service is started successfully.

9.6.5 Configuring Customized Service Parameters

Each component of MRS supports all open-source parameters. You can modify some parameters for key application scenarios on MRS Manager. Some component clients may not include all parameters with open-source features. For component parameters that cannot be directly modified on Manager, users can add new parameters for components by using the configuration customization function on Manager. Newly added parameters are saved in component configuration files and take effect after restart.

Impact on the System

- After the service attributes are configured, the service needs to be restarted and cannot be accessed.
- You need to download and update the client configuration files after configuring HBase, HDFS, Hive, Spark, Yarn, and MapReduce service properties.

Prerequisites

You have understood the meanings of parameters to be added, configuration files that have taken effect, and the impact on components.

Procedure

Step 1 On MRS Manager, click **Services**.

Step 2 Select the target service from the service list.

Step 3 Click **Service Configuration**.

Step 4 Set **Type** to **All**.

Figure 9-2 Service Configuration

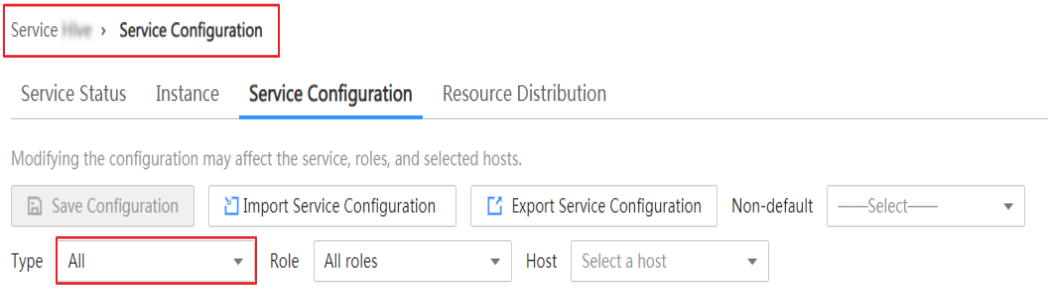

Step 5 In the navigation tree, select **Customization**. The customized parameters of the current component are displayed on Manager.

The configuration files that save the newly added customized parameters are displayed in the **Parameter File** column. Different configuration files may have same open source parameters. After the parameters in different files are set to different values, whether the configuration takes effect depends on the loading sequence of the configuration files by components. You can customize parameters for services and roles as required. Adding customized parameters for a single role instance is not supported.

Figure 9-3 Customization configurations

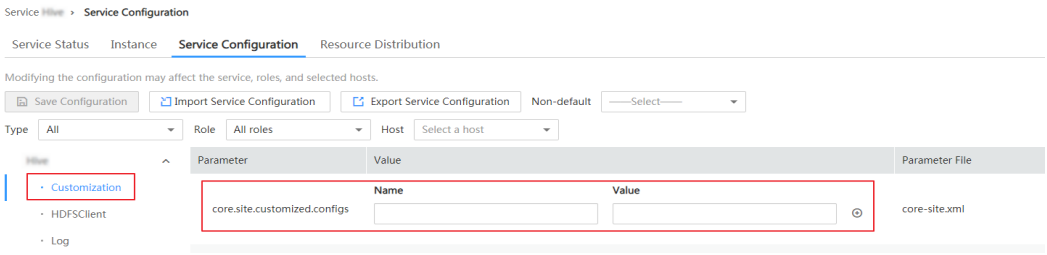

Step 6 Based on the configuration files and parameter functions, locate the row where a specified parameter resides, enter the parameter name supported by the component in the **Name** column and enter the parameter value in the **Value** column.

- You can click  or  to add or delete a user-defined parameter. You can delete a customized parameter only after you click  for the first time.
- If you want to cancel the modification of a parameter value, click  to restore it.

Step 7 Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK** to restart the services.

After **Operation successful.** is displayed, click **Finish**. The service is started successfully.

----End

Task Example

Configuring Customized Hive Parameters

Hive depends on HDFS. By default, Hive accesses the HDFS client. The configuration parameters to take effect are controlled by HDFS in a unified manner. For example, the HDFS parameter **ipc.client.rpc.timeout** affects the RPC timeout period for all clients to connect to the HDFS server. If you need to modify the timeout period for Hive to connect to HDFS, you can use the configuration customization function. After this parameter is added to the **core-site.xml** file of Hive, this parameter can be identified by the Hive service and its configuration overwrites the parameter configuration in HDFS.

Step 1 On MRS Manager, choose **Services** > **Hive** > **Service Configuration**.

Step 2 Set **Type** to **All**.

Figure 9-4 Hive Service Configuration

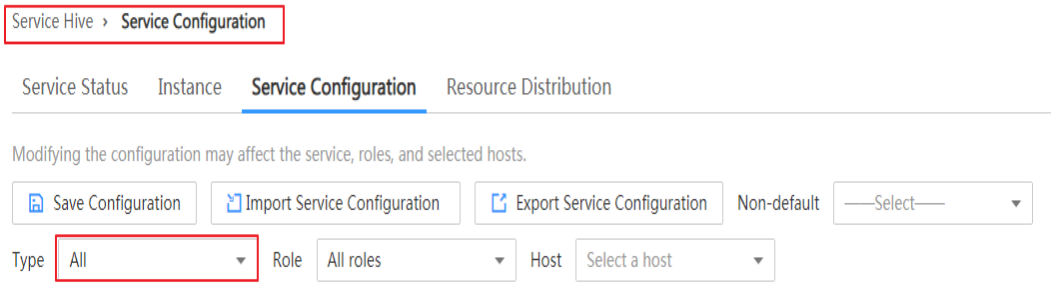

Step 3 In the navigation tree on the left, select **Customization** for the Hive service. The system displays the customized service parameters supported by Hive.

Step 4 In **core-site.xml**, locate the row that contains the **core.site.customized.configs** parameter, enter **ipc.client.rpc.timeout** in the **Name** column, and enter a new value in the **Value** column, for example, **150000**. The unit is millisecond.

Figure 9-5 Hive customization configurations

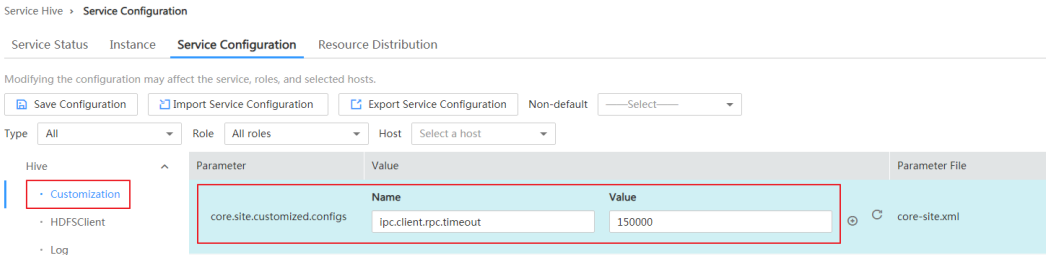

Step 5 Click **Save Configuration** and select **Restart the affected services or instances**. Click **OK** to restart the service.

After **Operation successful.** is displayed, click **Finish**. The service is started successfully.

Figure 9-6 Saving Hive configurations

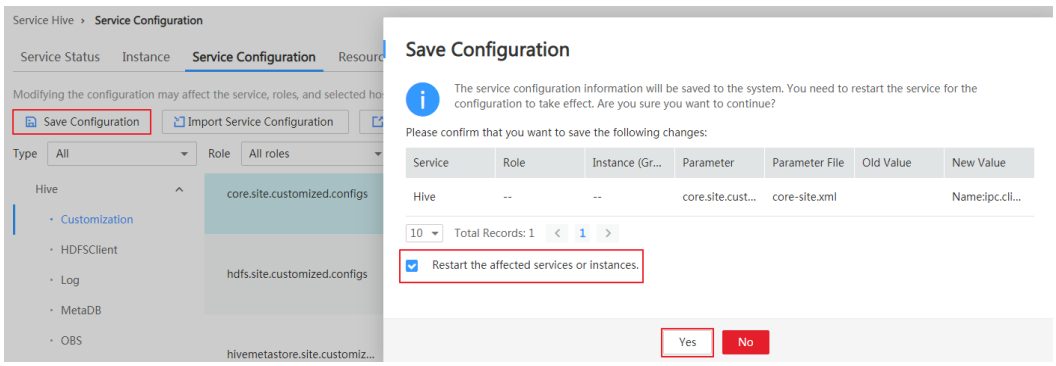

----End

9.6.6 Synchronizing Service Configurations

Scenario

If **Configuration Status** of a service is **Expired** or **Failed**, synchronize configurations for the cluster or service to restore its configuration status. If all services in the cluster are in the **Failed** state, synchronize the cluster configuration with the background configuration.

Impact on the System

After synchronizing service configurations, you need to restart the services whose configurations have expired. These services are unavailable during restart.

Procedure

Step 1 On MRS Manager page, click **Services**.

Step 2 Select the target service from the service list.

Step 3 In the upper part of the service status and metric information, choose **More** > **Synchronize Configuration**.

Figure 9-7 Synchronization configurations

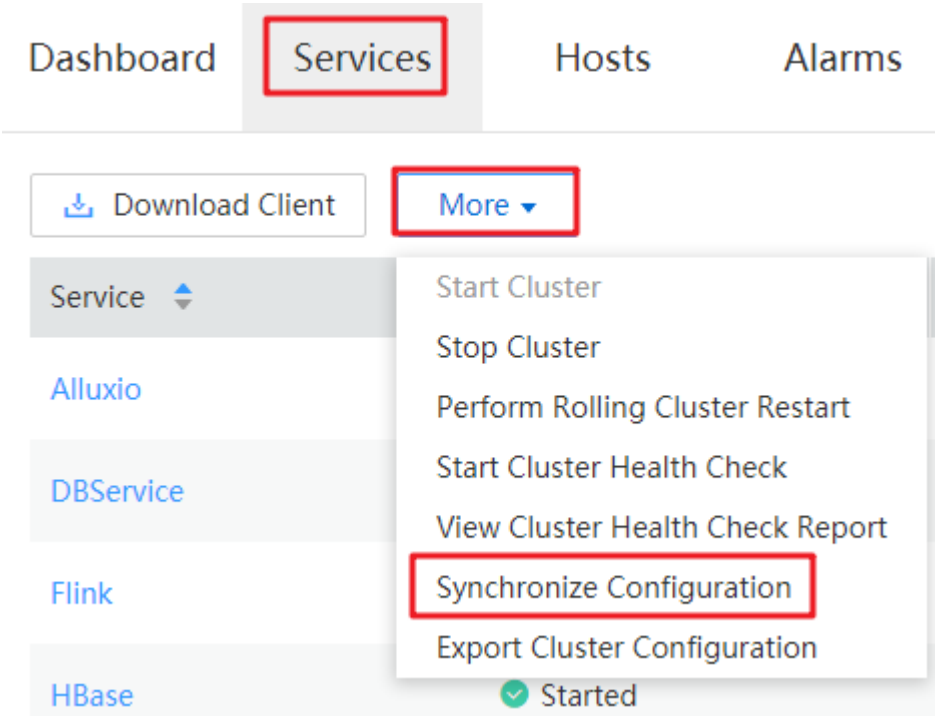

Step 4 In the dialog box that is displayed, enter the password as prompted and click **OK**. After your identity is verified, select **Restart the service or instance whose configuration has expired**. and click **OK** to restart the service whose configuration has expired.

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

----End

9.6.7 Managing Role Instances

Scenario

You can start a role instance that is in the **Stopped**, **Failed to stop** or **Failed to start** status, stop an unused or abnormal role instance or restart an abnormal role instance to recover its functions.

Procedure

Step 1 On MRS Manager page, click **Services**.

Step 2 Select the target service from the service list.

Step 3 Click the **Instances** tab.

Step 4 Select the check box on the left of the target role instance.

Step 5 Choose **More** > **Start Instance**, **Stop Instance**, or **Restart Instance** accordingly.

Figure 9-8 Instance operations

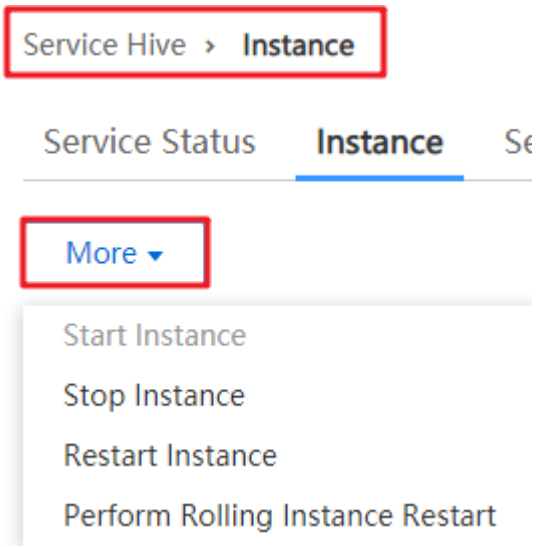

----End

9.6.8 Configuring Role Instance Parameters

Scenario

You can view and modify default role instance configurations on MRS Manager based on site requirements. The configurations can be imported and exported.

Impact on the System

You need to download and update the client configuration files after configuring HBase, HDFS, Hive, Spark, Yarn, and MapReduce service properties.

Procedure

- Modifying role instance configurations
  a. Click **Services**.
  b. Select the target service from the service list.
  c. Click the **Instances** tab.
  d. Click the target role instance from the role instance list.
  e. Click **Instance Configuration**.
  f. Set **Type** to **All**. The navigation tree of all configuration parameters of the role instance is displayed.
  g. In the navigation tree, select a specified parameter and change its value. You can also enter the parameter name in the **Search** box to search for the parameter and view the result.

     If you want to cancel the modification of a parameter value, click  to restore it.
  h. Click **Save Configuration**, select **Restart the role instance**, and click **OK** to restart the role instance.

     After **Operation successful.** is displayed, click **Finish**. The role instance is started successfully.
- Exporting Configuration Parameters of a Role Instance
  a. Click **Services**.
  b. Select a service.
  c. Select a role instance or click the **Instances** tab.
  d. Select a role instance on a specified host.
  e. Click **Instance Configuration**.
  f. Click **Export Instance Configuration** to export the configuration data of a specified role instance, and choose a path for saving the configuration file.
- Import configuration data of a role instance.
  a. Click **Services**.
  b. Select a service.
  c. Select a role instance or click the **Instances** tab.
  d. Select a role instance on a specified host.
  e. Click **Instance Configuration**.
  f. Click **Import Instance Configuration** to import the configuration data of the specified role instance.
  g. Click **Save Configuration** and select **Restart the role instance**. Click **OK**.

     After **Operation successful.** is displayed, click **Finish**. The role instance is started successfully.

9.6.9 Synchronizing Role Instance Configuration

Scenario

When **Configuration Status** of a role instance is **Expired** or **Failed**, you can synchronize the configuration data of the role instance with the background configuration.

Impact on the System

After synchronizing a role instance configuration, you need to restart the role instance whose configuration has expired. The role instance is unavailable during restart.

Procedure

Step 1 On MRS Manager, click **Services** and select a service name.

Step 2 Click the **Instances** tab.

Step 3 Click the target role instance from the role instance list.

Step 4 Choose **More** > **Synchronize Configuration** above the role instance status and indicator information.

Step 5 In the displayed dialog box, select **Restart services and instances whose configuration have expired**, and click **OK** to restart the role instance.

After **Operation successful** is displayed, click **Finish**. The role instance is started successfully.

Figure 9-9 Synchronizing role instance configurations

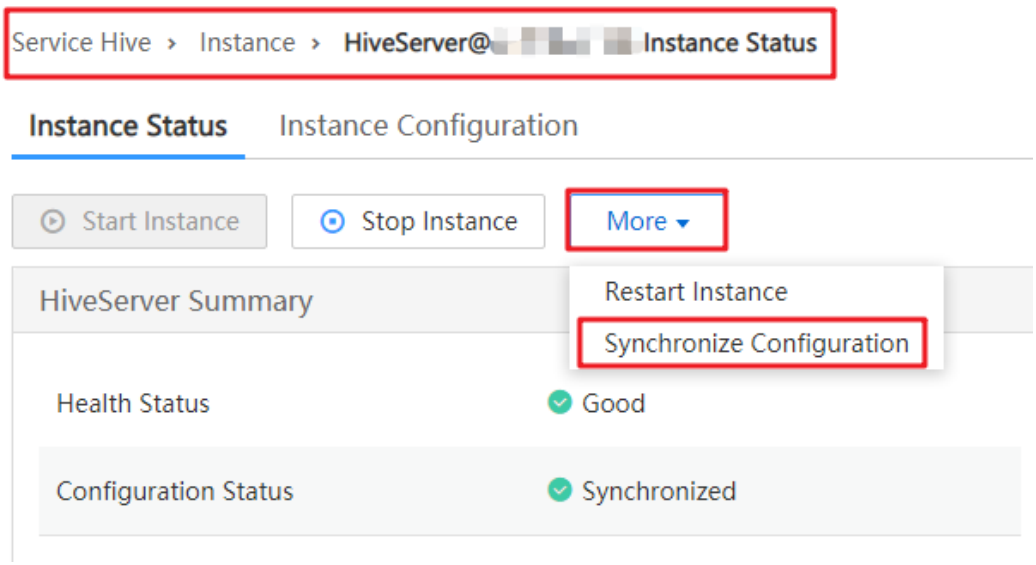

----End

9.6.10 Decommissioning and Recommissioning a Role Instance

Scenario

If a Core or Task node is faulty, the cluster status may be displayed as **Abnormal**. In an MRS cluster, data can be stored on different Core nodes. Users can decommission the specified role instance on MRS Manager to stop the role instance from providing services. After fault rectification, you can recommission the role instance.

The following role instances can be decommissioned and recommissioned.

- DataNode role instance on HDFS
- NodeManager role instance on Yarn
- RegionServer role instance on HBase
- Broker role instance on Kafka

Restrictions:

- If the number of the DataNodes is less than or equal to that of HDFS copies, decommissioning cannot be performed. If the number of HDFS copies is three and the number of DataNodes is less than four in the system, decommissioning cannot be performed. In this case, an error will be reported and the decommissioning will be stopped 30 minutes after the decommissioning attempt is performed on Manager.
- If the number of Kafka Broker instances is less than or equal to that of copies, decommissioning cannot be performed. For example, if the number of Kafka copies is two and the number of nodes is less than three in the system, decommissioning cannot be performed. Instance decommissioning will fail on Manager and exit.
- If a role instance is out of service, you must recommission the instance to start it before using it again.

Procedure

Step 1 On MRS Manager page, click **Services**.

Step 2 Click a service in the service list.

Step 3 Click the **Instances** tab.

Step 4 Select an instance.

Step 5 Choose **More** > **Decommission** or **Recommission** to perform the corresponding operation.

NOTE

During the instance decommissioning, if the service corresponding to the instance is restarted in the cluster using another browser, MRS Manager displays a message indicating that the instance decommissioning is stopped, but the **Operating Status** of the instance is displayed as **Started**. In this case, the instance has been decommissioned on the background. You need to decommission the instance again to synchronize the operating status.

----End

9.6.11 Managing a Host

Scenario

When a host is abnormal or faulty, you need to stop all roles of the host on MRS Manager to check the host. After the host fault is rectified, start all roles running on the host to recover host services.

Procedure

Step 1 Click **Hosts**.

Step 2 Select the check box of the target host.

Step 3 Choose **More** > **Start All Roles** or **Stop All Roles** accordingly.

----End

9.6.12 Isolating a Host

Scenario

If a host is found to be abnormal or faulty, affecting cluster performance or preventing services from being provided, you can temporarily exclude that host from the available nodes in the cluster. In this way, the client can access other available nodes. In scenarios where patches are to be installed in a cluster, you can also exclude a specified node from patch installation.

Users can isolate a host manually on MRS Manager based on the actual service requirements or O&M plan. Only non-management nodes can be isolated.

Impact on the System

- After a host is isolated, all role instances on the host will be stopped. You cannot start, stop, or configure the host and any instances on the host.
- After a host is isolated, statistics about the monitoring status and indicator data of the host hardware and instances on the host cannot be collected or displayed.

Procedure

Step 1 On MRS Manager, click **Hosts**.

Step 2 Select the check box of the host to be isolated.

Step 3 Choose **More** > **Isolate Host**,

Step 4 and click **OK** in the displayed dialog box.

After **Operation successful.** is displayed, click **Finish**. The host is isolated successfully, and the value of **Operating Status** becomes **Isolated**.

NOTE

For isolated hosts, you can cancel the isolation and add them to the cluster again. For details, see **Canceling Host Isolation**.

----End

9.6.13 Canceling Host Isolation

Scenario

After the exception or fault of a host is handled, you must cancel the isolation of the host for proper usage.

Users can cancel the isolation of a host on MRS Manager.

Prerequisites

- The host is in the **Isolated** state.
- The exception or fault of the host has been rectified.

Procedure

Step 1 On MRS Manager, click **Hosts**.

Step 2 Select the check box of the host to be de-isolated.

Step 3 Choose **More** > **Cancel Host Isolation**,

Step 4 and click **OK** in the displayed dialog box.

After **Operation successful.** is displayed, click **Finish**. The host is de-isolated successfully, and the value of **Operating Status** becomes **Normal**.

Step 5 Click the name of the de-isolated host to show its status, and click **Start All Roles**.

----End

9.6.14 Starting or Stopping a Cluster

Scenario

A cluster is a collection of service components. You can start or stop all services in a cluster.

Procedure

Step 1 On MRS Manager page, click **Services**.

Step 2 In the upper part of the service list, choose **More** > **Start Cluster** or **Stop Cluster** accordingly.

----End

9.6.15 Synchronizing Cluster Configurations

Scenario

If **Configuration Status** of all services or some services is **Expired** or **Failed**, synchronize configuration for the cluster or service to restore its configuration status.

- If all services in the cluster are in the **Failed** state, synchronize the cluster configuration with the background configuration.
- If all services in the cluster are in the **Failed** state, synchronize the service configuration with the background configuration.

Impact on the System

After synchronizing cluster configurations, you need to restart the services whose configurations have expired. These services are unavailable during restart.

Procedure

Step 1 On MRS Manager page, click **Services**.

Step 2 In the upper part of the service list, choose **More** > **Synchronize Configuration**.

Step 3 In the dialog box that is displayed, enter the password of the current login user for identity verification. After the verification is complete, select **Restart the service or instance whose configuration has expired**. and click **OK** to restart the service whose configuration has expired.

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

----End

9.6.16 Exporting Configuration Data of a Cluster

Scenario

You can export all configuration data of a cluster on MRS Manager to meet site requirements. The exported configuration data is used to rapidly update service configuration.

Procedure

Step 1 On MRS Manager page, click **Services**.

Step 2 Choose **More** > **Export Cluster Configuration**.

The exported file is used to update service configurations. For details, see **Import service configuration parameters** in **Configuring Service Parameters**.

----End

9.7 Log Management

9.7.1 About Logs

Log Description

MRS cluster logs are stored in the **/var/log/Bigdata** directory. The following table lists the log types.

Table 9-17 Log types

Type	Description
Installation log	Installation logs record information about FusionInsight Manager, cluster, and service installation to help users locate installation errors.
Run logs	Run logs record the running track information, debugging information, status changes, potential problems, and error information generated during the running of services.
Audit logs	Audit logs record information about users' activities and operation instructions, which can be used to locate fault causes in security events and determine who are responsible for these faults.

The following table lists the MRS log directories.

Table 9-18 Log directories

File Directory	Log Content
/var/log/Bigdata/audit	Component audit log.
/var/log/Bigdata/controller	Log collecting script log. Controller process log. Controller monitoring log.
/var/log/Bigdata/dbservice	DBService log.
/var/log/Bigdata/flume	Flume log.
/var/log/Bigdata/hbase	HBase log.
/var/log/Bigdata/hdfs	HDFS log.
/var/log/Bigdata/hive	Hive log.
/var/log/Bigdata/httpd	HTTPD log.
/var/log/Bigdata/hue	Hue log.

File Directory	Log Content
/var/log/Bigdata/kerberos	Kerberos log.
/var/log/Bigdata/ldapclient	LDAP client log.
/var/log/Bigdata/ldapserver	LDAP server log.
/var/log/Bigdata/loader	Loader log.
/var/log/Bigdata/logman	logman script log management log.
/var/log/Bigdata/mapreduce	MapReduce log.
/var/log/Bigdata/nodeagent	NodeAgent log.
/var/log/Bigdata/okerberos	OMS Kerberos log.
/var/log/Bigdata/oldapserver	OMS LDAP log.
/var/log/Bigdata/omm	**oms**: complex event processing log, alarm service log, HA log, authentication and authorization management log, and monitoring service run log of the omm server. **oma**: installation log and run log of the omm agent. **core**: dump log generated when the omm agent and the HA process are suspended.
/var/log/Bigdata/spark	Spark log.
/var/log/Bigdata/sudo	Log generated when the **sudo** command is executed by user **omm**.
/var/log/Bigdata/timestamp	Time synchronization management log.
/var/log/Bigdata/tomcat	Tomcat log.
/var/log/Bigdata/yarn	Yarn log.
/var/log/Bigdata/zookeeper	ZooKeeper log.
/var/log/Bigdata/kafka	Kafka log.
/var/log/Bigdata/storm	Storm log.
/var/log/Bigdata/patch	Patch log.

Run logs

Table 9-19 describes the running information recorded in run logs.

Table 9-19 Running information

Run Log	Description
Installation preparation log	Records information about preparations for the installation, such as the detection, configuration, and feedback operation information.
Process startup log	Records information about the commands executed during the process startup.
Process startup exception log	Records information about exceptions during process startup, such as dependent service errors and insufficient resources.
Process run log	Records information about the process running track information and debugging information, such as function entries and exits as well as cross-module interface messages.
Process running exception log	Records errors that cause process running errors, for example, the empty input objects or encoding or decoding failure.
Process running environment log	Records information about the process running environment, such as resource status and environment variables.
Script logs	Records information about the script execution process.
Resource reclamation log	Records information about the resource reclaiming process.
Uninstallation clearing logs	Records information about operations performed during service uninstallation, such as directory deletion and execution time

Audit logs

Audit information recorded in audit logs includes FusionInsight Manager audit information and component audit information.

Table 9-20 Audit information of FusionInsight Manager

Audit Log	Operation Type	Operation
Manager audit log	User management	Creating a user Modifying a user Deleting a user Creating a user group Modifying a user group Deleting a user group Adding a role Modifying a role Deleting a role Changing a password policy Changing a password Resetting a password User login User logout Unlocking the screen Downloading the authentication credential Unauthorized operation Unlocking a user account Locking a user account Locking the screen Exporting user information Exporting a user group Exporting a role

Audit Log	Operation Type	Operation
	Tenant management	Saving the static configuration Adding a tenant Deleting a tenant Associating a service with a tenant Deleting a service from a tenant Configuring resources Creating resources Deleting resources Adding a resource pool Modifying a resource pool Deleting a resource pool Restoring tenant data

Audit Log	Operation Type	Operation
	Cluster management	Starting a cluster Stopping a cluster Saving configurations Synchronizing cluster configurations Customizing cluster monitoring indicators Saving monitoring thresholds Downloading a client configuration file Configuring the northbound API Configuring the northbound SNMP API Creating a threshold template Deleting a threshold template Applying a threshold template Saving cluster monitoring configuration data Exporting configuration data Importing cluster configuration data Exporting an installation template Modifying a threshold template Canceling the application of a threshold template Masking alarms Sending an alarm Changing the OMS database password Changing the component database password Starting the health check of a cluster

Audit Log	Operation Type	Operation
		Updating the health check configuration Exporting cluster health check results Importing a certificate file Deleting historical health check reports Exporting historical health check reports Customizing report monitoring indicators Exporting report monitoring data Customizing monitoring indicators for static resource pools Exporting monitoring data of a static resource pool
	Service management	Starting a service Stopping a service Synchronizing service configurations Refreshing a service queue Customizing service monitoring indicators Restarting a service Exporting service monitoring data Importing service configuration data Starting the health check of a service Exporting service health check results Configuring the service Uploading a configuration file Downloading a configuration file

Audit Log	Operation Type	Operation
	Instance management	Synchronizing instance configurations Commissioning an instance Decommissioning an instance Starting an instance Stopping an instance Customizing instance monitoring indicators Restarting an instance Exporting instance monitoring data Importing instance configuration data
	Host management	Setting a node rack Starting all roles Stopping all roles Isolating a host Canceling host isolation Customizing host monitoring indicators Exporting host monitoring data Starting the health check of a host Exporting the health check result of a host

Audit Log	Operation Type	Operation
	Maintenance management	Exporting alarms Clearing alarms Exporting events Clearing alarms in batches Clearing alarm through SNMP Adding a trap target through SNMP Deleting a trap target through SNMP Checking alarms through SNMP Synchronizing alarms through SNMP Modifying audit dump configurations Exporting audit logs Collecting log files Downloading log files Uploading a file Deleting an uploaded file Creating a backup task Executing a backup task Stopping a backup task Deleting a backup task Modifying a backup task Locking a backup task Unlocking a backup task Creating a restoration task Executing a backup restoration task Stopping a restoration task Retrying a restoration task Deleting a restoration task

Table 9-21 Component audit information

Audit Log	Operation Type	Operation
DBService audit log	Maintenance management	Performing backup restoration operations
HBase audit log	Data definition language (DDL) statement	Creating a table Deleting a table Modifying a table Adding a column family Modifying a column family Deleting a column family Enabling a table Disabling a table Modify the user information Changing a password User login
	Data manipulation language (DML) statement	Putting data (to the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables) Deleting data (from the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables) Checking and putting data (to the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables) Checking and deleting data (from the **hbase:meta**, **_ctmeta_**, and **hbase:acl** tables)
	Permission control	Assigning permissions to a user Canceling permission assigning

Audit Log	Operation Type	Operation
Hive audit logs	Metadata operation	Defining metadata, such as creating databases and tables Deleting metadata, such as deleting databases and tables Modifying metadata, such as adding columns and renaming tables Importing and exporting metadata
	Data maintenance	Loading data to a table Inserting data into a table
	Permissions management	Creating or deleting roles Granting/Reclaiming roles Granting/Reclaiming permissions
HDFS audit log	Permissions management	Managing permissions on files or folders Managing permissions on owner information files or folders
	File operation	Creating a folder Creating a file Opening a file Appending file content Changing a file name Deleting a file or folder Setting time property of a file Setting the number of file copies Merging files Checking the file system File links

Audit Log	Operation Type	Operation
MapReduce audit log	Application running	Starting a Container request Stopping a Container request After Container request is completed, the status of the request is displayed as succeeded. After Container request is completed, the status of the request is displayed as failed. After Container request is completed, the status of the request is displayed as suspended. Submitting a task Ending a task
LdapServer audit log	Maintenance management	Adding an operating system user Adding a user group Adding a user to user group Deleting a user Deleting a group
KrbServer audit log	Maintenance management	Changing the password of a Kerberos account Adding a Kerberos account Deleting a Kerberos account Authenticating a user
Loader audit log	Security management	User login
	Metadata management	Querying connector information Querying a framework Querying step information

Audit Log	Operation Type	Operation
	Managing data source connections	Querying a data source connection Adding a data source connection Updating a data source connection Deleting a data source connection Activating a data source connection Disabling a data source connection
	Job management	Querying a job Creating a Job Updating a Job Deleting a job Activating a job Disabling a job Querying all execution records of a job Querying the latest execution record of a job Submitting a job Stopping a job
Hue audit log	Service startup	Starting Hue
	User operation	User login User logout
	Task operation	Creating a job Modifying a job Deleting a job Submitting a task Saving a task Updating the status of a task
ZooKeeper audit log	Permissions management	Setting the access permission to Znode
	Znode operation	Creating a Znode Deleting a Znode Configuring Znode data

Audit Log	Operation Type	Operation
Storm audit log	Nimbus	Submitting a topology Stopping a topology Reallocating a topology Deactivating a topology Activating a topology
	UI	Stopping a topology Reallocating a topology Deactivating a topology Activating a topology

MRS audit logs are stored in the database. You can view and export audit logs on the **Audit** page.

The following table lists the directories to store component audit logs. Audit log files of some components are stored in **/var/log/Bigdata/audit**, such as HDFS, HBase, MapReduce, Hive, Hue, Yarn, Storm, and ZooKeeper. The component audit logs are automatically compressed and backed up to **/var/log/Bigdata/audit/bk** at 03: 00 every day. A maximum of latest 90 compressed backup files are retained, and the backup time cannot be changed.

Audit log files of other components are stored in the component log directory.

Table 9-22 Directory for storing component audit logs

Component	Audit Log Directory
DBService	/var/log/Bigdata/audit/dbservice/dbservice_audit.log
HDFS	/var/log/Bigdata/audit/hdfs/nn/hdfs-audit-namenode.log /var/log/Bigdata/audit/hdfs/dn/hdfs-audit-datanode.log /var/log/Bigdata/audit/hdfs/jn/hdfs-audit-journalnode.log /var/log/Bigdata/audit/hdfs/zkfc/hdfs-audit-zkfc.log /var/log/Bigdata/audit/hdfs/httpfs/hdfs-audit-httpfs.log /var/log/Bigdata/audit/hdfs/router/hdfs-audit-router.log
MapReduce	/var/log/Bigdata/audit/mapreduce/jobhistory/mapred-audit-jobhistory.log

Component	Audit Log Directory
Hive	/var/log/Bigdata/audit/hive/hiveserver/hive-audit.log /var/log/Bigdata/audit/hive/metastore/metastore-audit.log /var/log/Bigdata/audit/hive/webhcat/webhcat-audit.log
Loader	/var/log/Bigdata/loader/audit/default.audit
Hue	/var/log/Bigdata/audit/hue/hue-audits.log
ZooKeeper	/var/log/Bigdata/audit/zookeeper/quorumpeer/zk-audit-quorumpeer.log
Spark	/var/log/Bigdata/audit/spark/jdbcserver/jdbcserver-audit.log /var/log/Bigdata/audit/spark/jobhistory/jobhistory-audit.log
Yarn	/var/log/Bigdata/audit/yarn/rm/yarn-audit-resourcemanager.log /var/log/Bigdata/audit/yarn/nm/yarn-audit-nodemanager.log
Storm	/var/log/Bigdata/audit/storm/nimbus/audit.log /var/log/Bigdata/audit/storm/ui/audit.log

9.7.2 Manager Log List

Log Description

Log path: The default storage path of Manager log files is **/var/log/Bigdata/***Manager component*.

- ControllerService: **/var/log/Bigdata/controller/** (operation & maintenance system (OMS) installation and run logs)
- Httpd: **/var/log/Bigdata/httpd** (httpd installation and run logs)
- logman: **/var/log/Bigdata/logman** (log packaging tool logs)
- NodeAgent: **/var/log/Bigdata/nodeagent** (NodeAgent installation and run logs)
- okerberos: **/var/log/Bigdata/okerberos** (okerberos installation and run logs)
- oldapserver: **/var/log/Bigdata/oldapserver** (oldapserver installation and run logs)
- MetricAgent: **/var/log/Bigdata/metric_agent** (MetricAgent run log)
- omm: **/var/log/Bigdata/omm** (omm installation and run logs)
- timestamp: **/var/log/Bigdata/timestamp** (NodeAgent startup time logs)
- tomcat: **/var/log/Bigdata/tomcat** (Web process logs)

- Patch: **/var/log/Bigdata/patch** (patch installation log)
- Sudo: **/var/log/Bigdata/sudo** (sudo script execution log)
- OS: **/var/log/***message file* (OS system log)
- OS Performance: **/var/log/osperf** (OS performance statistics log)
- OS Statistics: **/var/log/osinfo/statistics** (OS parameter configuration log)

Log archiving rule:

The automatic compression and archiving function is enabled for Manager logs. By default, when the size of a log file exceeds 10 MB, the log file is automatically compressed. The naming rule of a compressed log file is as follows: <*Original log name*>-<*yyyy-mm-dd_hh-mm-ss*>.[*ID*].**log.zip** A maximum of 20 latest compressed files are reserved.

Table 9-23 Manager logs

Type	Log File Name	Description
Controller run log	controller.log	Log that records component installation, upgrade, patch installation, configuration, monitoring, alarms, and routine O&M operations
	controller_client.log	Run log of the Representational State Transfer (REST) API
	acs.log	ACS run log file
	acs_spnego.log	spnego user log in ACS
	aos.log	AOS run log
	plugin.log	AOS plug-in log
	backupplugin.log	Log that records the backup and restoration operations
	controller_config.log	Configuration run log
	controller_nodesetup.log	Controller loading task log
	controller_root.log	System log of the Controller process
	controller_trace.log	Log that records the remote procedure call (RPC) communication between Controller and NodeAgent

Type	Log File Name	Description
	controller_monitor.log	Monitoring log
	controller_fsm.log	State machine log
	controller_alarm.log	Controller alarm log
	controller_backup.log	Controller backup and recovery log
	install.log, distributeAdapterFiles.log, install_os_optimization.log	OMS installation log
	oms_ctl.log	OMS startup and stop log
	installntp.log	NTP installation log
	modify_manager_param.log	Manager parameter modification log
	backup.log	OMS backup script run log
	supressionAlarm.log	Alarm script run log
	om.log	OM certificate generation log
	backupplugin_ctl.log	Startup log of the backup and restoration plug-in process
	getLogs.log	Run log of the collection log script
	backupAuditLogs.log	Run log of the audit log backup script
	certStatus.log	Log that records regular certificate checks
	distribute.log	Certificate distribution log
	ficertgenetrate.log	Certificate replacement logs, including logs of level-2 certificates, CAS certificates, and httpd certificates
	genPwFile.log	Log that records the generation of certificate password files

Type	Log File Name	Description
	modifyproxyconf.log	Log that records the modification of the HTTPD proxy configuration
	importTar.log	Log that records the process of importing certificates into the trust library
Httpd	install.log	Httpd installation log
	access_log, error_log	Httpd run log
logman	logman.log	Log packaging tool log
NodeAgent	install.log, install_os_optimization.log	NodeAgent installation log
	installntp.log	NTP installation log
	start_ntp.log	NTP startup log
	ntpChecker.log	NTP check log
	ntpMonitor.log	NTP monitoring log
	heartbeat_trace.log	Log that records heartbeats between NodeAgent and Controller
	alarm.log	Alarm log
	monitor.log	Monitoring log
	nodeagent_ctl.log, start-agent.log	NodeAgent startup log
	agent.log	NodeAgent run log
	cert.log	Certificate log
	agentplugin.log	Agent plug-in running status monitoring log
	omaplugin.log	OMA plug-in run log
	diskhealth.log	Disk health check log
	supressionAlarm.log	Alarm script run log
	updateHostFile.log	Host list update log
	collectLog.log	Run log of the node log collection script

Type	Log File Name	Description
	host_metric_collect.log	Host index collection run log
	checkfileconfig.log	Run log file of file permission check
	entropycheck.log	Entropy check run log
	timer.log	Log of periodic node scheduling
	pluginmonitor.log	Component monitoring plug-in log
	agent_alarm_py.log	Log that records alarms upon insufficient NodeAgent file permission
okerberos	addRealm.log, modifyKerberosRealm.log	Domain handover log
	checkservice_detail.log	Okerberos health check log
	genKeytab.log	keytab generation log
	KerberosAdmin_genConfigDetail.log	Run log that records the generation of kadmin.conf when starting the kadmin process
	KerberosServer_genConfigDetail.log	Run log that records the generation of krb5kdc.conf when starting the krb5kdc process
	oms-kadmind.log	Run log of the kadmin process
	oms_kerberos_install.log, postinstall_detail.log	Okerberos installation log
	oms-krb5kdc.log	Run log of the krbkdc process
	start_detail.log	Okerberos startup log
	realmDataConfigPro-cess.log	Log rollback for domain handover failure
	stop_detail.log	Okerberos stop log
oldapserver	ldapserver_backup.log	Oldapserver backup log

Type	Log File Name	Description
	ldapserver_chk_service.log	Oldapserver health check log
	ldapserver_install.log	Oldapserver installation log
	ldapserver_start.log	Oldapserver startup log
	ldapserver_status.log	Log that records the status of the Oldapserver process
	ldapserver_stop.log	Oldapserver stop log
	ldapserver_wrap.log	Oldapserver service management log
	ldapserver_uninstall.log	Oldapserver uninstallation log
	restart_service.log	Oldapserver restart log
	ldapserver_unlockUser.log	Log that records information about unlocking LDAP users and managing accounts
omm	omsconfig.log	OMS configuration log
	check_oms_heartbeat.log	OMS heartbeat log
	monitor.log	OMS monitoring log
	ha_monitor.log	HA_Monitor operation log
	ha.log	HA operation log
	fms.log	Alarm log
	fms_ha.log	HA alarm monitoring log
	fms_script.log	Alarm control log
	config.log	Alarm configuration log
	iam.log	IAM log
	iam_script.log	IAM control log
	iam_ha.log	IAM HA monitoring log
	config.log	IAM configuration log
	operatelog.log	IAM operation log

Type	Log File Name	Description
	heartbeatcheck_ha.log	OMS heartbeat HA monitoring log
	install_oms.log	OMS installation log
	pms_ha.log	HA monitoring log
	pms_script.log	Monitoring control log
	config.log	Monitoring configuration log
	plugin.log	Monitoring plug-in run log
	pms.log	Monitoring log
	ha.log	HA run log
	cep_ha.log	CEP HA monitoring log
	cep_script.log	CEP control log
	cep.log	CEP log
	config.log	CEP configuration log
	omm_gaussdba.log	GaussDB HA monitoring log
	gaussdb-<SERIAL>.log	GaussDB run log
	gs_ctl-<DATE>.log	GaussDB control log archive log
	gs_ctl-current.log	GaussDB control log
	gs_guc-current.log	GaussDB operation log
	encrypt.log	Omm encryption log
	omm_agent_ctl.log	OMA control log
	oma_monitor.log	OMA monitoring log
	install_oma.log	OMA installation log
	config_oma.log	OMA configuration log
	omm_agent.log	OMA run log
	acs.log	ACS resource log
	aos.log	AOS resource log
	controller.log	Controller resource log

Type	Log File Name	Description
	feed_watchdog.log	feed_watchdog resource log
	floatip.log	Floating IP address resource log
	ha_ntp.log	NTP resource log
	httpd.log	Httpd resource log
	okerberos.log	Okerberos resource log
	oldap.log	OLdap resource log
	tomcat.log	Tomcat resource log
	send_alarm.log	Run log of the HA alarm sending script of the management node
timestamp	restart_stamp	NodeAgent start time log
tomcat	cas.log, localhost_access_cas_log.log	CAS run log
	catalina.log, catalina.out, host-manager.log, localhost.log, manager.log	Tomcat run log
	localhost_access_web_log.log	Log that records the access to REST APIs of FusionInsight Manager
	web.log	Run log of the web process
	northbound_ftp_sftp.log, snmp.log	Northbound log
watchdog	watchdog.log, feed_watchdog.log	watchdog run log
patch	oms_installPatch.log	OMS patch installation log
	agent_installPatch.log	Agent patch installation log
	agent_uninstallPatch.log	Agent patch uninstallation log
	NODE_AGENT_restoreFile.log	Agent patch restoration log

Type	Log File Name	Description
	NODE_AGENT_updateFile.log	Agent patch update log
	OMA_restoreFile.log	OMA patch restoration file log
	OMA_updateFile.log	OMA patch update file log
	CONTROLLER_restoreFile.log	CONTROLLER patch restoration file log
	CONTROLLER_updateFile.log	CONTROLLER patch update file log
	OMS_restoreFile.log	OMS patch restoration file log
	oms_uninstallPatch.log	OMS patch uninstallation log
	OMS_updateFile.log	OMS patch update file log
	createStackConf.log, decompress.log, decompress_OMS.log, distrExtractPatchO-nOMS.log, slimReduction.log, switch_adapter.log	Patch installation log
sudo	sudo.log	Sudo script execution log

Log Levels

Table 9-24 describes the log levels provided by Manager. The priorities of log levels are FATAL, ERROR, WARN, INFO, and DEBUG in descending order. Logs whose levels are higher than or equal to the specified level are printed. The number of printed logs decreases as the specified log level increases.

Table 9-24 Log levels

Level	Description
FATAL	Logs of this level record fatal error information about the current event processing that may result in a system crash.
ERROR	Logs of this level record error information about the current event processing, which indicates that system running is abnormal.

Level	Description
WARN	Abnormal information about the current event processing. These abnormalities will not result in system faults.
INFO	Normal running status information about the system and events.
DEBUG	Logs of this level record the system information and system debugging information.

Log Formats

The following table lists the Manager log formats.

Table 9-25 Log formats

Type	Component	Format	Example
Controller, Httpd, logman, NodeAgent, okerberos, oldapserver, omm, tomcat, upgrade	Controller, Httpd, logman, NodeAgent, okerberos, oldapserver, omm, tomcat, upgrade	<*yyyy-MM-dd HH:mm:ss,SSS*>|<*Log level*>|<*Name of the thread that generates the log*>|<*Message in the log*>|<*Location where the log event occurs*>	2015-06-30 00:37:09,067 INFO [pool-1-thread-1] Completed Discovering Node. com.huawei.hadoop.om.controller.tasks.nodesetup.DiscoverNodeTask.execute(DiscoverNodeTask.java:299)

9.7.3 Viewing and Exporting Audit Logs

Scenario

This section describes how to view and export audit logs on MRS Manager. The audit logs can be used to trace security events, locate fault causes, and determine responsibilities.

The system record the following log information:

- User activity information, such as user login and logout, system user information modification, and system user group information modification
- User operation instruction information, such as cluster startup, stop, and software upgrade.

Procedure

- Viewing audit logs

a. On MRS Manager, click **Audit** to view the default audit logs.

   If an audit log contains more than 256 characters, click the expand button to view audit log details.

   - By default, records are sorted in descending order by the **Occurred** column. You can click **Operation Type**, **Severity**, **Occurred**, **User**, **Host**, **Service**, **Instance**, or **Operation Result** to change the sorting mode.
   - All alarms of the same severity can be filtered by **Severity**. The results include cleared and uncleared alarms.

   Exported audit logs contain the following information:

   - **Sno**: indicates the number of audit logs generated by MRS Manager. The number is incremented by 1 when a new audit log is generated.
   - **Operation Type**: indicates the operation type of a user operation. There are nine scenarios: **Alarm**, **Auditlog**, **Backup And Restoration**, **Cluster**, **Collect Log**, **Host**, **Service**, **Tenant** and **User_Manager**. **User_Manager** is supported only in clusters with Kerberos authentication enabled. Each scenario contains different operation types. For example, **Alarm** includes **Export alarms**; **Cluster** includes **Start cluster**, and **Tenant** include **Add tenant**.
   - **Severity**: indicates the security level of each audit log, including **Critical**, **Major**, **Minor** and **Informational**.
   - **Start Time**: indicates the time when the operation starts. The time is CET or CEST.
   - **End Time**: indicates the time when the operation ends. The time is CET or CEST.
   - **User IP Address**: indicates the IP address used by a user to perform operations.
   - **User**: indicates the name of the user who performs the operation.
   - **Host**: indicates the node where the user operation is performed. The information is not saved if the operation does not involve a node.
   - **Service**: indicates the service in the cluster where the user operation is performed. The information is not saved if the operation does not involve a service.
   - **Instance**: indicates the role instance in the cluster where the user operation is performed. The information is not saved if the operation does not involve a role instance.
   - **Operation Result**: indicates the operation result, including **Successful**, **Failed** and **Unknown**.
   - **Content**: indicates execution information of the user operation.

b. Click **Advanced Search**. In the search area, set search criteria and click **Search** to view audit logs of the specified type. Click **Reset** to clear the search criteria.

NOTE

Start Time and **End Time** specify the start time and end time of the time range. You can search for alarms generated within the time range.

- Exporting audit logs
  a. In the audit log list, click **Export All** to export all logs.
  b. In the audit log list, select the check box of a log and click **Export** to export the log.

9.7.4 Exporting Service Logs

Scenario

This section describes how to export logs generated by each service role from MRS Manager.

Prerequisites

- You have obtained the access key ID (AK) and secret access key (SK) of the account.
- A parallel file system has been created in OBS.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 Click **Export Log** under **Maintenance**.

Step 3 Set a service for **Service**. Set **Host** to the IP address of the host where the service is deployed. Select the corresponding time for **Start Time** and **End Time**.

Step 4 In **Export To**, select a path for saving logs. This parameter is available only for clusters with Kerberos authentication enabled.

- **Local PC**: indicates that logs are saved to the local environment. Then go to **Step 8**.
- **OBS**: indicates that logs are saved to OBS. This is the default option. Then go to **Step 5**.

Step 5 Set **OBS Path** to the path for storing service logs on OBS.

The value must be a complete path and cannot start with a slash (/). The path can be nonexistent and will be automatically created by the system. The full path of OBS can contain a maximum of 900 bytes.

Step 6 In **Bucket**, enter the name of the created OBS file system.

Step 7 Set **AK** and **SK** to the access key ID and secret access key of the user.

Step 8 Click **OK**.

----End

9.7.5 Configuring Audit Log Exporting Parameters

Scenario

If MRS audit logs are stored in the system for a long time, the disk space of the data directory may be insufficient. Therefore, you can set export parameters to automatically export audit logs to a specified directory on the OBS server timely, facilitating audit log management.

NOTE

Audit logs exported to the OBS server include service audit logs and management audit logs.

- Service audit logs are automatically compressed and stored in the **/var/log/Bigdata/audit/bk/** directory on the active management node at 03:00 every day. The file name format is *<yyyy-MM-dd_HH-mm-ss>***.tar.gz**. By default, a maximum of seven log files can be stored. If more than seven log files are stored, the system automatically deletes the log files generated seven days ago.
- The data range of management audit logs exported to OBS each time is from the last date when the logs are successfully exported to OBS to the date when the task is executed. When the number of management audit logs reaches 100,000, the system automatically dumps the first 90,000 audit logs to a local file and retains 10,000 audit logs in the database. The dumped log files are saved in the **${BIGDATA_DATA_HOME}/dbdata_om/dumpData/iam/operatelog** directory on the active management node. The file name format is **OperateLog_store**_*YY_MM_DD_HH_MM_SS*.**csv**. A maximum of 50 historical audit log files can be saved.

Prerequisites

- You have obtained the access key ID (AK) and secret access key (SK) of the account.
- A parallel file system has been created in OBS.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 Choose **Export Audit Log** under **Maintenance**.

Table 9-26 Parameters for exporting audit logs

Parameter	Value	Description
Export Audit Log	• On • Off	(Mandatory) Specifies whether to enable the audit log export function. • On: enables audit log exporting. • Off: disables audit log exporting.
Start Time	07/24/2017 09:00:00 (example)	(Mandatory) Specifies the start time for exporting audit logs.

Parameter	Value	Description
Period (days)	1 day (example value)	(Mandatory) Specifies the interval for exporting audit logs. The interval ranges from 1 to 5 days.
Bucket	mrs-bucket (example value)	(Mandatory) Specifies the name of the OBS file system to which audit logs are exported.
OBS path	**/opt/omm/oms/auditLog** (example value)	(Mandatory) Specifies the OBS path to which audit logs are exported.
AK	*XXX* (example value)	(Mandatory) Specifies the user's access key ID.
SK	*XXX* (example value)	(Mandatory) Specifies the user's secret access key.

NOTE

Audit logs are stored in **service_auditlog** and **manager_auditlog** on OBS. They are used to store service audit logs and management audit logs, respectively.

----End

9.8 Health Check Management

9.8.1 Performing a Health Check

Scenario

To ensure that cluster parameters, configurations, and monitoring are correct and that the cluster can run stably for a long time, you can perform a health check during routine maintenance.

NOTE

A system health check includes MRS Manager, service-level, and host-level health checks:

- MRS Manager health checks focus on whether the unified management platform can provide management functions.
- Service-level health checks focus on whether components can provide services properly.
- Host-level health checks focus on whether host indicators are normal.

The system health check includes three types of check items: health status, related alarms, and customized monitoring indicators for each check object. The health check results are not always the same as the **Health Status** on the portal.

Procedure

- Manually perform the health check for all services.
  a. Click **Services**.
  b. Choose **More** > **Start Service Health Check** to start the health check for the service.

  Figure 9-10 Starting the service health check

  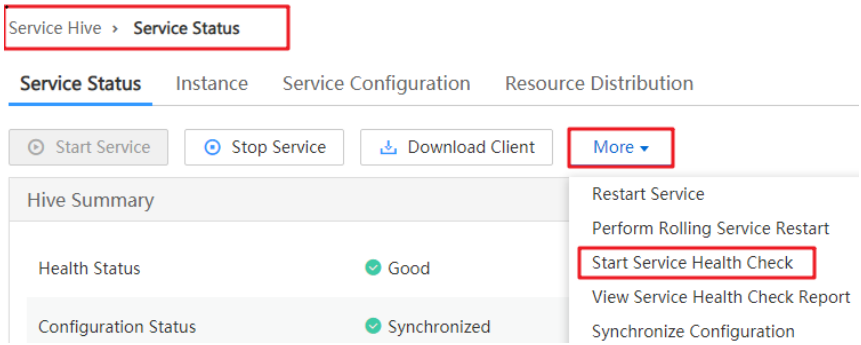

  NOTE

  - The cluster health check includes Manager, service, and host status checks.
  - To perform cluster health checks, you can also choose **System** > **Maintenance** > **Check Health Check** > **Start Cluster Health Check** on MRS Manager.
  - To export the health check result, click **Export Report** in the upper left corner.
- Manually perform the health check for a service.
  a. Click **Services**. In the services list, click the desired service name.
  b. Choose **More** > **Start Service Health Check** to start the health check for the service.
- Manually perform the health check for a host.
  a. Click **Hosts**.
  b. Select the check box of the host for which you want to check the health status.
  c. Choose **More** > **Start Host Health Check** to start the health check for the host.
- Automatically performing a health check
  a. Click **System**.
  b. Click **Check Health Status** under **Maintenance**.
  c. Click **Configure Health Check** to configure automatic health check items.

  Max. Number of Health Check Reports: maximum number of health check reports. The value must be an integer ranging from 1 to 100.

  Periodic Health Check: specifies whether to enable automatic health check. The **Periodic Health Check** function is disabled by default. You can click to enable the function and select **Daily**, **Weekly**, or **Monthly** based on management requirements.
  d. Click **OK** to save the settings. The **Health check configuration saved successfully** is displayed in the upper right corner.

9.8.2 Viewing and Exporting a Health Check Report

Scenario

You can view the health check result in MRS Manager and export the health check results for further analysis.

NOTE

A system health check includes MRS Manager, service-level, and host-level health checks:

- MRS Manager health checks focus on whether the unified management platform can provide management functions.
- Service-level health checks focus on whether components can provide services properly.
- Host-level health checks focus on whether host indicators are normal.

The system health check includes three types of check items: health status, related alarms, and customized monitoring indicators for each check object. The health check results are not always the same as the **Health Status** on the portal.

Prerequisites

You have performed a health check.

Procedure

Step 1 Click **Services**.

Step 2 Choose **More > View Cluster Health Check Report** to view the health check report of a cluster.

Step 3 Click **Export Report** on the health check report pane to export the report and view detailed information about check items.

NOTE

For details about how to rectify the faults of the check items, see **DBService Health Check Indicators** to **ZooKeeper Health Check Indicators**.

----End

9.8.3 Configuring the Number of Health Check Reports to Be Reserved

Scenario

Health check reports of MRS clusters, services, and hosts may vary with the time and scenario. You can modify the number of health check reports to be reserved on MRS Manager for later comparison.

This setting is valid for health check reports of clusters, services, and hosts. Report files are saved in **$BIGDATA_DATA_HOME/Manager/healthcheck** on the active management node by default and are automatically synchronized to the standby management node.

Prerequisites

Users have specified service requirements and planned the save time and health check frequency, and the disk space of the active and standby management nodes is sufficient.

Procedure

Step 1 Choose **System** > **Maintenance** > **Check Health Status** > **Configure Health Check**.

Step 2 Set **Max. Number of Health Check Reports** to the number of health check reports to be reserved. The value ranges from 1 to 100. The default value is 50.

Step 3 Click **OK** to save the settings. Message "Health check configuration saved successfully" is displayed in the upper right corner.

----End

9.8.4 Managing Health Check Reports

Scenario

On MRS Manager, users can manage historical health check reports, for example, viewing, downloading, and deleting historical health check reports.

Procedure

- Download a specified health check report.
  a. Choose **System** > **Maintenance** > **Check Health Status**.
  b. Locate the row that contains the target health check report and click **Download** to download the report file.
- Download specified health check reports in batches.
  a. Choose **System** > **Maintenance** > **Check Health Status**.
  b. Select multiple health check reports and click **Download File** to download them.
- Delete a specified health check report.
  a. Choose **System** > **Maintenance** > **Check Health Status**.
  b. Locate the row that contains the target health check report and click **Delete** to delete the report file.
- Delete specified health check reports in batches.
  a. Choose **System** > **Maintenance** > **Check Health Status**.
  b. Select multiple health check reports and click **Delete File** to delete them.

9.8.5 DBService Health Check Indicators

Service Health Check

Indicator: Service Status

Description: This indicator is used to check whether the DBService service status is normal. If the status is abnormal, the service is unhealthy.

Handling method: If the indicator is abnormal, rectify the fault by referring to ALM-27001.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist on the host. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.6 Flume Health Check Indicators

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the Flume service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If the indicator is abnormal, rectify the fault by referring to ALM-24000.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist on the host. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.7 HBase Health Check Indicators

Normal RegionServer Count

Indicator: Normal RegionServer Count

Description: This indicator is used to check the number of RegionServers that are running properly in an HBase cluster.

Recovery Guide: If the indicator is abnormal, check whether the status of RegionServer is normal. If the status is abnormal, resolve the problem and check that the network is normal.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the HBase service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If the indicator is abnormal, check whether the status of HMaster and RegionServer is normal. If the status is abnormal, resolve the problem. Then, check whether the status of the ZooKeeper service is faulty. On the HBase client, check whether the data in the HBase table can be correctly read and locate the data reading failure cause. Handle the alarm following instructions in the alarm processing document.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.8 Host Health Check Indicators

Swap Usage

Indicator: Swap Usage

Description: Swap usage of the system. The value is calculated using the following formula: Swap usage = Used swap size/Total swap size. Assume that the current threshold is set to 75.0%. If the usage of the file handles in the system exceeds the threshold, the system is unhealthy.

Recovery Guide:

1. Check the swap usage of the node.

   Log in to the unhealthy node and run the **free -m** command to check the total swap space and used swap space. If the swap space usage exceeds the threshold, go to **2**.
2. If the swap usage exceeds the threshold, you are advised to expand the system capacity, for example, add nodes.

Host File Handle Usage

Indicator: Host File Handle Usage

Description: This indicator indicates the file handle usage in the system. Host file handle usage = Number of used handles/Total number of handles. If the usage exceeds the threshold, the system is unhealthy.

Recovery Guide:

1. Check the file handle usage of the host.

   Log in to the unhealthy node and run the **cat /proc/sys/fs/file-nr** command. In the command output, the first and third columns indicate the number of used handles and the total number of handles, respectively. If the usage exceeds the threshold, go to **2**.
2. If the file handle usage of the host exceeds the threshold, you are advised to check the system and analyze the file handle usage.

NTP Offset

Indicator: NTP Offset

Description: This indicator indicates the NTP time offset. If the time deviation exceeds the threshold, the system is unhealthy.

Recovery Guide:

1. Check the NTP time offset.

   Log in to the unhealthy node and run the **/usr/sbin/ntpq -np** command to view the information. In the command output, the **Offset** column indicates the time offset. If the time offset is greater than the threshold, go to **2**.

2. If the indicator is abnormal, check whether the clock source configuration is correct. Contact O&M personnel.

Average Load

Indicator: Average Load

Description: Average system load, indicating the average number of processes in the running queue in a specified period. The system average load is calculated using the load value obtained by the uptime command. Calculation method: (Load of 1 minute + Load of 5 minutes + Load of 15 minutes)/(3 x Number of CPUs). Assume that the current threshold is set to 2. If the average load exceeds 2, the system is unhealthy.

Recovery Guide:

1. Log in to the unhealthy node and run the **uptime** command. The last three columns in the command output indicate the load in 1 minute, 5 minutes, and 15 minutes, respectively. If the average system load exceeds the threshold, go to **2**.
2. If the system average load exceeds the threshold, you are advised to perform system capacity expansion, such as adding nodes.

D State Process

Indicator: D State Process

Description: This indicator indicates the unstoppable sleep process, that is, the process in the D state. A process that is in the D state is waiting for I/O, such as disk I/O and network I/O, and experiences an I/O exception. If any process in the D state exists in the system, the system is unhealthy.

Recovery Guide: If the indicator is abnormal, the system generates an alarm. You are advised to handle the alarm by referring to ALM-12028.

Hardware Status

Indicator: Hardware Status

Description: This indicator is used to check the system hardware status, including the CPU, memory, disk, power supply, and fan. This indicator obtains related hardware information using **ipmitool sdr elist**. If the hardware status is abnormal, the hardware is unhealthy.

Recovery Guide:

1. Log in to the node where the check result is unhealthy. Run the **ipmitool sdr elist** command to check system hardware status. The last column in the command output indicates the hardware status. If the status is included in the following fault description table, the check result is unhealthy.

Module	Symptom
Processor	IERR Thermal Trip FRB1/BIST failure FRB2/Hang in POST failure FRB3/Processor startup/init failure Configuration Error SM BIOS Uncorrectable CPU-complex Error Disabled Throttled Uncorrectable machine check exception
Power Supply	Failure detected Predictive failure Power Supply AC lost AC lost or out-of-range AC out-of-range, but present Config Error: Vendor Mismatch Config Error: Revision Mismatch Config Error: Processor Missing Config Error: Power Supply Rating Mismatch Config Error: Voltage Rating Mismatch Config Error
Power Unit	240VA power down Interlock power down AC lost Soft-power control failure Failure detected Predictive failure

Module	Symptom
Memory	Uncorrectable ECC Parity Memory Scrub Failed Memory Device Disabled Correctable ECC logging limit reached Configuration Error Throttled Critical Overtemperature
Drive Slot	Drive Fault Predictive Failure Parity Check In Progress In Critical Array In Failed Array Rebuild In Progress Rebuild Aborted
Battery	Low Failed

2. If the indicator is abnormal, contact O&M personnel.

Host Name

Indicator: Host Name

Description: This indicator is used to check whether the host name is set. If the host name is not set, the system is unhealthy. If the indicator is abnormal, you are advised to set the host name properly.

Recovery Guide:

1. Log in to the node where the check result is unhealthy.
2. Run the hostname host name command to change the host name to ensure that the host name is consistent with the planned host name.

   hostname*host name* For example, to change the host name to **Bigdata-OM-01**, run the **hostname Bigdata-OM-01** command.
3. Modify the host name configuration file.

   Run the **vi /etc/HOSTNAME** command to edit the file. Change the file content to **Bigdata-OM-01**. Save the file, and exit.

Umask

Indicator: Umask

Description: This indicator is used to check whether the umask setting of user **omm** is correct. If Umask is not 0077, the system is unhealthy.

Recovery Guide:

1. If the indicator is abnormal, you are advised to set umask of user **omm** to 0077. Log in to the unhealthy node and run the **su - omm** command to switch to user **omm**.
2. Run the **vi ${BIGDATA_HOME}/.om_profile** command and change the value of **umask** to **0077**. Save and exit.

OMS HA Status

Indicator: OMS HA Status

Description: This indicator is used to check whether the OMS two-node cluster resources are normal. You can run the **${CONTROLLER_HOME}/sbin/status-oms.sh** command to view the detailed information about the status of the OMS two-node cluster resources. If any module is abnormal, the OMS is unhealthy.

Recovery Guide:

1. Log in to the active management node and run the **su - omm** command to switch to user **omm**. Run the **${CONTROLLER_HOME}/sbin/status-oms.sh** command to check the OMS status.
2. If floatip, okerberos, and oldap are abnormal, handle the problems by referring to ALM-12002, ALM-12004, and ALM-12005 respectively.
3. If other resources are abnormal, you are advised to view the logs of the faulty modules.

   If controller resources are abnormal, view **/var/log/Bigdata/controller/controller.log** of the faulty node.

   If CEP resources are abnormal, view **/var/log/Bigdata/omm/oms/cep/cep.log** of the faulty node.

   If AOS resources are abnormal, view **/var/log/Bigdata/controller/aos/aos.log** of the faulty node.

   If feed_watchdog resources are abnormal, view **/var/log/Bigdata/watchdog/watchdog.log** of the abnormal node.

   If HTTPD resources are abnormal, view **/var/log/Bigdata/httpd/error_log** of the abnormal node.

   If FMS resources are abnormal, view **/var/log/Bigdata/omm/oms/fms/fms.log** of the abnormal node.

   If PMS resources are abnormal, view **/var/log/Bigdata/omm/oms/pms/pms.log** of the abnormal node.

   If IAM resources are abnormal, view **/var/log/Bigdata/omm/oms/iam/iam.log** of the abnormal node.

   If the GaussDB resource is abnormal, check the **/var/log/Bigdata/omm/oms/db/omm_gaussdba.log** of the abnormal node.

   If NTP resources are abnormal, view **/var/log/Bigdata/omm/oms/ha/scriptlog/ha_ntp.log** of the abnormal node.

   If Tomcat resources are abnormal, view **/var/log/Bigdata/tomcat/catalina.log** of the abnormal node.
4. If the fault cannot be rectified based on the logs, contact O&M personnel and send the collected fault logs.

Checking the Installation Directory and Data Directory

Indicator: Installation Directory and Data Directory Check

Description: This indicator checks the **lost+found** directory in the root directory of the disk partition where the installation directory (**/opt/Bigdata** by default) is located. If the directory contains the files of user **omm**, there are exceptions. When a node is abnormal, related files are stored in the **lost+found** directory. This indicator is used to check whether files are lost in such scenarios. Check the installation directory (for example, **/opt/Bigdata**) and data directory (for example, **/srv/BigData**). If any files of non-omm users exist in the two directories, the system is unhealthy.

Recovery Guide:

1. Log in to the unhealthy node and run the **su - omm** command to switch to user **omm**. Check whether files or folders of user omm exist in the **lost+found** directory.

   If the **omm** user file exists, you are advised to restore it and check again. If the **omm** user file does not exist, go to **2**.

2. Check the installation directory and data directory. Check whether the files or folders of other users exist in the installation directory and data directory. If the files and folders are manually generated temporary files, you are advised to delete them and check again.

CPU Usage

Indicator: CPU Usage

Description: This indicator is used to check whether the CPU usage exceeds the threshold. If the disk usage exceeds the threshold, the system is unhealthy.

Recovery Guide: If the indicator is abnormal, the system generates an alarm. You are advised to handle the alarm by referring to ALM-12016.

Memory Usage

Indicator: Memory Usage

Description: This indicator is used to check whether the memory usage exceeds the threshold. If the disk usage exceeds the threshold, the system is unhealthy.

Recovery Guide: If the indicator is abnormal, the system generates an alarm. You are advised to handle the alarm by referring to ALM-12018.

Host Disk Usage

Indicator: Host Disk Usage

Description: This indicator is used to check whether the host disk usage exceeds the threshold. If the disk usage exceeds the threshold, the system is unhealthy.

Recovery Guide: If the indicator is abnormal, the system generates an alarm. You are advised to handle the alarm by referring to ALM-12017.

Host Disk Write Rate

Indicator: Host Disk Write Rate

Description: This indicator is used to check the disk write rate of a host. The write rate of the host disk may vary according to the service scenario. Therefore, the value of this indicator reflects only the specified value. You need to determine whether the indicator is normal in specified service scenarios.

Recovery Guide: Determine whether the current disk write rate is normal based on the service scenario.

Host Disk Read Rate

Indicator: Host Disk Read Rate

Description: This indicator is used to check the disk read rate of a host. The read rate of the host disk may vary by service scenario. Therefore, the value of this indicator reflects only the specified value. You need to determine whether the indicator is normal in specified service scenarios.

Recovery Guide: Determine whether the current disk read rate is normal based on the service scenario.

Host Service Plane Network Status

Indicator: Host Service Plane Network Status

Description: This indicator is used to check the connectivity of the service plane network of the cluster host. If the hosts are disconnected, the cluster is unhealthy.

Recovery Guide: If the single-plane networking is used, check the IP address of the single plane. For a dual-plane network, the operation procedure is as follows:

1. Check the network connectivity between the service plane IP addresses of the active and standby management nodes.

   If the network is abnormal, go to **3**.

   If the network is normal, go to **2**.
2. Check the network connectivity between the IP address of the active management node and the IP address of the abnormal node in the cluster.
3. If the network is disconnected, contact O&M personnel to rectify the network fault to ensure that the network meets service requirements.

Host Status

Indicator: Host Status

Description: This indicator is used to check whether the host status is normal. If a node is faulty, the host is unhealthy.

Recovery Guide: If the indicator is abnormal, rectify the fault by referring to ALM-12006.

Alarm Check

Indicator: Alarm Check

Description: This indicator is used to check whether alarms exist on the host. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.9 HDFS Health Check Indicators

Average Packet Sending Time

Indicator: Average Packet Sending Time

Description: This indicator is used to collect statistics on the average time for the DataNode in the HDFS to execute SendPacket each time. If the average time is greater than 2,000,000 ns, the DataNode is unhealthy.

Recovery Guide: If the indicator is abnormal, check whether the network speed of the cluster is normal and whether the memory or CPU usage is too high. Check whether the HDFS load in the cluster is high.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the HDFS service status is normal. If a node is faulty, the host is unhealthy.

Recovery Guide: If the indicator is abnormal, check whether the health status of the KrbServer, LdapServer and ZooKeeper services are faulty. If yes, rectify the fault. Then, check whether the file writing failure is caused by HDFS SafeMode ON. Use the client to check whether data cannot be written into HDFS and locate the cause of the HDFS data writing failure. Handle the alarm following instructions in the alarm processing document.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.10 Hive Health Check Indicators

Maximum Number of Sessions Allowed by HiveServer

Indicator: Maximum Number of Sessions Allowed by HiveServer

Description: This indicator is used to check the maximum number of sessions that can be connected to Hive.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Number of Sessions Connected to HiveServer

Indicator: Number of Sessions Connected to HiveServer

Description: This indicator is used to check the number of Hive connections.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the Hive service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist on the host. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.11 Kafka Health Check Indicators

Number of Available Broker Nodes

Indicator: Number of Brokers

Description: This indicator is used to check the number of available Broker nodes in a cluster. If the number of available Broker nodes in a cluster is less than 2, the cluster is unhealthy.

Recovery Guide: If the indicator is abnormal, go to the Kafka service instance page and click the host name of the unavailable Broker instance. View the host health status in the **Overview** area. If the host health status is **Good**, rectify the fault by referring to the alarm handling suggestions in **Process Fault**. If the status is not **Good**, rectify the fault by referring to the handling procedure of the **Node Fault** alarm.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the Kafka service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If the indicator is abnormal, rectify the fault by referring to the alarm "Kafka Service Unavailable".

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.12 KrbServer Health Check Indicators

KerberosAdmin Service Availability

Indicator: KerberosAdmin Service Availability

Description: The system checks the KerberosAdmin service status. If the check result is abnormal, the KerberosAdmin service is unavailable.

Recovery Guide: If the indicator check result is abnormal, the possible cause is that the node where the KerberosAdmin service is located is faulty or the SlapdServer service is unavailable. During the KerberosAdmin service recovery, try the following operations:

1. Check whether the node where the KerberosAdmin service locates is faulty.
2. Check whether the SlapdServer service is unavailable.

KerberosServer Service Availability

Indicator: KerberosServer Service Availability

Description: The system checks the KerberosServer service status. If the check result is abnormal, the KerberosServer service is unavailable.

Recovery Guide: If the indicator check result is abnormal, the possible cause is that the node where the KerberosServer service is located is faulty or the SlapdServer service is unavailable. During the KerberosServer service recovery, try the following operations:

1. Check whether the node where the KerberosServer service locates is faulty.
2. Check whether the SlapdServer service is unavailable.

Service Health Status

Indicator: Service Status

Description: The system checks the KrbServer service status. If the check result is abnormal, the KrbServer service is unavailable.

Recovery Guide: If the indicator check result is abnormal, the possible cause is that the node where the KrbServer service resides is faulty or the LdapServer service is unavailable. For details, see the handling procedure of ALM-25500.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check the alarm information about the KrbServer service. If any alarms exist, the KrbServer service may be abnormal.

Recovery Guide: If this indicator check result is abnormal, see the related alarm document to handle the alarms.

9.8.13 LdapServer Health Check Indicators

SlapdServer Service Availability

Indicator: SlapdServer Service Availability

Description: The system checks the SlapdServer service status. If the status is abnormal, the SlapdServer service is unavailable.

Recovery Guide: If the indicator check result is abnormal, the possible cause is that the node where the SlapdServer service is located is faulty or the SlapdServer process is faulty. During the SlapdServer service recovery, try the following operations:

1. Check whether the node where the SlapdServer service locates is faulty. For details, see ALM-12006.
2. Check whether the SlapdServer process is normal. For details, see ALM-12007.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check the alarm information about the LdapServer service. If the status is abnormal, the LdapServer service is unavailable.

Recovery Guide: If the indicator check result is abnormal, the possible cause is that the node where the active LdapServer service resides is faulty or the active LdapServer process is faulty. For details, see ALM-25000.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check the alarm information about the LdapServer service. If any alarms exist, the LdapServer service may be abnormal.

Recovery Guide: If this indicator check result is abnormal, see the related alarm document to handle the alarms.

9.8.14 Loader Health Check Indicators

ZooKeeper Health Status

Indicator: ZooKeeper health status

Description: This indicator is used to check whether the ZooKeeper health status is normal. If the status is abnormal, the ZooKeeper service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

HDFS Health Status

Indicator: HDFS health status

Description: This indicator is used to check whether the HDFS health status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

DBService Health Status

Indicator: DBService Health Status

Description: This indicator is used to check whether the DBService health status is normal. If the status is abnormal, the DBService service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Yarn Health Status

Indicator: Yarn health status

Description: This indicator is used to check whether the Yarn health status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

MapReduce Health Status

Indicator: MapReduce Health Status

Description: This indicator is used to check whether the MapReduce health status is normal. If the status is abnormal, the MapReduce service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Loader Process Status

Indicator: Loader Process Status

Description: This indicator is used to check whether the Loader process is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the Loader service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist for loader. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.15 MapReduce Health Check Indicators

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the MapReduce service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.16 OMS Health Check Indicators

OMS Status Check

Indicator: OMS Status Check

Description: The OMS status check includes the HA status check and resource status check. The HA status includes **active**, **standby**, and **NULL**, indicating the active node, standby node, and unknown, respectively. The resource status includes normal, abnormal, and NULL. If the HA status is NULL, the HA status is unhealthy. If the resource status is NULL or abnormal, the resource status is unhealthy.

Table 9-27 OMS status description

Name	Description
HA state	**active**: indicates the active node. **standby**: indicates the standby node. **NULL**: unknown
Resource status	**normal**: All resources are normal. **abnormal**: indicates that abnormal resources exist. **NULL**: unknown

Recovery Guide:

1. Log in to the active management node and run the **su - omm** command to switch to user **omm**. Run the **${CONTROLLER_HOME}/sbin/status-oms.sh** command to check the status of OMS.
2. If the HA status is NULL, the system may be restarting. NULL is an intermediate state, and the HA status will automatically change to a normal state.
3. If the resource status is abnormal, certain component resources of FusionInsight Manager are abnormal. Check whether the status of components such as acs, aos cep, controller, feed_watchdog, fms, guassDB, httpd, iam, ntp, okerberos, oldap, pms, and tomcat component is normal.
4. If any Manager component resource is abnormal, see Manager component status check to rectify the fault.

Manager Component Status Check

Indicator: Manager Component Status Check

Description: This indicator is used to check the running status and HA status of Manager components. The resource running status includes **Normal** and **Abnormal**, and the resource HA status includes **Normal** and **Exception**. Manager components include Acs, Aos, Cep, Controller, feed_watchdog, Floatip, Fms, GaussDB, HeartBeatCheck, httpd, IAM, NTP, Okerberos, OLDAP, PMS, and Tomcat. If the running status and HA status is not Normal, the check result is unhealthy.

Table 9-28 Manager status description

Name	Description
Resource running status:	**Normal**: The system is running properly. **Abnormal**: The running is abnormal. **Stopped**: The task is stopped. **Unknown**: The status is unknown. **Starting**: The process is being started. **Stopping**: The task is being stopped. **Active_normal**: The active node is running properly. **Standby_normal**: The standby node is running properly. **Raising_active**: The node is being promoted to be the active node. **Lowing_standby**: The node is being set to be the standby node. **No_action**: the action does not exist. **Repairing**: The disk is being repaired. **NULL**: unknown
Resource HA status	**Normal**: the status is normal. **Exception**: indicates a fault. **Non_steady**: indicates the non-steady state. **Unknown**: unknown **NULL**: unknown

Recovery Guide:

1. Log in to the active management node and run the **su - omm** command to switch to user **omm**. Run the **${CONTROLLER_HOME}/sbin/status-oms.sh** command to check the status of OMS.
2. If floatip, okerberos, and oldap are abnormal, handle the problems by referring to ALM-12002, ALM-12004, and ALM-12005 respectively.
3. If other resources are abnormal, you are advised to view the logs of the faulty modules.

   If controller resources are abnormal, view **/var/log/Bigdata/controller/controller.log** of the faulty node.

   If CEP resources are abnormal, view **/var/log/Bigdata/omm/oms/cep/cep.log** of the faulty node.

   If AOS resources are abnormal, view **/var/log/Bigdata/controller/aos/aos.log** of the faulty node.

   If feed_watchdog resources are abnormal, view **/var/log/Bigdata/watchdog/watchdog.log** of the abnormal node.

If HTTPD resources are abnormal, view **/var/log/Bigdata/httpd/error_log** of the abnormal node.

If FMS resources are abnormal, view **/var/log/Bigdata/omm/oms/fms/fms.log** of the abnormal node.

If PMS resources are abnormal, view **/var/log/Bigdata/omm/oms/pms/pms.log** of the abnormal node.

If IAM resources are abnormal, view **/var/log/Bigdata/omm/oms/iam/iam.log** of the abnormal node.

If the GaussDB resource is abnormal, check the **/var/log/Bigdata/omm/oms/db/omm_gaussdba.log** of the abnormal node.

If NTP resources are abnormal, view **/var/log/Bigdata/omm/oms/ha/scriptlog/ha_ntp.log** of the abnormal node.

If Tomcat resources are abnormal, view **/var/log/Bigdata/tomcat/catalina.log** of the abnormal node.

4. If the fault cannot be rectified based on the logs, contact O&M personnel and send the collected fault logs.

OMA Running Status

Indicator: OMA Running Status

Description: This indicator is used to check the running status of the OMA. The status can be **Running** or **Stopped**. If the OMA is **Stopped**, the OMA is unhealthy.

Recovery Guide:

1. Log in to the unhealthy node and run the **su - omm** command to switch to user **omm**.
2. Run **${OMA_PATH}/restart_oma_app** to manually start the OMA and check again. If the check result is still unhealthy, go to **3**.
3. If manually starting the OMA cannot resolve the problem, you are advised to check the OMA logs in **/var/log/Bigdata/omm/oma/omm_agent.log**.
4. If the fault cannot be rectified based on the logs, contact O&M personnel and send the collected fault logs.

SSH Trust Between Each Node and the Active Management Node

Indicator: SSH Trust Between Each Node and the Active Management Node

Description: This indicator is used to check whether the SSH mutual trust is normal. If you can switch to another node through SSH from the active OMS node as user omm without the need of entering the password, SSH communication is normal. Otherwise, SSH communication is abnormal. In addition, if you can switch to another node through SSH from the active OMS node but fail to switch to the active OMS node from the other nodes, SSH communication is abnormal.

Recovery Guide:

1. If the indicator check result is abnormal, the SSH trust relationships between the nodes and the active management node are abnormal. In this case, check whether the permission of the **/home/omm** directory is **omm**. If non-omm users have the directory permission, the SSH trust relationship may be

abnormal. You are advised to run **chown omm:wheel** to modify the permission and check again. If the permission on the **/home/omm** directory is normal, go to **2**.

2. The SSH trust relationship exception may cause heartbeat exceptions between Controller and NodeAgent, resulting in node fault alarms. In this case, rectify the fault by referring to the handling procedure of ALM-12006.

Process Running Time

Indicator: Running Time of NodeAgent, Controller, and Tomcat

Description: This indicator is used to check the running time of the NodeAgent, Controller, and Tomcat processes. If the time is less than half an hour (1,800s), the process may have been restarted. You are advised to check the process after half an hour. If multiple check results indicate that the process runs for less than half an hour, the process is abnormal.

Recovery Guide:

1. Log in to the unhealthy node and run the **su - omm** command to switch to user **omm**.
2. Run the following command to check the PID based on the process name:

   ps -ef | grep NodeAgent
3. Run the following command to check the process startup time based on the PID:

   ps -p pid -o lstart
4. Check whether the process start time is normal. If the process restarts repeatedly, go to **5**.
5. View the related logs and analyze restart causes.

   If the runtime of NodeAgent is abnormal, check **/var/log/Bigdata/nodeagent/agentlog/agent.log**.

   If the Controller running time is abnormal, check the **/var/log/Bigdata/controller/controller.log** file.

   If the Tomcat running time is abnormal, check the **/var/log/Bigdata/tomcat/web.log** file.
6. If the fault cannot be rectified based on the logs, contact O&M personnel and send the collected fault logs.

Account and Password Expiration Check

Indicator: Account and Password Expiration Check

Description: This indicator checks the two operating system users **omm** and **ommdba** of MRS. For OS users, both the account and password expiration time must be checked. If the validity period of the account or password is not greater than 15 days, the account is abnormal.

Recovery Guide: If the validity period of the account or password is less than or equal to 15 days, contact O&M engineers.

9.8.17 Spark Health Check Indicators

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the Spark service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If the indicator is abnormal, rectify the fault by referring to ALM-28001.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.18 Storm Health Check Indicators

Number of Working Nodes

Indicator: Number of Supervisors

Description: This indicator is used to check the number of available Supervisors in a cluster. If the number of available Supervisors in a cluster is less than 1, the cluster is unhealthy.

Recovery Guide: If the indicator is abnormal, go to the Streaming service instance page and click the host name of the unavailable Supervisor instance. View the host health status in the **Overview** area. If the host health status is **Good**, rectify the fault by referring to ALM-12007 Process Faults. If the status is not **Good**, rectify the fault by referring to the handling procedure of the ALM-12006 Node Faults.

Number of Idle Slots

Indicator: Number of Idle Slots

Description: This indicator is used to check the number of idle slots in a cluster. If the number of idle slots in a cluster is less than 1, the cluster is unhealthy.

Recovery Guide: If the indicator is abnormal, go to the Storm service instance page and check the health status of the Supervisor instance. If the health status of all Supervisor instances is **Good**, you need to expand the capacity of the Core node in the cluster. If not, rectify the fault by referring to ALM-12007 Process Faults.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the Storm service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If the indicator is abnormal, rectify the fault by referring to the alarm "ALM-26051 Storm Service Unavailable".

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.19 Yarn Health Check Indicators

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether the Yarn service status is normal. If the number of NodeManager nodes cannot be obtained, the system is unhealthy.

Recovery Guide: If this indicator is abnormal, you can handle the alarm by referring to the alarm handling guide and make sure that the network is normal.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.8.20 ZooKeeper Health Check Indicators

Average ZooKeeper Request Processing Latency

Indicator: Average ZooKeeper Service Request Processing Latency

Description: This indicator is used to check the average delay for the ZooKeeper service to process requests. If the average delay is greater than 300 ms, the ZooKeeper service is unhealthy.

Recovery Guide: If the indicator is abnormal, check whether the network speed of the cluster is normal and whether the memory or CPU usage is too high.

ZooKeeper Connections Usage

Indicator: ZooKeeper Connections Usage

Description: This indicator is used to check whether the ZooKeeper memory usage exceeds 80%. If the disk usage exceeds the threshold, the system is unhealthy.

Recovery Guide: If the indicator is abnormal, you are advised to increase the memory available for the ZooKeeper service. The method of increasing the memory is as follows: Increase the value of **-Xmx** in the **GC_OPTS** configuration item in the ZooKeeper service. After the modification, restart the ZooKeeper service for the configuration to take effect.

Service Health Status

Indicator: Service Status

Description: This indicator is used to check whether ZooKeeper service status is normal. If the status is abnormal, the service is unhealthy.

Recovery Guide: If the indicator is abnormal, check whether the health status of the KrbServer and LdapServer services is faulty. If yes, rectify the fault. Log in to the ZooKeeper client, check whether the ZooKeeper data writing fails. If yes, find the failure cause based on the error message and handle the fault according to error message. Rectify the fault by following the procedure for handling ALM-13000.

Alarm Check

Indicator: Alarm Information

Description: This indicator is used to check whether alarms exist. If alarms exist, the service is unhealthy.

Recovery Guide: If this indicator is abnormal, you can rectify the fault by referring to the alarm handling guide.

9.9 Static Service Pool Management

9.9.1 Viewing the Status of a Static Service Pool

Scenario

MRS Manager manages and isolates service resources that are not running on YARN through the static service resource pool. It dynamically manages the total CPU, I/O, and memory resources that can be used by HDFS and YARN on the deployment node. The system supports time-based automatic adjustment of static service resource pools. This enables the cluster to automatically adjust the parameter values at different periods to ensure more efficient resource utilization.

On MRS Manager, you can view the monitoring metrics of the resources used by each service in the static service pool. The monitoring metrics are as follows:

- Service Total CPU Usage
- Service Total Disk I/O Read Speed
- Service Total Disk I/O Write Speed

- Service Total Memory Usage

Procedure

Step 1 On MRS Manager, click **System**. In the **Resource** area, click **Configure Static Service Pool**.

Step 2 Click **Status**.

Step 3 Check the system resource adjustment base values.

- **System Resource Adjustment Base** indicates the maximum volume of resources that can be used by each node in the cluster. If a node has only one service, the service exclusively occupies the available resources on the node. If a node has multiple services, all services share the available resources on the node.
- **CPU(%)** indicates the maximum number of CPUs that can be used by services on a node.
- **Memory(%)** indicates the maximum memory that can be used by services on a node.

Step 4 Check the cluster service resource usage.

In the chart area, select **All services** from the service drop-down list box. The resource usage status of all services in the service pool is displayed.

NOTE

Effective Configuration Group indicates the resource control configuration group used by the cluster service. By default, the **default** configuration group is used at all time every day, indicating that the cluster service can use all CPUs and 70% memory of the node.

Step 5 View the resource usage of a single service.

In the chart area, select a service from the service drop-down list box. The resource usage status of the service is displayed.

Step 6 You can set the interval for automatically refreshing the page.

The following refresh interval options are supported:

- **Refresh every 30 seconds**
- **Refresh every 60 seconds**
- **Stop refreshing**

Step 7 In the **Period** area, select a time range for viewing service resources. The options are as follows:

- Real time
- Last 3 hours
- Last 6 hours
- Last 24 hours
- Last week
- Last month
- Last 3 months

- Last 6 months
- Customize: If you select this option, you can customize the period for viewing monitoring data.

Step 8 Click **View** to view the service resource data in the corresponding time range.

Step 9 Customize a service resource report.

1. Click **Customize** and select the service source indicators to be displayed.
   - Service Total Disk I/O Read Speed
   - Service Total Memory Usage
   - Service Total Disk I/O Write Speed
   - Service Total CPU Usage
2. Click **OK** to save the selected monitoring metrics for display.

   NOTE

   Click **Clear** to cancel all the selected monitoring metrics in a batch.

Step 10 Export a monitoring report.

Click **Export**. MRS Manager will generate a report about the selected service resources in a specified time of period. Save the report.

NOTE

To view the curve charts of monitoring metrics in a specified period, click **View**.

----End

9.9.2 Configuring a Static Service Pool

Scenario

If you need to control the node resources that can be used by the cluster service or the CPU usage of the node used by the cluster in different time periods, you can adjust the resource base on MRS Manager and customize the resource configuration groups.

Prerequisites

- After the static service pool is configured, the HDFS and YARN services need to be restarted. During the restart, the services are unavailable.
- After a static service pool is configured, the maximum number of resources used by each service and role instance cannot exceed the upper limit.

Procedure

Step 1 Modify the system resource adjustment base.

1. On MRS Manager, click **System**. In the **Resource** area, click **Configure Static Service Pool**.
2. Click **Configuration**. The service pool configuration group management page is displayed.

3. In the **System Resource Adjustment Base** area, change the values of **CPU(%)** and **Memory(%)**.

   Modifying **System Resource Adjustment Base** limits the maximum physical CPU and memory resource percentage of nodes that can be used by the Flume, HBase, HDFS, Impala and YARN services. If multiple services are deployed on the same node, the maximum physical resource usage of all services cannot exceed the adjusted CPU or memory usage.

4. Click **Next**.

   If you need to modify the parameters again, click **Previous** in the lower part of the page.

Step 2 Modify the **default** configuration group of the service pool.

1. In the **Service Pool Configuration** table, set **CPU LIMIT(%)**, **CPU SHARE(%)**, **I/O(%)**, and **Memory(%)** for the Flume, HBase, HDFS, Impala, and YARN services.

   NOTE

   - The sum of **CPU LIMIT(%)** used by all services can exceed 100%.
   - The sum of **CPU SHARE(%)** and **I/O(%)** used by all services must be 100%. For example, if CPU resources are allocated to the HDFS and Yarn services, the total CPU resources allocated to the two services are 100%.
   - The sum of **Memory(%)** used by all services can be greater than, smaller than, or equal to 100%.
   - **Memory(%)** cannot take effect dynamically and can only be modified in the default configuration group.

2. Click in the blank area of the page to complete the editing. MRS Manager generates the correct values of service pool parameters in the **Detailed Configuration** area based on the cluster hardware resources and allocation information.
3. You can click the edit icon on the right of **Detailed Configuration** to modify the parameter values of the service pool based on service requirements.

   In the **Service Pool Configuration** area, click the specified service name. The **Detailed Configuration** area displays only the parameters of the service. Manual changing of parameter values does not refresh the service resource usage. In added configuration groups, the configuration group numbers of the parameters that take effect dynamically will be displayed. For example, **HBase: RegionServer: dynamic-config1.RES_CPUSET_PERCENTAGE**. The parameter functions do not change.

   Table 9-29 Parameters of the static service pool

Parameter	Description
– RES_CPUSET_PERCENTAGE – dynamic-configX.RES_CPUSET_PERCENTAGE	Configures the service CPU percentage.

Parameter	Description
– RES_CPU_SHARE – dynamic-configX.RES_CPU_SHARE	Configures the service CPU share.
– RES_BLKIO_WEIGHT – dynamic-configX.RES_BLKIO_WEIGHT	Configures service I/O usage.
HBASE_HEAPSIZE	Configures the maximum JVM memory for RegionServer.
HADOOP_HEAPSIZE	Configures the maximum JVM memory of a DataNode.
yarn.nodemanager.resource.memory-mb	Configures the memory that can be used by NodeManager on the current node.
dfs.datanode.max.locked.memory	Configures the maximum memory that can be used by a DataNode as the HDFS cache.
FLUME_HEAPSIZE	Configures the maximum JVM memory that can be used by each Flume instance.
IMPALAD_MEM_LIMIT	Configures the maximum memory that can be used by an Impalad instance.

Step 3 Add a customized resource configuration group.

1. Determine whether to automatically adjust resource configurations based on the time.

   If yes, go to **Step 3.2**.

   If no, go to **Step 4**.

2. Click  to add a resource configuration group. In the **Scheduling Time** area, click the edit icon. The time policy configuration page is displayed.

   Modify the following parameters based on service requirements and click **OK**.

   – **Repeat**: If selected, the resource configuration group runs repeatedly based on the scheduling period. If not selected, set the date and time when the configuration of the group of resources can be applied.

   – **Repeat Policy**: can be set to **Daily**, **Weekly**, and **Monthly**. This parameter is valid only when **Repeat** is selected.

   – **Between**: indicates the time period between the start time and end time when the resource configuration is applied. Set a unique time range. If the time range overlaps with that of an existing group of resource configuration, the time range cannot be saved. This parameter is valid only when **Repeat** is selected.

NOTE

- The **default** group of resource configuration takes effect in all undefined time segments.
- The newly added resource group is a parameter set that takes effect dynamically in a specified time range.
- The newly added resource group can be deleted. A maximum of four resource configuration groups that take effect dynamically can be added.
- Select a repetition policy. If the end time is earlier than the start time, the next day is labeled by default. For example, if a validity period ranges from 22:00 to 06:00, the customized resource configuration takes effect from 22:00 on the current day to 06:00 on the next day.
- If the repeat policy types of multiple configuration groups are different, the time ranges can overlap. The policy types are listed as follows by priority from low to high: daily, weekly, and monthly. The following is an example. There are two resource configuration groups using the monthly and daily policies, respectively. Their application time ranges in a day overlap as follows: [04:00 to 07:00] and [06:00 to 08:00]. In this case, the configuration of the group that uses the monthly policy prevails.
- If the repeat policy types of multiple resource configuration groups are the same, the time ranges of different dates can overlap. For example, if there are two weekly scheduling groups, you can set the same time range on different day for them, such as to 04:00 to 07:00, on Monday and Wednesday, respectively.

3. On the **Service Pool Configuration** page, modify the resource configuration of each service. Click the blank area on the page to complete the editing, and go to **Step 4**.

   You can click the edit icon on the right of **Service Pool Configuration** to modify the parameters. Click the edit icon in the **Detailed Configuration** area to manually update the parameter values generated by the system based on service requirements.

Step 4 Saves the settings.

Click **Save**. In the **Save Configuration** dialog box, select **Restart the affected services or instances**. Click **OK** to save the settings and restart related services.

Operation succeeded is displayed. Click **Finish**. The service is started.

----End

9.10 Tenant Management

9.10.1 Overview

Definition

An MRS cluster provides various resources and services for multiple organizations, departments, or applications to share. The cluster provides tenants as a logical entity to use these resources and services. A mode involving different tenants is called multi-tenant mode. Currently, only the analysis cluster supports tenant management.

Principles

The MRS cluster provides the multi-tenant function. It supports a layered tenant model and allows dynamic adding or deleting of tenants to isolate resources. It dynamically manages and configures tenants' computing and storage resources.

The computing resources indicate tenants' Yarn task queue resources. The task queue quota can be modified, and the task queue usage status and statistics can be viewed.

The storage resources can be stored on HDFS. You can add and delete the HDFS storage directories of tenants, and set the quotas of file quantity and the storage space of the directories.

As the unified tenant management platform of MRS clusters, MRS Manager provides enterprises with time-tested multi-tenant management models, enabling centralized tenant and service management. Tenants can create and manage tenants in a cluster based on service requirements.

- Roles, computing resources, and storage resources are automatically created when tenants are created. By default, all permissions of the new computing resources and storage resources are allocated to a tenant's roles.
- Permissions to view the current tenant's resources, add a subtenant, and manage the subtenant's resources are granted to the tenant's roles by default.
- After you have modified the tenant's computing or storage resources, permissions of the tenant's roles are automatically updated.

MRS Manager supports a maximum of 512 tenants. The tenants that are created by default in the system contain **default**. Tenants that are in the topmost layer with the default tenant are called level-1 tenants.

Resource Pools

Yarn task queues support only the label-based scheduling policy. This policy enables Yarn task queues to associate NodeManagers that have specific node labels. In this way, Yarn tasks run on specified nodes so that tasks are scheduled and certain hardware resources are utilized. For example, Yarn tasks requiring a large memory capacity can run on nodes with a large memory capacity by means of label association, preventing poor service performance.

In an MRS cluster, the tenant logically divides Yarn cluster nodes to combine multiple NodeManagers into a resource pool. Yarn task queues can be associated with specified resource pools by configuring queue capacity policies, ensuring efficient and independent resource utilization in the resource pools.

MRS Manager supports a maximum of 50 resource pools. The system has a **Default** resource pool.

9.10.2 Creating a Tenant

Scenario

You can create a tenant on MRS Manager to specify the resource usage.

Prerequisites

- A tenant name has been planned. The name must not be the same as that of a role or Yarn queue that exists in the current cluster.
- If a tenant requires storage resources, a storage directory has been planned based on service requirements, and the planned directory does not exist under the HDFS directory.
- The resources that can be allocated to the current tenant have been planned and the sum of the resource percentages of direct sub-tenants under the parent tenant at every level does not exceed 100%.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 Click **Create Tenant**. On the page that is displayed, configure tenant properties.

Table 9-30 Tenant parameters

Parameter	Description
Name	Specifies the name of the current tenant. The value consists of 1 to 20 characters, and can contain letters, numbers, and underscores (_).
Tenant Type	The options include **Leaf** and **Non-leaf**. If **Leaf** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. If **Non-leaf** is selected, sub-tenants can be added to the current tenant.
Dynamic Resources	Specifies the dynamic computing resources for the current tenant. The system automatically creates a task queue named after the tenant name in Yarn. When dynamic resources are not **Yarn**, the system does not automatically create a task queue.
Default Resource Pool Capacity (%)	Specifies the percentage of the computing resources used by the current tenant in the **default** resource pool.
Default Resource Pool Max. Capacity (%)	Specifies the maximum percentage of the computing resources used by the current tenant in the **default** resource pool.
Storage Resource	Specifies storage resources for the current tenant. The system automatically creates a file folder named after the tenant name in the **/tenant** directory. When a tenant is created for the first time, the system automatically creates the **/tenant** directory in the HDFS root directory. If storage resources are not **HDFS**, the system does not create a storage directory under the root directory of HDFS.

Parameter	Description
Space Quota (MB)	Specifies the quota for HDFS storage space used by the current tenant. The value ranges from **1** to **8796093022208**. The unit is MB. This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used. If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the oHDFS physical disk. **NOTE** To ensure data reliability, one copy of a file is automatically generated when the file is stored in HDFS. That is, two copies of the same file are stored by default. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value of **Storage Space Quota** is set to **500**, the actual space for storing files is about 250 MB (500/2 = 250).
Storage Path	Specifies the tenant's HDFS storage directory. The system automatically creates a file folder named after the tenant name in the **/tenant** directory by default. For example, the default HDFS storage directory for tenant **ta1** is **tenant/ta1**. When a tenant is created for the first time, the system automatically creates the **/tenant** directory in the HDFS root directory. The storage path is customizable.
Service	Specifies other service resources associated with the current tenant. HBase is supported. To configure this parameter, click **Associate Services**. In the dialog box that is displayed, set **Service** to **HBase**. If **Association Mode** is set to **Exclusive**, service resources are occupied exclusively. If **share** is selected, service resources are shared.
Description	Specifies the description of the current tenant.

Step 3 Click **OK** to save the settings.

It takes a few minutes to save the settings. If the **Tenant created successfully** is displayed in the upper-right corner, the tenant is added successfully.

NOTE

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. The role and its permissions are controlled by the system automatically and cannot be controlled manually under **Manage Role**.
- If you want to use the tenant, create a system user and assign the Manager_tenant role and the role corresponding to the tenant to the user. For details, see **Creating a User**.

----End

Related Tasks

Viewing an added tenant

Step 1 On MRS Manager, click **Tenant**.

Step 2 In the tenant list on the left, click the name of the added tenant.

The **Summary** tab is displayed on the right by default.

Step 3 View **Basic Information**, **Resource Quota**, and **Statistics** of the tenant.

If HDFS is in the **Stopped** state, **Available** and **Used** of **Space** in **Resource Quota** are **unknown.**

----End

9.10.3 Creating a Sub-tenant

Scenario

You can create a sub-tenant on MRS Manager if the resources of the current tenant need to be further allocated.

Prerequisites

- A parent tenant has been added.
- A tenant name has been planned. The name must not be the same as that of a role or Yarn queue that exists in the current cluster.
- If a sub-tenant requires storage resources, a storage directory has been planned based on service requirements, and the planned directory does not exist under the storage directory of the parent tenant.
- The resources that can be allocated to the current tenant have been planned and the sum of the resource percentages of direct sub-tenants under the parent tenant at every level does not exceed 100%.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 In the tenant list on the left, move the cursor to the tenant node to which a sub-tenant is to be added. Click **Create sub-tenant**. On the displayed page, configure the sub-tenant attributes according to the following table:

Table 9-31 Sub-tenant parameters

Parameter	Description
Parent tenant	Specifies the name of the parent tenant.
Name	Specifies the name of the current tenant. The value consists of 1 to 20 characters, and can contain letters, numbers, and underscores (_).

Parameter	Description
Tenant Type	The options include **Leaf** and **Non-leaf**. If **Leaf** is selected, the current tenant is a leaf tenant and no sub-tenant can be added. If **Non-leaf** is selected, sub-tenants can be added to the current tenant.
Dynamic Resources	Specifies the dynamic computing resources for the current tenant. The system automatically creates a task queue named after the sub-tenant name in the Yarn parent queue. When dynamic resources are not **Yarn**, the system does not automatically create a task queue. If the parent tenant does not have dynamic resources, the sub-tenant cannot use dynamic resources.
Default Resource Pool Capacity (%)	Specifies the percentage of the resources used by the current tenant. The base value is the total resources of the parent tenant.
Default Resource Pool Max. Capacity (%)	Specifies the maximum percentage of the computing resources used by the current tenant. The base value is the total resources of the parent tenant.
Storage Resource	Specifies storage resources for the current tenant. The system automatically creates a file in the HDFS parent tenant directory. The file is named the same as the name of the sub-tenant. If storage resources are not **HDFS**, the system does not create a storage directory under the root directory of HDFS. If the parent tenant does not have storage resources, the sub-tenant cannot use storage resources.
Space Quota (MB)	Specifies the quota for HDFS storage space used by the current tenant. The minimum value is 1, and the maximum value is the total storage quota of the parent tenant. The unit is MB. This parameter indicates the maximum HDFS storage space that can be used by a tenant, but does not indicate the actual space used. If the value is greater than the size of the HDFS physical disk, the maximum space available is the full space of the oHDFS physical disk. If the quota is greater than the quota of the parent tenant, the actual storage capacity is subject to the quota of the parent tenant. **NOTE** To ensure data reliability, one copy of a file is automatically generated when the file is stored in HDFS. That is, two copies of the same file are stored by default. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value is set to **500**, the actual space for storing files is about 250 MB (500/2 = 250).

Parameter	Description
Storage Path	Specifies the tenant's HDFS storage directory. The system automatically creates a file folder named after the sub-tenant name in the directory of the parent tenant by default. For example, if the sub-tenant is **ta1s** and the parent directory is **tenant/ta1**, the system sets this parameter for the sub-tenant to **tenant/ta1/ta1s**. The storage path is customizable in the parent directory. The parent directory for the storage path must be the storage directory of the parent tenant.
Service	Specifies other service resources associated with the current tenant. HBase is supported. To configure this parameter, click **Associate Services**. In the dialog box that is displayed, set **Service** to **HBase**. If **Association Mode** is set to **Exclusive**, service resources are occupied exclusively. If **share** is selected, service resources are shared.
Description	Specifies the description of the current tenant.

Step 3 Click **OK** to save the settings.

It takes a few minutes to save the settings. If the **Tenant created successfully** is displayed in the upper-right corner, the tenant is added successfully. The tenant is created successfully.

NOTE

- Roles, computing resources, and storage resources are automatically created when tenants are created.
- The new role has permissions on the computing and storage resources. The role and its permissions are controlled by the system automatically and cannot be controlled manually under **Manage Role**.
- When using this tenant, create a system user and assign the user a related tenant role. For details, see **Creating a User**.

----End

9.10.4 Deleting a tenant

Scenario

You can delete a tenant that is not required on MRS Manager.

Prerequisites

- A tenant has been added.
- You have checked whether the tenant to be deleted has sub-tenants. If the tenant has sub-tenants, delete them; otherwise, you cannot delete the tenant.

- The role of the tenant to be deleted cannot be associated with any user or user group. For details about how to cancel the binding between a role and a user, see **Modifying User Information**.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 In the tenant list on the left, move the cursor to the tenant node to be deleted and click **Delete**.

The **Delete Tenant** dialog box is displayed. If you want to save the tenant data, select **Reserve the data of this tenant**. Otherwise, the tenant's storage space will be deleted.

Step 3 Click OK to save the settings.

It takes a few minutes to save the configuration. After the tenant is deleted successfully, the role and storage space of the tenant are also deleted.

NOTE

- After the tenant is deleted, the task queue of the tenant still exists in Yarn.
- If you choose not to reserve data when deleting the parent tenant, data of sub-tenants is also deleted if the sub-tenants use storage resources.

----End

9.10.5 Managing a Tenant Directory

Scenario

You can manage the HDFS storage directory used by a specific tenant on MRS Manager. The management operations include adding a tenant directory, modifying the directory file quota, modifying the storage space, and deleting a directory.

Prerequisites

A tenant associated with HDFS storage resources has been added.

Procedure

- Viewing a tenant directory
  a. On MRS Manager, click **Tenant**.
  b. In the tenant list on the left, click the target tenant.
  c. Click the **Resource** tab.
  d. View the **HDFS Storage** table.
     - The Quota column indicates the quantity quotas of files and directories.
     - The **Storage Space Quota** column indicates the storage space size of the tenant directory.

- Adding a tenant directory
  a. On MRS Manager, click **Tenant**.
  b. In the tenant list on the left, click the tenant whose HDFS storage directory needs to be added.
  c. Click the **Resource** tab.
  d. In the **HDFS Storage** table, click **Create Directory**.
     - In **Parent Directory**, select a storage directory of a parent tenant.
       This parameter applies only to sub-tenants. If the parent tenant has multiple directories, select any of them.
     - Set **Path** to a tenant directory path.
       NOTE
       - If the current tenant is not a sub-tenant, the new path is created in the HDFS root directory.
       - If the current tenant is a sub-tenant, the new path is created in the specified directory.

       A complete HDFS storage directory can contain a maximum of 1,023 characters. An HDFS directory name contains digits, letters, spaces, and underscores (_). The name cannot start or end with a space.
     - Set **Quota** to the quotas of file and directory quantity.
       Maximum Number of Files/Directories is optional. Its value ranges from **1** to **9223372036854775806**.
     - Set **Storage Space Quota** to the storage space size of the tenant directory.
       The value of **Storage Space Quota** ranges from **1** to **8796093022208**.
       NOTE
       To ensure data reliability, one copy of a file is automatically generated when the file is stored in HDFS. That is, two copies of the same file are stored by default. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value of **Storage Space Quota** is set to **500**, the actual space for storing files is about 250 MB (500/2 = 250).
  e. Click **OK**. The system creates tenant directories in the HDFS root directory.
- Modify a tenant directory.
  a. On MRS Manager, click **Tenant**.
  b. In the tenant list on the left, click the tenant whose HDFS storage directory needs to be modified.
  c. Click the **Resource** tab.
  d. In the **HDFS Storage** table, click **Modify** in the **Operation** column of the specified tenant directory.
     - Set **Quota** to the quotas of file and directory quantity.
       Maximum Number of Files/Directories is optional. Its value ranges from **1** to **9223372036854775806**.

- Set **Storage Space Quota** to the storage space size of the tenant directory.

  The value of **Storage Space Quota** ranges from **1** to **8796093022208**.

  NOTE

  To ensure data reliability, one copy of a file is automatically generated when the file is stored in HDFS. That is, two copies of the same file are stored by default. The HDFS storage space indicates the total disk space occupied by all these copies. For example, if the value of **Storage Space Quota** is set to **500**, the actual space for storing files is about 250 MB (500/2 = 250).

e. Click **OK**.

- Delete a tenant directory.

  a. On MRS Manager, click **Tenant**.

  b. In the tenant list on the left, click the tenant whose HDFS storage directory needs to be deleted.

  c. Click the **Resource** tab.

  d. In the **HDFS Storage** table, click **Delete** in the **Operation** column of the specified tenant directory.

  The default HDFS storage directory set during tenant creation cannot be deleted. Only the newly added HDFS storage directory can be deleted.

  e. Click **OK**.

9.10.6 Restoring Tenant Data

Scenario

Tenant data is stored on Manager and in cluster components by default. When components are restored from faults or reinstalled, some tenant configuration data may be abnormal. In this case, you can manually restore the tenant data.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 In the tenant list on the left, click a tenant node.

Step 3 Check the status of the tenant data.

1. In **Summary**, check the color of the circle on the left of **Basic Information**. Green indicates that the tenant is available and gray indicates that the tenant is unavailable.
2. Click **Resources** and check the circle on the left of **Yarn** or **HDFS Storage**. Green indicates that the resource is available, and gray indicates that the resource is unavailable.
3. Click **Service Association** and check the **Status** column of the associated service table. **Good** indicates that the component can provide services for the associated tenant. **Bad** indicates that the component cannot provide services for the tenant.

4. If any check result is abnormal, go to **Step 4** to restore tenant data.

Step 4 Click **Restore Tenant Data**.

Step 5 In the **Restore Tenant Data** window, select one or more components whose data needs to be restored. Click **OK**. The system automatically restores the tenant data.

----End

9.10.7 Creating a Resource Pool

Scenario

In an MRS cluster, users can logically divide Yarn cluster nodes to combine multiple NodeManagers into a Yarn resource pool. Each NodeManager belongs to one resource pool only. The system contains a **Default** resource pool by default. All NodeManagers that are not added to customized resource pools belong to this resource pool.

You can create a customized resource pool on MRS Manager and add hosts that have not been added to other customized resource pools to it.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 Click the **Resource Pools** tab.

Step 3 Click **Add Resource Pool**.

Step 4 In **Create Resource Pool**, set the properties of the resource pool.

- **Name**: Enter a name for the resource pool. The name of the newly created resource pool cannot be **Default**.

  The name consists of 1 to 20 characters and can contain digits, letters, and underscores (_) but cannot start with an underscore (_).

- **Hosts**: In the host list on the left, select the name of a specified host and click 

to add the selected host to the resource pool. Only hosts in the cluster can be selected. The host list of a resource pool can be left blank.

Step 5 Click **OK**.

Step 6 After a resource pool is created, users can view the **Name**, **Members**, **Type**, **vCore** and **Memory** in the resource pool list. Hosts that are added to the customized resource pool are no longer members of the **Default** resource pool.

----End

9.10.8 Modifying a Resource Pool

Scenario

You can modify members of an existing resource pool on MRS Manager.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 Click the **Resource Pools** tab.

Step 3 Locate the row that contains the specified resource pool, and click **Modify** in the **Operation** column.

Step 4 In **Modify Resource Pool**, modify **Added Hosts**.

- Adding a host: Select the name of a specified host in host list on the left and click 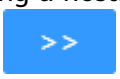

to add the selected host to the resource pool.
- Deleting a host: In the host list on the right, select the name of a specified host and click 

to add the selected host to the resource pool. The host list of a resource pool can be left blank.

Step 5 Click **OK**.

----End

9.10.9 Deleting a Resource Pool

Scenario

You can delete an existing resource pool on MRS Manager.

Prerequisites

- Any queue in a cluster cannot use the resource pool to be deleted as the default resource pool. Before deleting the resource pool, cancel the default resource pool. For details, see **Configuring a Queue**.
- Resource distribution policies of all queues have been cleared from the resource pool being deleted. For details, see **Clearing Configuration of a Queue**.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 Click the **Resource Pools** tab.

Step 3 Locate the row that contains the specified resource pool, and click **Delete** in the **Operation** column.

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

----End

9.10.10 Configuring a Queue

Scenario

This section describes how to modify the queue configuration for a specified tenant on MRS Manager.

Prerequisites

A tenant associated with Yarn and allocated dynamic resources has been added.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 Click the **Dynamic Resource Plan** tab.

Step 3 Click the **Queue Configuration** tab.

Step 4 In the tenant queue table, click **Modify** in the **Operation** column of the specified tenant queue.

NOTE

In the tenant list on the left of the **Tenant Management** tab, click the target tenant. In the window that is displayed, choose **Resource**. On the page that is displayed, click the edit icon to open the queue modification page.

Table 9-32 Queue configuration parameters

Parameter	Description
Maximum Application	Specifies the maximum number of applications. The value ranges from 1 to 2147483647.
Maximum AM Resource Percent	Specifies the maximum percentage of resources that can be used to run the ApplicationMaster in a cluster. The value ranges from 0 to 1.
Minimum User Limit Percent (%)	Specifies the minimum percentage of resources consumed by a user. The value ranges from 0 to 100.
User Limit Factor	Specifies the limit factor of the maximum user resource usage. The maximum user resource usage percentage can be obtained by multiplying the limit factor with the percentage of the tenant's actual resource usage in the cluster. The minimum value is **0**.
Status	Specifies the current status of a resource plan. The values are **Running** and **Stopped**.
Default Resource Pool	Specifies the resource pool used by a queue. The default value is **Default**. If you want to change the resource pool, configure the queue capacity first. For details, see **Configuring the Queue Capacity Policy of a Resource Pool**.

----End

9.10.11 Configuring the Queue Capacity Policy of a Resource Pool

Scenario

After a resource pool is added, the capacity policies of available resources need to be configured for Yarn task queues. This ensures that tasks in the resource pool are running properly. Each queue can be configured with the queue capacity policy of only one resource pool. Users can view the queues in any resource pool and configure queue capacity policies. After the queue policies are configured, Yarn task queues and resource pools are associated.

You can configure queue policies on MRS Manager.

Prerequisites

- A resource pool has been added.
- The task queues are not associated with other resource pools. By default, all queues are associated with the **Default** resource pool.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 Click the **Dynamic Resource Plan** tab.

Step 3 In **Resource Pools**, select a specified resource pool.

Available Resource Quota: indicates that all resources in each resource pool are available for queues by default.

Step 4 Locate the specified queue in the **Resource Allocation** table, and click **Modify** in the **Operation** column.

Step 5 In **Modify Resource Allocation**, configure the resource capacity policy of the task queue in the resource pool.

- **Capacity (%)**: specifies the percentage of the current tenant's computing resource usage.
- **Maximum Capacity (%)**: specifies the percentage of the current tenant's maximum computing resource usage.

Step 6 Click **OK** to save the settings.

----End

9.10.12 Clearing Configuration of a Queue

Scenario

Users can clear the configuration of a queue on MRS Manager when the queue does not need resources from a resource pool or if a resource pool needs to be

disassociated from the queue. Clearing queue configurations means that the resource capacity policy of the queue is canceled.

Prerequisites

If a queue is to be unbound from a resource pool, this resource pool cannot serve as the default resource pool of the queue. Therefore, you must first change the default resource pool of the queue to another one. For details, see **Configuring a Queue**.

Procedure

Step 1 On MRS Manager, click **Tenant**.

Step 2 Click the **Dynamic Resource Plan** tab.

Step 3 In **Resource Pools**, select a specified resource pool.

Step 4 Locate the specified queue in the **Resource Allocation** table, and click **Clear** in the **Operation** column.

In the **Clear Queue Configuration** dialog box, click **OK** to clear the queue configuration in the current resource pool.

NOTE

If no resource capacity policy is configured for a queue, the clearing function is unavailable for the queue by default.

----End

9.11 Backup and Restoration

9.11.1 Introduction

Purpose

MRS Manager provides backup and restoration for user data and system data. The backup function is provided based on components to back up Manager data (including OMS data and LdapServer data), Hive user data, component metadata saved in DBService, and HDFS metadata.

Backup and restoration tasks are performed in the following scenarios:

- Routine backup is performed to ensure the data security of the system and components.
- If the system is faulty, the data backup can be used to recover the system.
- If the active cluster is completely faulty, a mirror cluster identical to the active cluster needs to be created. You can use the backup data to restore the active cluster.

Table 9-33 Backing up metadata

Backup Type	Backup Content
OMS	Database data (excluding alarm data) and configuration data in the cluster management system to be backed up by default
LdapServer	User information, including the username, password, key, password policy, and group information
DBService	Metadata of the components (Hive) managed by DBService
NameNode	HDFS metadata.

Principles

Task

Before backup or restoration, you need to create a backup or restoration task and set task parameters, such as the task name, backup data source, and type of backup file save path. Data backup and restoration can be performed by executing backup and restoration tasks. When the Manager is used to recover the data of HDFS, HBase, Hive, and NameNode, no cluster can be accessed.

Each backup task can back up data of different data sources and generates an independent backup file for each data source. All the backup files generated in each backup task form a backup file set, which can be used in restoration tasks. Backup data can be stored on Linux local disks, local cluster HDFS, and standby cluster HDFS. The backup task provides the full backup or incremental backup policies. HDFS and Hive backup tasks support the incremental backup policy, while OMS, LdapServer, DBService, and NameNode backup tasks support only the full backup policy.

NOTE

Task execution rules:

- If a task is being executed, the task cannot be executed repeatedly and other tasks cannot be started at the same time.
- The interval at which a periodical task is automatically executed must be greater than 120s; otherwise, the task is postponed and will be executed in the next period. Manual tasks can be executed at any interval.
- When a periodic task is to be automatically executed, the current time cannot be 120s later than the task start time; otherwise, the task is postponed and executed in the next period.
- When a periodic task is locked, it cannot be automatically executed and needs to be manually unlocked.
- Before an OMS, LdapServer, DBService, or NameNode backup task starts, ensure that the LocalBackup partition on the active management node has more than 20 GB available space. Otherwise, the backup task cannot be started.
- When you are planning backup and restoration tasks, select the data to be backed up or restored strictly based on the service logic, data store structure, and database or table association. The system creates a default periodic backup task **default** whose execution interval is 24 hours to perform full backup of OMS, LdapServer, DBService, and NameNode data to the Linux local disk.

Specifications

Table 9-34 Backup and restoration feature specifications

Item	Specifications
Maximum number of backup or restoration tasks	100
Number of concurrent running tasks	1
Maximum number of waiting tasks	199
Maximum size of backup files on a Linux local disk (GB)	600

Table 9-35 Specifications of the **default** task

Item	OMS	LdapServer	DBService	NameNode
Backup period	1 hour			
Maximum number of copies	2			
Maximum size of a backup file	10 MB	20 MB	100 MB	1.5 GB

Item	OMS	LdapServer	DBService	NameNode
Maximum size of disk space used	20 MB	40 MB	200 MB	3 GB
Save path of backup data	*Data save path*/**LocalBackup/** of the active and standby management nodes			

NOTE

The backup data of the **default** task must be periodically transferred and saved outside the cluster based on the enterprise O&M requirements.

9.11.2 Backing Up Metadata

Scenario

To ensure the security of metadata either on a routine basis or before and after performing critical metadata operations (such as scale-out, scale-in, patch installation, upgrades, and migration), metadata must be backed up. The backup data can be used to recover the system if an exception occurs or if the operation has not achieved the expected result. This minimizes the adverse impact on services. Metadata includes data of OMS, LdapServer, DBService, and NameNode. MRS Manager data to be backed up includes OMS data and LdapServer data.

By default, metadata backup is supported by the **default** task. This section describes how to create a backup task and back up metadata on MRS Manager. Both automatic backup tasks and manual backup tasks are supported.

Prerequisites

- A standby cluster for backing up data has been created, and the network is connected. The inbound rules of the two security groups on the peer cluster have been added to the two security groups in each cluster to allow all access requests of all protocols and ports of all ECSs in the security groups.
- The backup type, period, policy, and other specifications have been planned based on the service requirements and you have checked whether *Data storage path*/**LocalBackup/** has sufficient space on the active and standby management nodes.

Procedure

Step 1 Create a backup task.

1. On MRS Manager, choose **System** > **Back Up Data**.
2. Click **Create Backup Task**.

Step 2 Configure a backup policy.

1. Set **Task Name** to the name of the backup task.

2. Set **Backup Mode** to the type of the backup task. **Periodic** indicates that the backup task is periodically executed. **Manual** indicates that the backup task is manually executed.

   To create a periodic backup task, set the following parameters:

   - **Started**: indicates the time when the task is started for the first time.
   - **Period**: indicates the task execution interval. The options include **By hour** and **By day**.
   - **Backup Policy**: indicates the volume of data to be backed up in each task execution. The options include **Full backup at the first time and incremental backup later**, **Full backup every time**, and **Full backup once every n times**. If you select **Full backup once every n times**, you need to specify the value of **n**.

Step 3 Select backup sources.

In the **Configuration** area, select **OMS** and **LdapServer** under **Metadata**.

Step 4 Set backup parameters.

1. Set **Path Type** of **OMS** and **LdapServer** to a backup directory type.

   The following backup directory types are supported:

   - **LocalDir**: indicates that the backup files are stored on the local disk of the active management node and the standby management node automatically synchronizes the backup files. By default, the backup files are stored in *Data storage path***/LocalBackup/**. If you select **LocalDir**, you need to set the maximum number of copies to specify the number of backup files that can be retained in the backup directory.
   - **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster. If you select **SFTP**, set the following parameters:

     - **Target Path**: indicates the HDFS directory for storing the backup files. The save path cannot be an HDFS hidden directory, such as a snapshot or recycle bin directory, or a default system directory.
     - **Max Number of Backup Copies**: indicates the number of backup files that can be retained in the backup directory.
     - **Target Instance Name**: indicates the NameService name of the backup directory. The default value is **hacluster**.

2. Click **OK**.

Step 5 Execute the backup task.

In the **Operation** column of the created task in the backup task list, click **Back Up Now** if **Backup Mode** is set to **Periodic** or click **Start** If **Backup Mode** is set to **Manual** to execute the backup task.

After the backup task is executed, the system automatically creates a subdirectory for each backup task in the backup directory. The format of the subdirectory name is *Backup task name_Task creation time*, and the subdirectory is used to save data source backup files. The format of the backup file name is *Version_Data source_Task execution time*.**tar.gz**.

----End

9.11.3 Restoring Metadata

Scenario

You need to restore metadata in the following scenarios: A user modifies or deletes data unexpectedly, data needs to be retrieved, system data becomes abnormal or does not achieve the expected result, all modules are faulty, and data is migrated to a new cluster.

This section describes how to restore metadata on MRS Manager. Only manual restoration tasks are supported.

NOTICE

- Data restoration can be performed only when the system version is consistent with that during data backup.
- To restore the data when services are normal, manually back up the latest management data first and then restore the data. Otherwise, the data that is generated after the data backup and before the data restoration will be lost.
- Use the OMS data and LdapServer data backed up at the same time to restore data. Otherwise, the service and operation may fail.
- By default, MRS clusters use DBService to store Hive metadata.

Impact on the System

- After the data is restored, the data generated between the backup time and restoration time is lost.
- After the data is restored, the configuration of the components that depend on DBService may expire and these components need to be restarted.

Prerequisites

- The data in the OMS and LdapServer backup files has been backed up at the same time.
- The status of the OMS resources and the LdapServer instances is normal. If the status is abnormal, data restoration cannot be performed.
- The status of the cluster hosts and services is normal. If the status is abnormal, data restoration cannot be performed.
- The cluster host topologies during data restoration and data backup are the same. If the topologies are different, data restoration cannot be performed and you need to back up data again.
- The services added to the cluster during data restoration and data backup are the same. If the services are different, data restoration cannot be performed and you need to back up data again.
- The status of the active and standby DBService instances is normal. If the status is abnormal, data restoration cannot be performed.
- The upper-layer applications depending on the MRS cluster have been stopped.

- On MRS Manager, you have stopped all the NameNode role instances whose data is to be recovered. Other HDFS role instances are running properly. After data is recovered, the NameNode role instances need to be restarted and cannot be accessed before the restart.
- You have checked whether NameNode backup files have been stored in the *Data save path*/**LocalBackup/** directory on the active management node.

Procedure

Step 1 Check the location of backup data.

1. On MRS Manager, choose **System** > **Back Up Data**.
2. In the row where the specified backup task resides, choose **More** > **View History** in the **Operation** column to display the historical execution records of the backup task. In the window that is displayed, select a success record and click **View Backup Path** in the corresponding column to view its backup path information. Find the following information:
   - **Backup Object**: indicates the backup data source.
   - **Backup Path**: indicates the full path where backup files are stored.
3. Select the correct path, and manually copy the full path of backup files in **Backup Path**.

Step 2 Create a restoration task.

1. On MRS Manager, choose System > Recovery Management.
2. On the page that is displayed, click **Create Restoration Task**.
3. Set **Task Name** to the name of the restoration task.

Step 3 Select restoration sources.

In **Configuration**, select the metadata component whose data is to be restored.

Step 4 Set the restoration parameters.

1. Set **Path Type** to a backup directory type.
2. The settings vary according to backup directory types:
   - **LocalDir**: indicates that the backup files are stored on the local disk of the active management node. If you select **LocalDir**, you need to set **Source Path** to specify the full path of the backup file. For example, *Data storage path*/**LocalBackup/***Backup task name_Task creation time*/*Data source_Task execution time*/*Version number_Data source_Task execution time*.**tar.gz**.
   - **LocalHDFS**: indicates that the backup files are stored in the HDFS directory of the current cluster. If you select **SFTP**, set the following parameters:
     - **Source Path**: indicates the full HDFS path of a backup file. for example, *Backup path/Backup task name_Task creation time/Version_Data source_Task execution time*.**tar.gz**.
     - **Source Instance Name**: indicates the name of NameService corresponding to the backup directory when a restoration task is being executed. The default value is **hacluster**.

3. Click **OK**.

Step 5 Execute the restoration task.

In the restoration task list, locate the row where the created task resides, and click **Start** in the **Operation** column.

- After the restoration is successful, the progress bar is in green.
- After the restoration is successful, the restoration task cannot be executed again.
- If the restoration task fails during the first execution, rectify the fault and try to execute the task again by clicking **Start**.

Step 6 Determine what metadata has been restored.

- If the OMS and LdapServer metadata is restored, go to **Step 7**.
- If DBService data is restored, no further action is required.
- Restore NameNode data. On MRS Manager, choose **Services > HDFS > More > Restart Service**. The task is complete.

Step 7 Restarting Manager for the recovered data to take effect

1. In MRS Manager, Choose **LdapServer** > **More** > **Restart Service** and click **OK**. Wait until the LdapServer service is restarted successfully.
2. Log in to the active management node. For details, see **Determining Active and Standby Management Nodes**.
3. Run the following command to restart OMS:

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

   The command has been executed successfully if the following information is displayed:

   ```
   start HA successfully.
   ```

4. On MRS Manager, choose **KrbServer > More > Synchronize Configuration**. Do not select Restart the services and instances whose configuration has expired. Click **OK** and wait until the KrbServer service configuration is synchronized and restarted successfully.
5. Choose **Services > More > Synchronize Configuration**. Do not select Restart the services and instances whose configuration has expired. Click **OK** and wait until the cluster is configured and synchronized successfully.
6. Choose **Services > More > Stop Cluster**. After the cluster is stopped, choose **Services > More> Start Cluster**.

----End

9.11.4 Modifying a Backup Task

Scenario

This section describes how to modify the parameters of a created backup task on MRS Manager to meet changing service requirements. The parameters of restoration tasks can be viewed but not modified.

Impact on the System

After a backup task is modified, the new parameters take effect when the task is executed next time.

Prerequisites

- A backup task has been created.
- A new backup task policy has been planned based on the actual situation.

Procedure

Step 1 On MRS Manager, choose **System** > **Back Up Data**.

Step 2 In the task list, locate a specified task, click **Modify** in the **Operation** column to go to the configuration modification page.

Step 3 Modify the following parameters on the displayed page:

- Manual backup:
  - Target Path
  - Max Number of Backup Copies
- Periodic backup:
  - Started
  - Period
  - Target Path
  - Max Number of Backup Copies

  NOTE

  - When **Path Type** is set to **LocalHDFS**, **Target Path** is valid for modifying a backup task.
  - After you change the value of **Target Path** for a backup task, full backup is performed by default when the task is executed for the first time.

Step 4 Click **OK**.

----End

9.11.5 Viewing Backup and Restoration Tasks

Scenario

This section describes how to view created backup and restoration tasks and check their running status on MRS Manager.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 Click **Back Up Data** or **Restore Data**.

Step 3 In the task list, obtain the previous execution result in the **Task Progress** column. Green indicates that the task is executed successfully, and red indicates that the execution fails.

Step 4 In the **Operation** column of a specified task in the task list, choose **More** > **View History** to view the historical record of backup and restoration execution.

In the displayed window, click **View** in the **Details** column. The task execution logs and paths are displayed.

----End

Related Tasks

- Modifying a backup task

  For details, see **Modifying a Backup Task**.

- Viewing a restoration task

  In the **Operation** column of the specified task in the task list, click **View Details** to view the restoration task. You can only view but cannot modify the parameters of a restoration task.

- Executing a backup or restoration task

  In the task list, locate a specified task and click **Start** in the **Operation** column to start a backup or restoration task that is ready or fails to be executed. Executed restoration tasks cannot be repeatedly executed.

- Stopping backup tasks

  In the task list, locate a specified task and click **More** > **Stop** in the **Operation** column to stop a backup task that is running.

- Deleting a backup or restoration task

  In the **Operation** column of the specified task in the task list, choose **More** > **Delete** to delete the backup or restoration task. After a task is deleted, the backup data is retained by default.

- Suspending a backup task

  In the **Operation** column of the specified task in the task list, choose **More** > **Suspend** to suspend the backup task. Only periodic backup tasks can be suspended. Suspended backup tasks are no longer executed automatically. When you suspend a backup task that is being executed, the task execution stops. To cancel the suspension status of a task, click **More** > **Resume**.

9.12 Security Management

9.12.1 Default Users of Clusters with Kerberos Authentication Disabled

User Classification

The MRS cluster provides the following two types of users. Users are advised to periodically change the passwords. It is not recommended to use the default passwords.

User Type	Description
System users	User who runs OMS processes
Database users	• User who manages OMS database and accesses data • User who runs the database of service components (Hive, Loader, and DBService)

System users

NOTE

- User **Idap** of the OS is required in the MRS cluster. Do not delete this account. Otherwise, the cluster may not work properly. Password management policies are maintained by the operation users.
- Reset the passwords when you change the passwords of user **ommdba** and user **omm** for the first time. Change the passwords periodically after retrieving them.

Operation	Username	Initial Password	Description
System administrator of the MRS cluster	admin	Specified by the user during the cluster creation	MRS Manager This user has the following permissions: • Common HDFS and ZooKeeper user permissions. • Permissions to submit and query MapReduce and Yarn tasks, manage Yarn queues, and access the Yarn web UI. • Permissions to submit, query, activate, deactivate, reassign, delete topologies, and operate all topologies of the Storm service. • Permissions to create, delete, authorize, reassign, consume, write, and query topics of the Kafka service.

Operation	Username	Initial Password	Description
MRS cluster node OS user	omm	Randomly generated by the system	Internal running user of the MRS cluster system. This user is an OS user generated on all node and does not require a unified password.
MRS cluster node OS user	root	Set by the user	User for logging in to the node in the MRS cluster. This user is an OS user generated on all nodes.

User Group Information

Default User Group	Description
supergroup	Primary group of user **admin**, which has no additional permissions in the cluster with Kerberos authentication disabled.
check_sec_ldap	Used to test whether the active LDAP works properly. This user group is generated randomly in a test and automatically deleted after the test is complete. which is an internal system user group used only between components.
Manager_tenant	Tenant system user group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.
System_administrator	MRS cluster system administrator group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.
Manager_viewer	MRS Manager system viewer group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.
Manager_operator	MRS Manager system operator group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.
Manager_auditor	MRS Manager system auditor group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.

Default User Group	Description
Manager_administrator	MRS Manager system administrator group, which is an internal system user group used only between components. It is used only in clusters with Kerberos authentication enabled.
compcommon	MRS cluster internal group, used to access public resources in the cluster. All system users and system running users are added to this user group by default.
default_1000	User group created for tenants, which is an internal system user group used only between components.
launcher-job	MRS internal group, which is used to submit jobs using V2 APIs.

OS User Group	Description
wheel	Primary group of MRS internal running user **omm**.
ficommon	MRS cluster common group that corresponds to **compcommon** for accessing public resource files stored in the OS of the cluster.

Database users

MRS cluster system database users include OMS database users and DBService database users.

NOTE

Do not delete database users. Otherwise, the cluster or components may not work properly.

Operation	Default User	Initial Password	Description
OMS database	ommdba	dbChangeMe@123456	OMS database administrator who performs maintenance operations, such as creating, starting, and stopping.
	omm	ChangeMe@123456	User for accessing OMS database data
DBService database	omm	dbserverAdmin@123	Administrator of the GaussDB database in the DBService component

Operation	Default User	Initial Password	Description
	hive	HiveUser@	User for Hive to connect to the DBService database
	hue	HueUser@123	User for Hue to connect to the DBService database
	sqoop	SqoopUser@	User for Loader to connect to the DBService database.

9.12.2 Default Users of Clusters with Kerberos Authentication Enabled

User Classification

The MRS cluster provides the following three types of users. Users are advised to periodically change the passwords. It is not recommended to use the default passwords.

User Type	Description
System user	• User created on Manager for MRS cluster O&M and service scenarios. There are two types of users: – **Human-machine** user: used for Manager O&M scenarios and component client operation scenarios. – **Machine-machine** user: used for MRS cluster application development scenarios. • User who runs OMS processes.
Internal system user	Internal user who performs process communications, saves user group information, and associates user permissions.
Database user	• User who manages OMS database and accesses data. • User who runs the database of service components (Hive, Hue, Loader, and DBService)

System User

NOTE

- User **Idap** of the OS is required in the MRS cluster. Do not delete this account. Otherwise, the cluster may not work properly. Password management policies are maintained by the operation users.
- Reset the passwords when you change the passwords of user **ommdba** and user **omm** for the first time. Change the passwords periodically after retrieving them.

Type	Username	Initial Password	Description
System administrator of the MRS cluster	admin	Specified by the user during the cluster creation.	Manager administrator with the following permissions: • Common HDFS and ZooKeeper user permissions. • Permissions to submit and query MapReduce and Yarn tasks, manage Yarn queues, and access the Yarn web UI. • Permissions to submit, query, activate, deactivate, reassign, delete topologies, and operate all topologies of the Storm service. • Permissions to create, delete, authorize, reassign, consume, write, and query topics of the Kafka service.
MRS cluster node OS user	omm	Randomly generated by the system.	Internal running user of the MRS cluster system. This user is an OS user generated on all nodes and does not require a unified password.
MRS cluster node OS user	root	Set by the user.	User for logging in to the node in the MRS cluster. This user is an OS user generated on all nodes.

Internal System Users

NOTE

Do not delete the following internal system users. Otherwise, the cluster or components may not work properly.

Type	Default User	Initial Password	Description
Component running user	hdfs	Hdfs@123	This user is the HDFS system administrator and has the following permissions: 1. File system operation permissions: ● Views, modifies, and creates files. ● Views and creates directories. ● Views and modifies the groups where files belong. ● Views and sets disk quotas for users. 2. HDFS management operation permissions: ● Views the web UI status. ● Views and sets the active and standby HDFS status. ● Enters and exits the HDFS in security mode. ● Checks the HDFS file system.

Type	Default User	Initial Password	Description
	hbase	Hbase@123	This user is the HBase system administrator and has the following permissions: • Cluster management permission: **Enable** and **Disable** operations on tables to trigger MajorCompact and ACL operations. • Grants and revokes permissions, and shuts down the cluster. • Table management permission: Creates, modifies, and deletes tables. • Data management permission: Reads and writes data in tables, column families, and columns. • Accesses the HBase web UI.
	mapred	Mapred@123	This user is the MapReduce system administrator and has the following permissions: • Submits, stops, and views the MapReduce tasks. • Modifies the Yarn configuration parameters. • Accesses the Yarn and MapReduce web UI.
	spark	Spark@123	This user is the Spark system administrator and has the following permissions: • Accesses the Spark web UI. • Submits Spark tasks.

User Group Information

Default User Group	Description
hadoop	Users added to this user group have the permission to submit tasks to all Yarn queues.
hbase	Common user group. Users added to this user group will not have any additional permission.
hive	Users added to this user group can use Hive.
spark	Common user group. Users added to this user group will not have any additional permission.
supergroup	Users added to this user group can have the administrator permission of HBase, HDFS, and Yarn and can use Hive.
check_sec_ldap	Used to test whether the active LDAP works properly. This user group is generated randomly in a test and automatically deleted after the test is complete. This is an internal system user group used only between components.
Manager_tenant	Tenant system user group, which is an internal system user group used only between components.
System_administrator	MRS cluster system administrator group, which is an internal system user group used only between components.
Manager_viewer	MRS Manager system viewer group, which is an internal system user group used only between components.
Manager_operator	MRS Manager system operator group, which is an internal system user group used only between components.
Manager_auditor	MRS Manager system auditor group, which is an internal system user group used only between components.
Manager_administrator	MRS Manager system administrator group, which is an internal system user group used only between components.
compcommon	Internal system group for accessing public resources in a cluster. All system users and system running users are added to this user group by default.
default_1000	User group created for tenants, which is an internal system user group used only between components.

Default User Group	Description
kafka	Kafka common user group. Users added to this group need to be granted with read and write permission by users in the **kafkaadmin** group before accessing the desired topics.
kafkasuperuser	Users added to this group have permissions to read data from and write data to all topics.
kafkaadmin	Kafka administrator group. Users added to this group have the permissions to create, delete, authorize, as well as read from and write data to all topics.
storm	Storm common user group. Users added to this group have the permissions to submit topologies and manage their own topologies.
stormadmin	Storm administrator user group. Users added to this group have the permissions to submit topologies and manage their own topologies.
opentsdb	Common user group. Users added to this user group will not have any additional permission.
presto	Common user group. Users added to this user group will not have any additional permission.
flume	Common user group. Users added to this user group will not have any additional permission.
launcher-job	MRS internal group, which is used to submit jobs using V2 APIs.

OS User Group	Description
wheel	Primary group of MRS internal running user **omm**.
ficommon	MRS cluster common group that corresponds to **compcommon** for accessing public resource files stored in the OS of the cluster.

Database User

MRS cluster system database users include OMS database users and DBService database users.

NOTE

Do not delete database users. Otherwise, the cluster or components may not work properly.

Type	Default User	Initial Password	Description
OMS database	ommdba	dbChangeMe@123456	OMS database administrator who performs maintenance operations, such as creating, starting, and stopping applications.
	omm	ChangeMe@123456	User for accessing OMS database data.
DBService database	omm	dbserverAdmin@123	Administrator of the GaussDB database in the DBService component.
	hive	HiveUser@	User for Hive to connect to the DBService database.
	hue	HueUser@123	User for Hue to connect to the DBService database.
	sqoop	SqoopUser@	User for Loader to connect to the DBService database.
	ranger	RangerUser@	User for Ranger to connect to the DBService database.

9.12.3 Changing the Password of an OS User

Scenario

This section describes how to periodically change the login passwords of the OS users **omm**, **ommdba**, and **root** on MRS cluster nodes to improve the system O&M security.

Passwords of users **omm**, **ommdba**, and **root** on each node can be different.

Procedure

Step 1 Log in to the **Master1** node and then log in to other nodes whose OS user passwords need to be changed.

Step 2 Run the following command to switch to user **root**:

sudo su - root

Step 3 Run the following command to change the passwords of users **omm**, **ommdba**, or **root**:

passwd omm

passwd ommdba

passwd root

For example, if you run the **omm:passwd** command, the system displays the following information:

```
Changing password for user omm.
New password:
```

Enter a new password. The password change policies for an OS vary according to the OS that is used.

```
Retype new password:
passwd: all authentication tokens updated successfully.
```

NOTE

The default password complexity requirements of the MRS cluster are as follows:

- The password must contain at least eight characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$%^&*()-_=+\|[{}];:'",<.>/?).
- The new password cannot be the same as last five historical passwords.

----End

9.12.4 Changing the password of user admin

This section describes how to periodically change the password of cluster user **admin** to improve the system O&M security.

If the password is changed, the downloaded user credential will be unavailable. Download the authentication credential again, and replace the old one.

Changing the Password of User admin on the Cluster Node

Step 1 Update the client of the active management node. For details, see **Updating a Client (Versions Earlier Than 3.x)**.

Step 2 Log in to the active management node.

Step 3 (Optional) To change the password as user **omm**, run the following command to switch the user:

sudo su - omm

Step 4 Run the following command to switch to the client directory, for example, **/opt/client**.

cd /opt/client

Step 5 Run the following command to configure environment variables:

source bigdata_env

Step 6 Run the following command to change the password of user **admin**: This operation takes effect in the whole cluster.

kpasswd admin

Enter the old password and then enter a new password twice.

For the cluster, the default password complexity requirements are as follows:

- The password must contain at least eight characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

Changing the Password of User admin on MRS Manager

You can change the password of user **admin** on MRS Manager only for clusters with Kerberos authentication enabled and clusters with Kerberos authentication disabled but the EIP function enabled.

Step 1 Log in to MRS Manager as user **admin**.

Step 2 Click the username in the upper right corner of the page and choose **Change Password**.

Step 3 On the **Change Password** page, set **Old Password**, **New Password**, and **Confirm Password**.

Figure 9-11 Changing the password of user **admin**

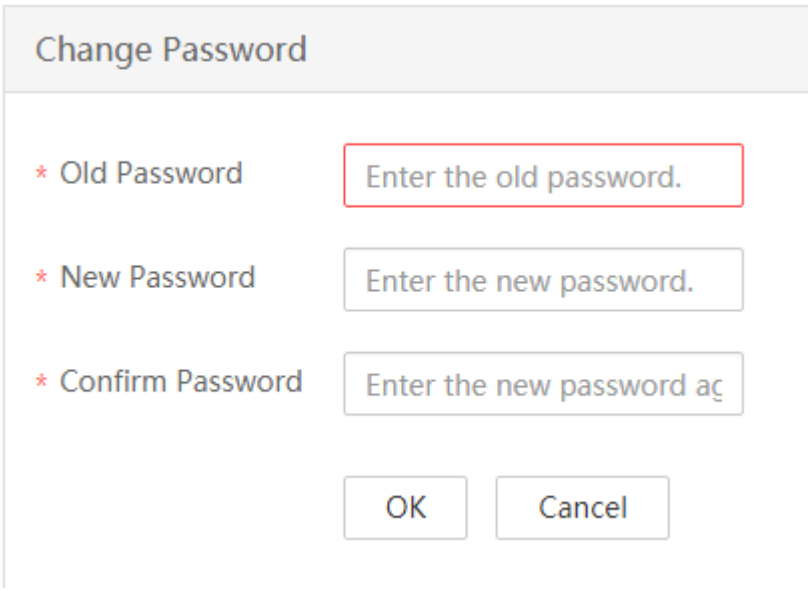

NOTE

The default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$%^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

Step 4 Click **OK**. Log in to MRS Manager with the new password.

----End

Resetting the Password for User admin

Step 1 Log in to the **Master1** node.

Step 2 (Optional) To change the password as user **omm**, run the following command to switch the user:

sudo su - omm

Step 3 Run the following command to switch to the client directory, for example, **/opt/client**:

cd /opt/client

Step 4 Run the following command to configure environment variables:

source bigdata_env

Step 5 Run the following command to log in to the console as user **kadmin/admin**:

kadmin -p kadmin/admin

NOTE

The default password of user **kadmin/admin** is **KAdmin@123**, which will expire upon your first login. Change the password as prompted. Keep the password secure because it cannot be retrieved once lost.

Step 6 Run the following command to reset the password of a component running user. This operation takes effect for all servers.

cpw *Component running user name*

For example, to reset the password of user admin, run the **cpw admin** command.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

9.12.5 Changing the Password of the Kerberos Administrator

Scenario

This section describes how to periodically change the password of the Kerberos administrator **kadmin** of the MRS cluster to improve the system O&M security.

If the password is changed, the downloaded user credential will be unavailable. Download the authentication credential again, and replace the old one.

Prerequisites

A client has been prepared on the **Master1** node.

Procedure

Step 1 Log in to the **Master1** node.

Step 2 (Optional) To change the password as user **omm**, run the following command to switch the user:

sudo su - omm

Step 3 Run the following command to switch to the client directory, for example, **/opt/client**.

cd /opt/client

Step 4 Run the following command to configure environment variables:

source bigdata_env

Step 5 Run the following command to change the password of **kadmin/admin**. This operation takes effect for all servers. Keep the password secure because it cannot be retrieved once lost.

kpasswd kadmin/admin

For the cluster, the default password complexity requirements are as follows:

- The password must contain at least eight characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

9.12.6 Changing the Passwords of the LDAP Administrator and the LDAP User

Scenario

This section describes how to periodically change the passwords of the LDAP administrator **rootdn:cn=root,dc=hadoop,dc=com** and the LDAP user **pg_search_dn:cn=pg_search_dn,ou=Users,dc=hadoop,dc=com** to improve the system O&M security.

Impact on the System

All services need to be restarted for the new password to take effect. The services are unavailable during the restart.

Procedure

Step 1 On MRS Manager, choose **Services** > **LdapServer** > **More**.

Step 2 Click **Change Password**.

Step 3 In the **Change Password** dialog box, select the user whose password needs to be modified in the **User Information** drop-down box.

Step 4 Enter the old password in the **Old Password** text box, and enter the new password in the **New Password** and **Confirm Password** text boxes.

The default password complexity requirements are as follows:

- The password contains 16 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters (`~!@#$%^&*()-_=+\|[{}];:",<.>/?).
- The password cannot be the username or the reverse username.
- The new password cannot be the same as the current password.

NOTE

The default password of the LDAP administrator **rootdn:cn=root,dc=hadoop,dc=com** is **LdapChangeMe@123**, and that of the LDAP user **pg_search_dn:cn=pg_search_dn,ou=Users,dc=hadoop,dc=com** is **pg_search_dn@123**. Periodically change the passwords and keep them secure.

Step 5 Select **I have read the information and understand the impact**, and click **OK** to confirm the modification and restart the service.

----End

9.12.7 Changing the Password of a Component Running User

Scenario

This section describes how to periodically change the password of the component running user of the MRS cluster to improve the system O&M security.

If the initial password is randomly generated by the system, reset the password.

If the password is changed, the downloaded user credential will be unavailable. Download the authentication credential again, and replace the old one.

Prerequisites

A client has been prepared on the **Master1** node.

Procedure

Step 1 Log in to the **Master1** node.

Step 2 (Optional) To change the password as user **omm**, run the following command to switch the user:

sudo su - omm

Step 3 Run the following command to switch to the client directory, for example, **/opt/client**:

cd /opt/client

Step 4 Run the following command to configure environment variables:

source bigdata_env

Step 5 Run the following command to log in to the console as user **kadmin/admin**:

kadmin -p kadmin/admin

NOTE

The default password of user **kadmin/admin** is **KAdmin@123**, which will expire upon your first login. Change the password as prompted. Keep the password secure because it cannot be retrieved once lost.

Step 6 Run the following command to reset the password of a component running user. This operation takes effect for all servers.

cpw *Component running user name*

For example, to reset the password of user admin, run the **cpw admin** command.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

9.12.8 Changing the Password of the OMS Database Administrator

Scenario

This section describes how to periodically change the password of the OMS database administrator to improve the system O&M security.

Procedure

Step 1 Log in to the active management node.

NOTE

The password of user **ommdba** cannot be changed on the standby management node. Otherwise, the cluster may not work properly. Change the password on the active management node only.

Step 2 Run the following command to switch the user:

sudo su - omm

Step 3 Run the following command to switch the directory:

cd $OMS_RUN_PATH/tools

Step 4 Run the following command to change the password of user **ommdba**:

mod_db_passwd ommdba

Step 5 Enter the old password of user **ommdba** and enter a new password twice.

The password complexity requirements are as follows:

- The password contains 16 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters ('~!@#$%^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.
- The password cannot be the same as the last 20 historical passwords.

If the following information is displayed, the password is changed successfully.

```
Congratulations, update [ommdba] password successfully.
```

----End

9.12.9 Changing the Password of the Data Access User of the OMS Database

Scenario

This section describes how to periodically change the password of the data access user of the OMS database to improve the system O&M security.

Impact on the System

The OMS service needs to be restarted for the new password to take effect. The service is unavailable during the restart.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Change OMS Database Password**.

Step 3 Locate the row that contains user **omm**, and click **Change password** in the **Operation** column.

The password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters ('~!@#$%^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.
- The password cannot be the same as the last 20 historical passwords.

Step 4 Click **OK**. When **Operation successful** is displayed, click **Finish**.

Step 5 Locate the row that contains user **omm**, and click **Restart the OMS service** in the **Operation** column to restart the OMS database.

NOTE

If the password is changed but the OMS database is not restarted, the status of user **omm** changes to **Waiting to restart** and the password cannot be changed until the OMS database is restarted.

Step 6 In the displayed dialog box, select **I have read the information and understand the impact**. Click **OK**, and restart the OMS service.

----End

9.12.10 Changing the Password of a Component Database User

Scenario

This section describes how to periodically change the password of the component database user to improve the system O&M security.

Impact on the System

The services need to be restarted for the new password to take effect. The services are unavailable during the restart.

Procedure

Step 1 On MRS Manager, click **Services** and click the name of the database user service to be modified.

Step 2 Determine the component database user whose password is to be changed.

- To change the password of the DBService database user, go to **Step 3**.
- To change the password of the Loader, Hive, or Hue database user, stop the service first and then execute **Step 3**.

Click **Stop Service**.

Step 3 Choose **More** > **Change Password**.

Step 4 Enter the old and new passwords as prompted.

The password complexity requirements are as follows:

- The password of the DBService database user contains 16 to 32 characters. The password of the Loader, Hive, or Hue database user contains 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters ('~!@#$%^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.
- The password cannot be the same as the last 20 historical passwords.

Step 5 Click **OK**. The system automatically restarts the corresponding service. When **Operation successful** is displayed, click **Finish**.

----End

9.12.11 Replacing the HA Certificate

Scenario

HA certificates are used to encrypt the communication between active/standby processes and HA processes to ensure the communication security. This section describes how to replace the HA certificates on the active and standby management nodes on MRS Manager to ensure the product security.

The certificate file and key file can be generated by the user.

Impact on the System

MRS Manager needs to be restarted during the replacement and cannot be accessed or provide services at that time.

Prerequisites

- You have obtained the **root-ca.crt** HA root certificate file and the **root-ca.pem** key file to be replaced.
- You have prepared a password, such as **Userpwd@123**, for accessing the key file.

  To avoid potential security risks, the password must meet the following complexity requirements:

  – The password must contain at least eight characters.

  – The password must contain at least four types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~`!?,.:;-_'(){}[]/<>@#$%^&*+|\=).

Procedure

Step 1 Log in to the active management node.

Step 2 Run the following commands to switch the user:

sudo su - root

su - omm

Step 3 Run the following commands to generate **root-ca.crt** and **root-ca.pem** in the **${OMS_RUN_PATH}/workspace0/ha/local/cert** directory on the active management node:

sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --root-ca --country=*country* **--state=***state* **--city=***city* **--company=***company* **--organize=***organize* **--common-name=***commonname* **--email=***Administrator email address* **--password=***password*

There can be security risks if a command contains the authentication password. You are advised to disable the command recording function (history) before running the command.

For example, run the following command: **sh ${OMS_RUN_PATH}/workspace/ha/module/hacom/script/gen-cert.sh --root-ca --country=CN --**

state=gd --city=sz --company=hw --organize=IT --common-name=HADOOP.COM --email=abc@hw.com --password=*xxx*

The command has been executed successfully if the following information is displayed:

```
Generate root-ca pair success.
```

Step 4 On the active management node, run the following command as user **omm** to copy **root-ca.crt** and **root-ca.pem** to the **${BIGDATA_HOME}/om-0.0.1/security/certHA** directory:

cp -arp ${OMS_RUN_PATH}/workspace0/ha/local/cert/root-ca.* ${BIGDATA_HOME}/om-0.0.1/security/certHA

Step 5 Copy **root-ca.crt** and **root-ca.pem** generated on the active management node to the **${BIGDATA_HOME}/om-0.0.1/security/certHA** directory on the standby management node as user **omm**.

Step 6 Run the following command to generate an HA certificate and perform the automatic replacement:

sh ${BIGDATA_HOME}/om-0.0.1/sbin/replacehaSSLCert.sh

Enter the password as prompted, and press **Enter**.

```
Please input ha ssl cert password:
```

The HA certificate is replaced successfully if the following information is displayed:

```
[INFO] Succeed to replace ha ssl cert.
```

Step 7 Run the following command to restart OMS:

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

The following information is displayed:

```
start HA successfully.
```

Step 8 Log in to the standby management node and switch to user **omm**. Repeat step **Step 6** to step **Step 7**.

Run the **sh ${BIGDATA_HOME}/om-0.0.1/sbin/status-oms.sh** command to check whether **HAAllResOK** of the management node is **Normal**. Access MRS Manager again. If MRS Manager can be accessed, the operation is successful.

----End

9.12.12 Updating Cluster Keys

Scenario

When a cluster is installed, an encryption key is generated automatically to store the security information in the cluster (such as all database user passwords and key file access passwords) in encryption mode. After the cluster is successfully installed, you are advised to periodically update the encryption key based on the following procedure.

Impact on the System

- After a cluster key is updated, a new key is generated randomly in the cluster. This key is used to encrypt and decrypt the newly stored data. The old key is not deleted, and it is used to decrypt data encrypted using the old key. After security information is modified, for example, a database user password is changed, the new password is encrypted using the new key.
- When the key is updated, the cluster is stopped and cannot be accessed.

Prerequisites

The upper-layer applications depending on the cluster are stopped.

Procedure

Step 1 Log in to MRS Manager and choose **Services** > **More** > **Stop Cluster**.

In the displayed dialog box, select **I have read the information and understand the impact.** Click **OK**. Wait until the system displays a message indicating that the operation is successful. Click **Finish**. The cluster is stopped successfully.

Step 2 Log in to the active management node.

Step 3 Run the following commands to switch the user:

sudo su - omm

Step 4 Run the following command to disable logout upon timeout:

TMOUT=0

Step 5 Run the following command to switch the directory:

cd ${BIGDATA_HOME}/om-0.0.1/tools

Step 6 Run the following command to update the cluster key:

sh updateRootKey.sh

Enter **y** as prompted.

```
The root key update is a critical operation.
Do you want to continue?(y/n):
```

The key is updated successfully if the following information is displayed:

```
...
Step 4-1: The key save path is obtained successfully.
...
Step 4-4: The root key is sent successfully.
```

Step 7 On MRS Manager, choose **Services > More > Start Cluster**.

In the displayed dialog box, click **OK**. After **Operation successful** is displayed, click **Finish**. The cluster is started.

----End

9.13 Permissions Management

9.13.1 Creating a Role

Scenario

This section describes how to create a role on MRS Manager and authorize and manage Manager and components.

Up to 1,000 roles can be created on MRS Manager.

Prerequisites

You have learned service requirements.

Procedure

Step 1 On MRS Manager, choose **System** > **Manage Role**.

Step 2 Click **Create Role** and fill in **Role Name** and **Description**.

Role Name is mandatory and contains 3 to 30 digits, letters, and underscores (_). **Description** is optional.

Step 3 In **Permission**, set role permission.

1. Click **Service Name** and select a name in **View Name**.
2. Select one or more permissions.

NOTE

- The **Permission** parameter is optional.
- If you select **View Name** to set component permissions, you can enter a resource name in the **Search** box in the upper right corner and click . The search result is displayed.
- The search scope covers only directories with current permissions. You cannot search subdirectories. Search by keywords supports fuzzy match and is case-insensitive. Results of the next page can be searched.

Table 9-36 Manager permission description

Resource Supporting Permission Management	Permission Setting
Alarm	Authorizes the Manager alarm function. You can select **View** to view alarms and **Management** to manage alarms.
Audit	Authorizes the Manager audit log function. You can select **View** to view audit logs and **Management** to manage audit logs.
Dashboard	Authorizes the Manager overview function. You can select **View** to view the cluster overview.
Hosts	Authorizes the node management function. You can select **View** to view node information and **Management** to manage nodes.

Resource Supporting Permission Management	Permission Setting
Services	Authorizes the service management function. You can select **View** to view service information and **Management** to manage services.
System_cluster_management	Authorizes the MRS cluster management function. You can select **Management** to use the MRS patch management function.
System_configuration	Authorizes the MRS cluster configuration function. You can select **Management** to configure MRS clusters on Manager.
System_task	Authorizes the MRS cluster task function. You can select **Management** to manage periodic tasks of MRS clusters on Manager.
Tenant	Authorizes the Manager multi-tenant management function. You can select **Management** to manage multi-tenants.

Table 9-37 HBase permission description

Resource Supporting Permission Management	Permission Setting
SUPER_USER_GROUP	Grants you HBase administrator rights.
Global	HBase resource type, indicating the whole HBase.
Namespace	HBase resource type, indicating namespace, which is used to store HBase tables. It has the following permissions: • **Admin** permission to manage the namespace • **Create**: permission to create HBase tables in the namespace • **Read**: permission to access the namespace • **Write**: permission to write data to the namespace • **Execute**: permission to execute the coprocessor (Endpoint)

Resource Supporting Permission Management	Permission Setting
Table	HBase resource type, indicating a data table, which is used to store data. It has the following permissions: • **Admin**: permission to manage a data table • **Create**: permission to create column families and columns in a data table • **Read**: permission to read a data table • **Write**: permission to write data to a data table • **Execute**: permission to execute the coprocessor (Endpoint)
ColumnFamily	HBase resource type, indicating a column family, which is used to store data. It has the following permissions: • **Create**: permission to create columns in a column family • **Read**: permission to read a column family • **Write**: permission to write data to a column family
Qualifier	HBase resource type, indicating a column, which is used to store data. It has the following permissions: • **Read**: permission to read a column • **Write**: permission to write data to a column

By default, permissions of an HBase resource type of each level are shared by resource types of sub-levels. However, the **Recursive** option is not selected by default. For example, if **Read** and **Write** permissions are added to the **default** namespace, they are automatically added to the tables, column families, and columns in the namespace. If a child resource is set after the parent resource, the permission of the child resource is the union of the permissions of the parent resource and the current child resource.

Table 9-38 HDFS permission description

Resource Supporting Permission Management	Permission Setting
Folder	HDFS resource type, indicating an HDFS directory, which is used to store files or subdirectories. It has the following permissions: • **Read**: permission to access the HDFS directory • **Write**: permission to write data to the HDFS directory • **Execute**: permission to perform an operation. It must be selected when you add access or write permission.
Files	HDFS resource type, indicating a file in HDFS. It has the following permissions: • **Read**: permission to access the file • **Write**: permission to write data to the file • **Execute**: permission to perform an operation. It must be selected when you add access or write permission.

Permissions of an HDFS directory of each level are not shared by directory types of sub-levels by default. For example, if **Read** and **Execute** permissions are added to the **tmp** directory, you must select **Recursive** at the same time to add permissions to subdirectories.

Table 9-39 Hive permission description

Resource Supporting Permission Management	Permission Setting
Hive Admin Privilege	Grants you Hive administrator rights.
Database	Hive resource type, indicating a Hive database, which is used to store Hive tables. It has the following permissions: • **Select**: permission to query the Hive database • **Delete**: permission to perform the deletion operation in the Hive database • **Insert**: permission to perform the insertion operation in the Hive database • **Create**: permission to perform the creation operation in the Hive database

Resource Supporting Permission Management	Permission Setting
Table	Hive resource type, indicating a Hive table, which is used to store data. It has the following permissions: • **Select**: permission to query the Hive table • **Delete**: permission to perform the deletion operation in the Hive table • **Update**: grants users the **Update** permission of the Hive table • **Insert**: permission to perform the insertion operation in the Hive table • **Grant of Select**: permission to grant the **Select** permission to other users using Hive statements • **Grant of Delete**: permission to grant the **Delete** permission to other users using Hive statements • **Grant of Update**: permission to grant the **Update** permission to other users using Hive statements • **Grant of Insert**: permission to grant the **Insert** permission to other users using Hive statements

By default, permissions of a Hive resource type of each level are shared by resource types of sub-levels. However, the **Recursive** option is not selected by default. For example, if **Select** and **Insert** permissions are added to the **default** database, they are automatically added to the tables and columns in the database. If a child resource is set after the parent resource, the permission of the child resource is the union of the permissions of the parent resource and the current child resource.

Table 9-40 Yarn permission description

Resource Supporting Permission Management	Permission Setting
Cluster Admin Operations	Grants you Yarn administrator rights.
root	Root queue of Yarn. It has the following permissions: • **Submit**: permission to submit jobs in the queue • **Admin**: permission to manage permissions of the current queue

Resource Supporting Permission Management	Permission Setting
Parent Queue	Yarn resource type, indicating a parent queue containing sub-queues. A root queue is a type of a parent queue. It has the following permissions: • **Submit**: permission to submit jobs in the queue • **Admin**: permission to manage permissions of the current queue
Leaf Queue	Yarn resource type, indicating a leaf queue. It has the following permissions: • **Submit**: permission to submit jobs in the queue • **Admin**: permission to manage permissions of the current queue

By default, permissions of a Yarn resource type of each level are shared by resource types of sub-levels. However, the **Recursive** option is not selected by default. For example, if the **Submit** permission is added to the **root** queue, it is automatically added to the sub-queue. Permissions inherited by sub-queues will not be displayed as selected in the **Permission** table. If a child resource is set after the parent resource, the permission of the child resource is the union of the permissions of the parent resource and the current child resource.

Table 9-41 Hue permission description

Resource Supporting Permission Management	Permission Setting
Storage Policy Admin	Grants you storage policy administrator rights.

Step 4 Click **OK**. Return to **Manage Role**.

----End

Related Tasks

Modifying a role

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage Role**.

Step 3 In the row of the role to be modified, click **Modify** to modify role information.

NOTE

If you change permissions assigned by the role, it takes 3 minutes to make new configurations take effect.

Step 4 Click **OK**. The modification is complete.

----End

Deleting a role

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage Role**.

Step 3 In the row of the role to be deleted, click **Delete**.

Step 4 Click **OK**. The role is deleted.

----End

9.13.2 Creating a User Group

Scenario

This section describes how to create user groups and specify their operation permissions on MRS Manager. Management of single or multiple users can be unified in the user groups. After being added to a user group, users can obtain operation permissions owned by the user group.

Up to 100 user groups can be created on MRS Manager.

Prerequisites

Administrators have learned service requirements and created roles required by service scenarios.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User Group**.

Step 3 Above the user group list, click **Create User Group**.

Step 4 Input **Group Name** and **Description**.

Group Name is mandatory and contains 3 to 20 digits, letters, and underscores (_). **Description** is optional.

Step 5 In **Role**, click **Select and Add Role** to select and add specified roles.

If you do not add the roles, the user group you are creating now does not have the permission to use MRS clusters.

Step 6 Click **OK**. The user group is created.

----End

Related Tasks

Modifying a user group

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User Group**.

Step 3 In the row of the user group to be modified, click **Modify**.

NOTE

If you change role permissions assigned to the user group, it takes 3 minutes to make new configurations take effect.

Step 4 Click **OK**. The modification is complete.

----End

Deleting a user group

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User Group**.

Step 3 In the row of the user group to be deleted, click **Delete**.

Step 4 Click **OK**. The user group is deleted.

----End

9.13.3 Creating a User

Scenario

This section describes how to create users on MRS Manager based on site requirements and specify their operation permissions to meet service requirements.

Up to 1,000 users can be created on MRS Manager.

If a new password policy needs to be used for a new user's password, follow instructions in **Modifying a Password Policy** to modify the password policy and then perform the following operations to create a user.

Prerequisites

Administrators have learned service requirements and created roles and role groups required by service scenarios.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User**.

Step 3 Above the user list, click **Create User**.

Step 4 Configure parameters as prompted and enter a username in **User Name**.

NOTE

- If a username exists, you cannot create another username that only differs from the existing username in case. For example, if **User1** has been created, you cannot create **user1**.
- When you use the user you created, enter the correct username, which is case-sensitive.
- **User Name** is mandatory and contains 3 to 20 digits, letters, and underscores (_).
- **root**, **omm**, and **ommdba** are reserved system user. Select another username.

Step 5 Set **User Type** to either **Human-Machine** or **Machine-Machine**.

- **Human-Machine** users: used for O&M on MRS Manager and operations on component clients. If you select this user type, you need to enter a password and confirm the password in **Password** and **Confirm Password** accordingly.
- **Machine-Machine** users: used for MRS application development. If you select this user type, you do not need to enter a password, because the password is randomly generated.

Step 6 In **User Group**, click **Select and Join User Group** to select user groups and add users to them.

NOTE

- If roles have been added to user groups, the users can be granted with permissions of the roles.
- If you want to grant new users with Hive permissions, add the users to the Hive group.
- If a user needs to manage tenant resources, the user group must be assigned the **Manager_tenant** role and the role corresponding to the tenant.

Step 7 In **Primary Group**, select a group as the primary group for users to create directories and files. The drop-down list contains all groups selected in **User Group**.

Step 8 In **Assign Rights by Role**, click **Select and Add Role** to add roles for users based on service requirements.

NOTE

- When you create a user, if permissions of a user group that is granted to the user cannot meet service requirements, you can assign other created roles to the user. It takes 3 minutes to make role permissions granted to the new user take effect.
- Adding a role when you create a user can specify the user rights.
- A new user can access WebUIs of HDFS, HBase, Yarn, Spark, and Hue even when roles are not assigned to the user.

Step 9 In **Description**, provide description based on onsite service requirements.

Description is optional.

Step 10 Click **OK**. The user is created.

If a new user is used in the MRS cluster for the first time, for example, used for logging in to MRS Manager or using the cluster client, the password must be changed. For details, see section **Changing the Password of an Operation User**.

----End

9.13.4 Modifying User Information

Scenario

This section describes how to modify user information on MRS Manager, including information about the user group, primary group, role, and description.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User**.

Step 3 In the row of the user to be modified, click **Modify**.

NOTE

If you change user groups for or assign role permissions to the user, it takes 3 minutes to make new configurations take effect.

Step 4 Click **OK**. The modification is complete.

----End

9.13.5 Locking a User

This section describes how to lock users in MRS clusters. A locked user cannot log in to MRS Manager or perform security authentication in the cluster.

A locked user can be unlocked by an administrator manually or until the lock duration expires. You can lock a user by using either of the following methods:

- Automatic lock: Set **Number of Password Retries** in **Configure Password Policy**. If user login attempts exceed the parameter value, the user is automatically locked. For details, see **Modifying a Password Policy** .
- Manual lock: The administrator manually locks a user.

The following describes how to manually lock a user. **Machine-Machine** users cannot be locked.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User**.

Figure 9-12 User management

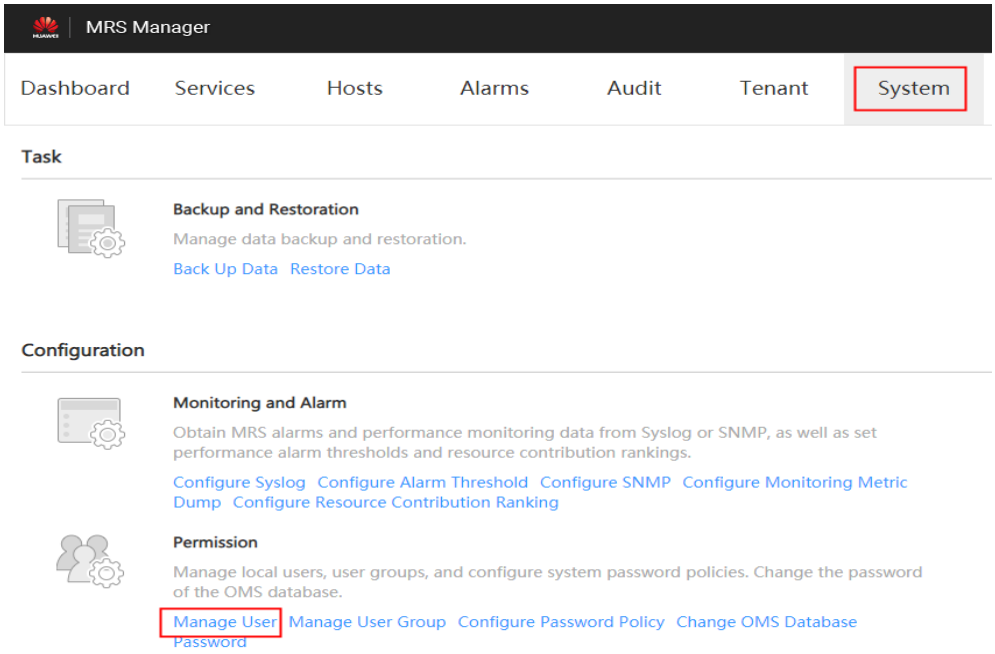

Step 3 In the row of the user to be locked, click **Lock User**.

Figure 9-13 Locking a user

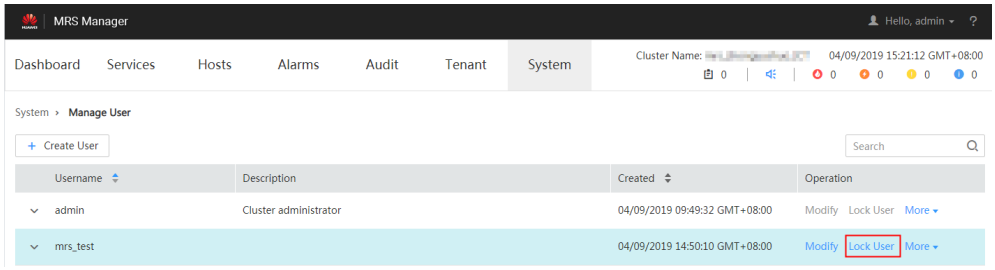

Step 4 In the window that is displayed, click **Yes** to lock the user.

----End

9.13.6 Unlocking a User

If a user is locked because the number of login attempts exceeds the value of **Number of Password Retries**, or the user is manually locked by the administrator, the administrator can unlock the user on MRS Manager.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User**.

Figure 9-14 User management

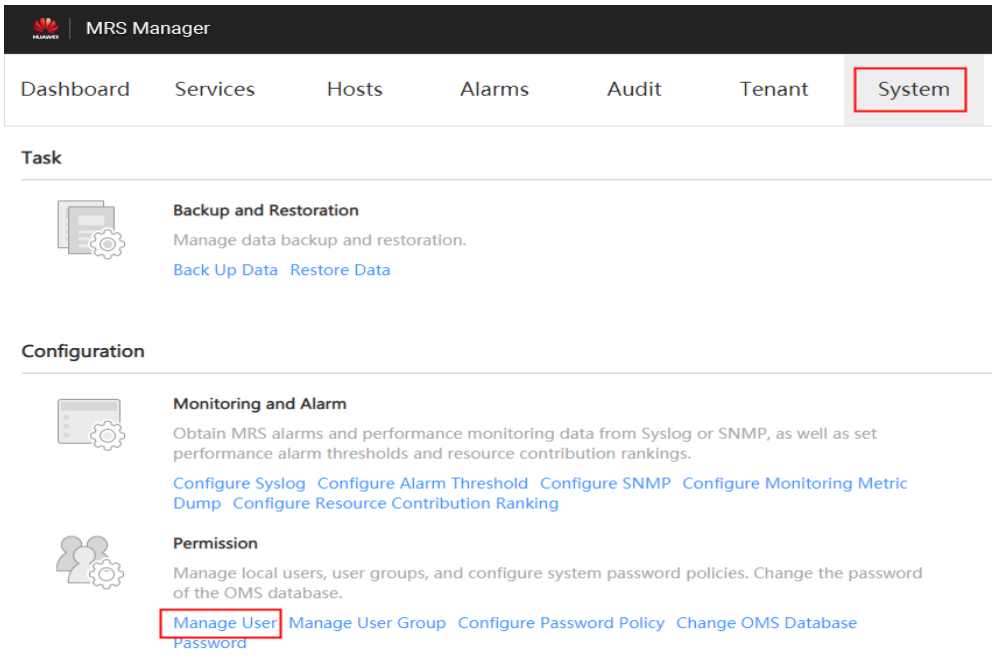

Step 3 In the row of the user to be unlocked, click **Unlock User**.

Figure 9-15 Unlocking a user

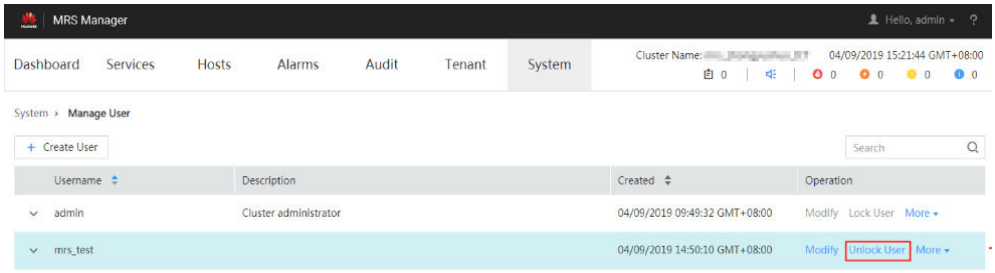

Step 4 In the window that is displayed, click **Yes** to unlock the user.

----End

9.13.7 Deleting a User

The administrator can delete an MRS cluster user that is not required on MRS Manager.

NOTE

If you want to create a new user with the same name as user A after deleting user A who has submitted a job on the client or MRS console, you need to delete user A's residual folders when deleting user A. Otherwise, the newly created user A may fail to submit a job.

To delete residual folders, log in to each Core node in the MRS cluster and run the following commands. In the following commands, **$user** indicates the folder named after the username.

cd /srv/BigData/hadoop/data1/nm/localdir/usercache/

rm -rf $user

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User**.

Figure 9-16 User management

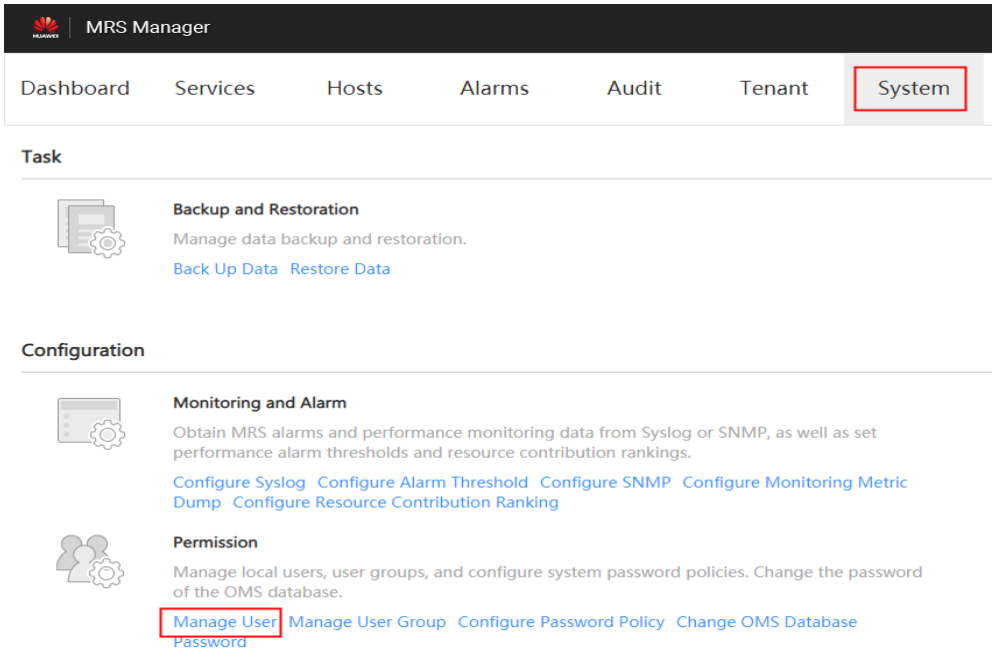

Step 3 In the row that contains the user to be deleted, choose **More** > **Delete**.

Figure 9-17 Deleting a User

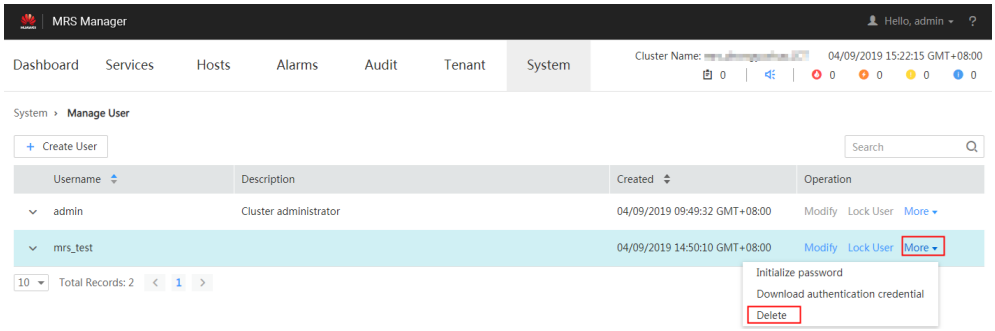

Step 4 Click **OK**.

----End

9.13.8 Changing the Password of an Operation User

Scenario

Passwords of **Human-Machine** system users must be regularly changed to ensure MRS cluster security. This section describes how to change your passwords on MRS Manager.

If a new password policy needs to be used for the password modified by the user, follow instructions in **Modifying a Password Policy** to modify the password policy and then perform the following operations to modify the password.

Impact on the System

If you have downloaded a user authentication file, download it again and obtain the keytab file after changing the password of the MRS cluster user.

Prerequisites

- You have obtained the current password policy.
- You have obtained the URL for accessing MRS Manager.

Procedure

Step 1 On MRS Manager, move the mouse cursor to  in the upper right corner.

On the menu that is displayed, select **Change Password**.

Step 2 Fill in the **Old Password**, **New Password**, and **Confirm Password**. Click **OK**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

9.13.9 Initializing the Password of a System User

Scenario

This section describes how to initialize a password on MRS Manager if a user forgets the password or the password of a public account needs to be changed regularly. After password initialization, the user must change the password upon the first login.

Impact on the System

If you have downloaded a user authentication file, download it again and obtain the keytab file after initializing the password of the MRS cluster user.

Initializing the Password of a Human-Machine User

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User**.

Step 3 Locate the row that contains the user whose password is to be initialized, choose **More** > **Initialize password**, and change the password as prompted.

In the window that is displayed, enter the password of the current administrator account and click **OK**. Then in **Initialize password**, click **OK**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

Initializing the Password of a Machine-Machine User

Step 1 Prepare a client based on service conditions and log in to the node where the client is installed.

Step 2 Run the following command to switch the user:

sudo su - omm

Step 3 Run the following command to switch to the client directory, for example, **/opt/client**:

cd /opt/client

Step 4 Run the following command to configure environment variables:

source bigdata_env

Step 5 Run the following command to log in to the console as user **kadmin/admin**:

kadmin -p kadmin/admin

NOTE

The default password of user **kadmin/admin** is **KAdmin@123**, which will expire upon your first login. Change the password as prompted and keep the new password secure.

Step 6 Run the following command to reset the password of a component running user. This operation takes effect for all servers.

cpw *Component running user name*

For example, **cpw oms/manager**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

9.13.10 Downloading a User Authentication File

Scenario

When a user develops big data applications and runs them in an MRS cluster that supports Kerberos authentication, the user needs to prepare a user authentication file for accessing the MRS cluster. The keytab file in the authentication file can be used for user authentication.

This section describes how to download a user authentication file and export the keytab file on MRS Manager.

NOTE

- Before downloading a **Human-machine** user authentication file, change the password for the user on MRS Manager to make the initial password set by the administrator invalid. Otherwise, the exported keytab file cannot be used. For details, see **Changing the Password of an Operation User**.
- After a user password is changed, the exported keytab file becomes invalid, and you need to export a keytab file again.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User**.

Step 3 In the row of the user for whom you want to export the keytab file, choose **More** > **Download authentication credential** to download the authentication file. After the file is automatically generated, save it to a specified path and keep it properly.

Step 4 Open the authentication file with a decompression program.

- **user.keytab** indicates a user keytab file used for user authentication.
- **krb5.conf** indicates the configuration file of the authentication server. The application connects to the authentication server according to the configuration file information when authenticating users.

----End

9.13.11 Modifying a Password Policy

Scenario

This section describes how to set password and user login security rules as well as user lock rules. Password policies set on MRS Manager take effect for **Human-machine** users only, because the passwords of **Machine-machine** users are randomly generated.

If a new password policy needs to be used for a new user's password or the password modified by the user, perform the following operations to modify the password policy first, and then create a user or change the password by following instructions in **Creating a User** or **Changing the Password of an Operation User**.

NOTICE

- Because password policies are critical to user management security, modify them based on service security requirements. Otherwise, security risks may be incurred.
- New password policies take effect only after both the password policies and user password are changed.

Procedure

Step 1 On MRS Manager, click **System**.

Step 2 Click **Configure Password Policy**.

Step 3 Modify password policies as prompted. For parameter details, see the following table:

Table 9-42 Password policy parameter description

Parameter	Description
Minimum Password Length	Indicates the minimum number of characters a password contains. The value ranges from 8 to 32. The default value is **8**.
Number of Character Types	Indicates the minimum number of character types a password contains. The character types are uppercase letters, lowercase letters, digits, spaces, and special characters (~`!?,.:;-_'(){}[]/<>@#$ %^&*+|\=). The value can be **3** or **4**. The default value **3** indicates that the password must contain at least three types of the following characters: uppercase letters, lowercase letters, digits, special characters, and spaces.
Password Validity Period (days)	Indicates the validity period (days) of a password. The value ranges from 0 to 90. 0 means that the password is permanently valid. The default value is **90**.

Parameter	Description
Password Expiration Notification Days	Indicates the number of days in advance users are notified that their passwords are about to expire. After the value is set, if the difference between the cluster time and the password expiration time is smaller than this value, the user receives password expiration notifications. When a user logs in to MRS Manager, a message is displayed, indicating that the password is about to expire and asking the user whether to change the password. The value ranges from **0** to *X* (*X* must be set to the half of the password validity period and rounded down). Value **0** indicates that no notification is sent. The default value is **5**.
Interval of Resetting Authentication Failure Count (min)	Indicates the interval of retaining incorrect password attempts, in minutes. The value ranges from 0 to 1440. 0 indicates that incorrect password attempts are permanently retained and 1440 indicates that incorrect password attempts are retained for one day. The default value is **5**.
Number of Password Retries	Indicates the number of consecutive wrong passwords allowed before the system locks the user. The value ranges from 3 to 30. The default value is **5**.
Account Lock Duration (min)	Indicates the time period for which a user is locked when the user lockout conditions are met. The value ranges from 5 to 120. The default value is **5**.

----End

9.14 MRS Multi-User Permission Management

9.14.1 Users and Permissions of MRS Clusters

Overview

- **MRS Cluster Users**

  Indicate the security accounts of Manager, including usernames and passwords. These accounts are used to access resources in MRS clusters. Each MRS cluster in which Kerberos authentication is enabled can have multiple users.

- **MRS Cluster Roles**

  Before using resources in an MRS cluster, users must obtain the access permission which is defined by MRS cluster objects. A cluster role is a set of one or more permissions. For example, the permission to access a directory in HDFS needs to be configured in the specified directory and saved in a role.

Manager provides the user permission management function for MRS clusters, facilitating permission and user management.

- Permission management: adopts the role-based access control (RBAC) mode. In this mode, permissions are granted by role to form a permission set. After one or more roles are allocated to a user, the user can obtain the permissions of the roles.
- User management: uses MRS Manager to uniformly manage users, adopts the Kerberos protocol for user identity verification, and employs Lightweight Directory Access Protocol (LDAP) to store user information.

Permission Management

Permissions provided by MRS clusters include the O&M permissions of Manager and components (such as HDFS, HBase, Hive, and Yarn). In actual application, permissions must be assigned to each user based on service scenarios. To facilitate permission management, Manager introduces the role function to allow administrators to select and assign specified permissions. Permissions are centrally viewed and managed in permission sets, enhancing user experience.

A role is a logical entity that contains one or more permissions. Permissions are assigned to roles, and users can be granted the permissions by obtaining the roles.

A role can have multiple permissions, and a user can be bound to multiple roles.

- Role 1: is assigned operation permissions A and B. After role 1 is allocated to users a and b, users a and b can obtain operation permissions A and B.
- Role 2: is assigned operation permission C. After role 2 is allocated to users c and d, users c and d can obtain operation permission C.
- Role 3: is assigned operation permissions D and F. After role 3 is allocated to user a, user a can obtain operation permissions D and F.

For example, if an MRS user is bound to the administrator role, the user becomes an administrator of the MRS cluster.

Table 9-43 lists the roles that are created by default on Manager.

Table 9-43 Default roles and description

Default Role	Description
default	Tenant role
Manager_administrator	Manager administrator: This role has the permission to manage MRS Manager.
Manager_auditor	Manager auditor: This role has the permission to view and manage auditing information.

Default Role	Description
Manager_operator	Manager operator: This role has all permissions except tenant, configuration, and cluster management permissions.
Manager_viewer	Manager viewer: This role has the permission to view the information about systems, services, hosts, alarms, and auditing logs.
System_administrator	System administrator: This role has the permissions of Manager administrators and all service administrators.
Manager_tenant	Manager tenant viewer: This role has the permission to view information on the **Tenant** page on MRS Manager.

When creating a role on Manager, you can perform rights management for Manager and components, as shown in **Table 9-44**.

Table 9-44 Manager and component permission management

Permission	Description
Manager	Manager access and login permission.
HBase	HBase administrator permission and permission for accessing HBase tables and column families.
HDFS	HDFS directory and file permission.
Hive	• Hive Admin Privilege Hive administrator permission. • Hive Read Write Privileges Hive data table management permission to set and manage the data of created tables.
Hue	Storage policy administrator permissions.
Yarn	• Cluster Admin Operations Yarn administrator permission. • Scheduler Queue Queue resource management permission.

User Management

MRS clusters that support Kerberos authentication use the Kerberos protocol and LDAP for user management.

- Kerberos verifies the identity of the user when a user logs in to Manager or uses a component client. Identity verification is not required for clusters with Kerberos authentication disabled.

- LDAP is used to store user information, including user records, user group information, and permission information.

MRS clusters can automatically update Kerberos and LDAP user data when users are created or modified on Manager. They can also automatically perform user identity verification and authentication and obtain user information when a user logs in to Manager or uses a component client. This ensures the security of user management and simplifies the user management tasks. Manager also provides the user group function for managing one or multiple users by type:

- A user group is a set of users, which can be used to manage users by type. Users in the system can exist independently or in a user group.
- After a user is added to a user group to which roles are allocated, the role permission of the user group is assigned to the user.

Table 9-45 lists the user groups that are created by default on MRS Manager in MRS 3.x or earlier.

For details about the default user groups displayed on FusionInsight Manager of MRS 3.*x* or later, see **User group**.

Table 9-45 Default user groups and description

User Group	Description
hadoop	Users added to this user group have the permission to submit tasks to all Yarn queues.
hbase	Common user group. Users added to this user group will not have any additional permission.
hive	Users added to this user group can use Hive.
spark	Common user group. Users added to this user group will not have any additional permission.
supergroup	Users added to this user group can have the administrator permission of HBase, HDFS, and Yarn and can use Hive.
flume	Common user group. Users added to this user group will not have any additional permission.
kafka	Kafka common user group. Users added to this group need to be granted with read and write permission by users in the **kafkaadmin** group before accessing the desired topics.
kafkasuperuser	Users added to this group have permissions to read data from and write data to all topics.
kafkaadmin	Kafka administrator group. Users added to this group have the permissions to create, delete, authorize, as well as read from and write data to all topics.

User Group	Description
storm	Storm common user group. Users added to this group have the permissions to submit topologies and manage their own topologies.
stormadmin	Storm administrator user group. Users added to this group have the permissions to submit topologies and manage their own topologies.

User **admin** is created by default for MRS clusters with Kerberos authentication enabled and is used for administrators to maintain the clusters.

Process Overview

In practice, MRS cluster users must understand the service scenarios of big data and plan user permissions. Then, create roles and assign permissions to the roles on MRS Manager to meet service requirements. Manager provides the user group function for administrators to create user groups for managing users of one or multiple service scenarios of the same type.

NOTE

If a role has the permission of HDFS, HBase, Hive, or Yarn respectively, the role can only use the corresponding functions of the component. To use Manager, the corresponding Manager permission must be added to the role.

Figure 9-18 Process of creating a user

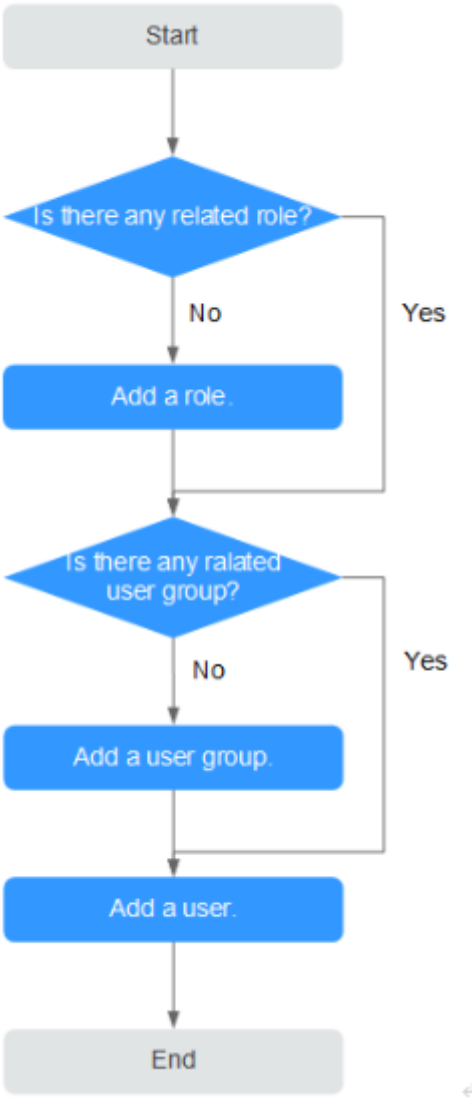

9.14.2 Default Users of Clusters with Kerberos Authentication Enabled

User Classification

The MRS cluster provides the following three types of users. Users are advised to periodically change the passwords. It is not recommended to use the default passwords.

User Type	Description
System user	• User created on Manager for MRS cluster O&M and service scenarios. There are two types of users: – **Human-machine** user: used for Manager O&M scenarios and component client operation scenarios. – **Machine-machine** user: used for MRS cluster application development scenarios. • User who runs OMS processes.
Internal system user	Internal user who performs process communications, saves user group information, and associates user permissions.
Database user	• User who manages OMS database and accesses data. • User who runs the database of service components (Hive, Hue, Loader, and DBService)

System User

NOTE

- User **Idap** of the OS is required in the MRS cluster. Do not delete this account. Otherwise, the cluster may not work properly. Password management policies are maintained by the operation users.
- Reset the passwords when you change the passwords of user **ommdba** and user **omm** for the first time. Change the passwords periodically after retrieving them.

Type	Username	Initial Password	Description
System administrator of the MRS cluster	admin	Specified by the user during the cluster creation.	Manager administrator with the following permissions: • Common HDFS and ZooKeeper user permissions. • Permissions to submit and query MapReduce and Yarn tasks, manage Yarn queues, and access the Yarn web UI. • Permissions to submit, query, activate, deactivate, reassign, delete topologies, and operate all topologies of the Storm service. • Permissions to create, delete, authorize, reassign, consume, write, and query topics of the Kafka service.
MRS cluster node OS user	omm	Randomly generated by the system.	Internal running user of the MRS cluster system. This user is an OS user generated on all nodes and does not require a unified password.
MRS cluster node OS user	root	Set by the user.	User for logging in to the node in the MRS cluster. This user is an OS user generated on all nodes.

Internal System Users

NOTE

Do not delete the following internal system users. Otherwise, the cluster or components may not work properly.

Type	Default User	Initial Password	Description
Component running user	hdfs	Hdfs@123	This user is the HDFS system administrator and has the following permissions: 1. File system operation permissions: ● Views, modifies, and creates files. ● Views and creates directories. ● Views and modifies the groups where files belong. ● Views and sets disk quotas for users. 2. HDFS management operation permissions: ● Views the web UI status. ● Views and sets the active and standby HDFS status. ● Enters and exits the HDFS in security mode. ● Checks the HDFS file system.

Type	Default User	Initial Password	Description
	hbase	Hbase@123	This user is the HBase system administrator and has the following permissions: • Cluster management permission: **Enable** and **Disable** operations on tables to trigger MajorCompact and ACL operations. • Grants and revokes permissions, and shuts down the cluster. • Table management permission: Creates, modifies, and deletes tables. • Data management permission: Reads and writes data in tables, column families, and columns. • Accesses the HBase web UI.
	mapred	Mapred@123	This user is the MapReduce system administrator and has the following permissions: • Submits, stops, and views the MapReduce tasks. • Modifies the Yarn configuration parameters. • Accesses the Yarn and MapReduce web UI.
	spark	Spark@123	This user is the Spark system administrator and has the following permissions: • Accesses the Spark web UI. • Submits Spark tasks.

User Group Information

Default User Group	Description
hadoop	Users added to this user group have the permission to submit tasks to all Yarn queues.
hbase	Common user group. Users added to this user group will not have any additional permission.
hive	Users added to this user group can use Hive.
spark	Common user group. Users added to this user group will not have any additional permission.
supergroup	Users added to this user group can have the administrator permission of HBase, HDFS, and Yarn and can use Hive.
check_sec_ldap	Used to test whether the active LDAP works properly. This user group is generated randomly in a test and automatically deleted after the test is complete. This is an internal system user group used only between components.
Manager_tenant	Tenant system user group, which is an internal system user group used only between components.
System_administrator	MRS cluster system administrator group, which is an internal system user group used only between components.
Manager_viewer	MRS Manager system viewer group, which is an internal system user group used only between components.
Manager_operator	MRS Manager system operator group, which is an internal system user group used only between components.
Manager_auditor	MRS Manager system auditor group, which is an internal system user group used only between components.
Manager_administrator	MRS Manager system administrator group, which is an internal system user group used only between components.
compcommon	Internal system group for accessing public resources in a cluster. All system users and system running users are added to this user group by default.
default_1000	User group created for tenants, which is an internal system user group used only between components.

Default User Group	Description
kafka	Kafka common user group. Users added to this group need to be granted with read and write permission by users in the **kafkaadmin** group before accessing the desired topics.
kafkasuperuser	Users added to this group have permissions to read data from and write data to all topics.
kafkaadmin	Kafka administrator group. Users added to this group have the permissions to create, delete, authorize, as well as read from and write data to all topics.
storm	Storm common user group. Users added to this group have the permissions to submit topologies and manage their own topologies.
stormadmin	Storm administrator user group. Users added to this group have the permissions to submit topologies and manage their own topologies.
opentsdb	Common user group. Users added to this user group will not have any additional permission.
presto	Common user group. Users added to this user group will not have any additional permission.
flume	Common user group. Users added to this user group will not have any additional permission.
launcher-job	MRS internal group, which is used to submit jobs using V2 APIs.

OS User Group	Description
wheel	Primary group of MRS internal running user **omm**.
ficommon	MRS cluster common group that corresponds to **compcommon** for accessing public resource files stored in the OS of the cluster.

Database User

MRS cluster system database users include OMS database users and DBService database users.

NOTE

Do not delete database users. Otherwise, the cluster or components may not work properly.

Type	Default User	Initial Password	Description
OMS database	ommdba	dbChangeMe@123456	OMS database administrator who performs maintenance operations, such as creating, starting, and stopping applications.
	omm	ChangeMe@123456	User for accessing OMS database data.
DBService database	omm	dbserverAdmin@123	Administrator of the GaussDB database in the DBService component.
	hive	HiveUser@	User for Hive to connect to the DBService database.
	hue	HueUser@123	User for Hue to connect to the DBService database.
	sqoop	SqoopUser@	User for Loader to connect to the DBService database.
	ranger	RangerUser@	User for Ranger to connect to the DBService database.

9.14.3 Creating a Role

Scenario

This section describes how to create a role on Manager and authorize and manage Manager and components.

Up to 1000 roles can be created on Manager.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Managing Roles**.

Prerequisites

- You have learned service requirements.
- You have obtained a cluster with Kerberos authentication enabled or a common cluster with the EIP function enabled.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, choose **System** > **Manage Role**.

Step 3 Click **Create Role** and fill in **Role Name** and **Description**.

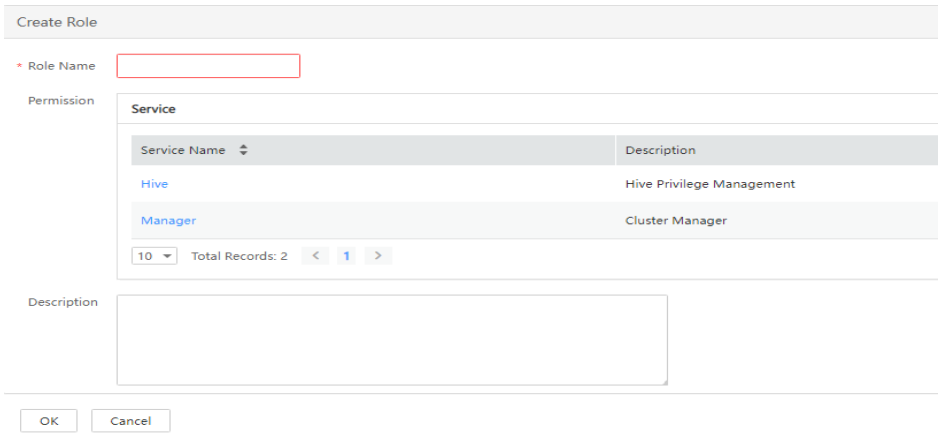

Role Name is mandatory and contains 3 to 30 characters. Only digits, letters, and underscores (_) are allowed. **Description** is optional.

Step 4 In **Permission**, set role permission.

1. Click **Service Name** and select a name in **View Name**.
2. Select one or more permissions.

NOTE

- The **Permission** parameter is optional.
- If you select **View Name** to set component permissions, you can enter a resource name in the **Search** box in the upper right corner and click . The search result is displayed.
- The search scope covers only directories with current permissions. You cannot search subdirectories. Search by keywords supports fuzzy match and is case-insensitive. Results of the next page can be searched.

Table 9-46 Manager permission description

Resource Supporting Permission Management	Permission Setting
Alarm	Authorizes the Manager alarm function. You can select **View** to view alarms and **Management** to manage alarms.
Audit	Authorizes the Manager audit log function. You can select **View** to view audit logs and **Management** to manage audit logs.
Dashboard	Authorizes the Manager overview function. You can select **View** to view the cluster overview.
Hosts	Authorizes the node management function. You can select **View** to view node information and **Management** to manage nodes.

Resource Supporting Permission Management	Permission Setting
Services	Authorizes the service management function. You can select **View** to view service information and **Management** to manage services.
System_cluster_management	Authorizes the MRS cluster management function. You can select **Management** to use the MRS patch management function.
System_configuration	Authorizes the MRS cluster configuration function. You can select **Management** to configure MRS clusters on Manager.
System_task	Authorizes the MRS cluster task function. You can select **Management** to manage periodic tasks of MRS clusters on Manager.
Tenant	Authorizes the Manager multi-tenant management function. You can select **Management** to manage multi-tenants.

Table 9-47 HBase permission description

Resource Supporting Permission Management	Permission Setting
SUPER_USER_GROUP	Grants you HBase administrator permissions.
Global	HBase resource type, indicating the whole HBase.
Namespace	HBase resource type, indicating namespace, which is used to store HBase tables. It has the following permissions: • **Admin** permission to manage the namespace • **Create**: permission to create HBase tables in the namespace • **Read**: permission to access the namespace • **Write**: permission to write data to the namespace • **Execute**: permission to execute the coprocessor (Endpoint)

Resource Supporting Permission Management	Permission Setting
Table	HBase resource type, indicating a data table, which is used to store data. It has the following permissions: • **Admin**: permission to manage a data table • **Create**: permission to create column families and columns in a data table • **Read**: permission to read a data table • **Write**: permission to write data to a data table • **Execute**: permission to execute the coprocessor (Endpoint)
ColumnFamily	HBase resource type, indicating a column family, which is used to store data. It has the following permissions: • **Create**: permission to create columns in a column family • **Read**: permission to read a column family • **Write**: permission to write data to a column family
Qualifier	HBase resource type, indicating a column, which is used to store data. It has the following permissions: • **Read**: permission to read a column • **Write**: permission to write data to a column

By default, permissions of an HBase resource type of each level are shared by resource types of sub-levels. However, the **Recursive** option is not selected by default. For example, if **Read** and **Write** permissions are added to the **default** namespace, they are automatically added to the tables, column families, and columns in the namespace. If a child resource is set after the parent resource, the permission of the child resource is the union of the permissions of the parent resource and the current child resource.

Table 9-48 HDFS permission description

Resource Supporting Permission Management	Permission Setting
Folder	HDFS resource type, indicating an HDFS directory, which is used to store files or subdirectories. It has the following permissions: • **Read**: permission to access the HDFS directory • **Write**: permission to write data to the HDFS directory • **Execute**: permission to perform an operation. It must be selected when you add access or write permission.
Files	HDFS resource type, indicating a file in HDFS. It has the following permissions: • **Read**: permission to access the file • **Write**: permission to write data to the file • **Execute**: permission to perform an operation. It must be selected when you add access or write permission.

Permissions of an HDFS directory of each level are not shared by directory types of sub-levels by default. For example, if **Read** and **Execute** permissions are added to the **tmp** directory, you must select **Recursive** for permissions to be added to subdirectories.

Table 9-49 Hive permission description

Resource Supporting Permission Management	Permission Setting
Hive Admin Privilege	Grants you Hive administrator permissions.
Database	Hive resource type, indicating a Hive database, which is used to store Hive tables. It has the following permissions: • **Select**: permission to query the Hive database • **Delete**: permission to perform the deletion operation in the Hive database • **Insert**: permission to perform the insertion operation in the Hive database • **Create**: permission to perform the creation operation in the Hive database

Resource Supporting Permission Management	Permission Setting
Table	Hive resource type, indicating a Hive table, which is used to store data. It has the following permissions: • **Select**: permission to query the Hive table • **Delete**: permission to perform the deletion operation in the Hive table • **Update**: permission to perform the update operation in the Hive table • **Insert**: permission to perform the insertion operation in the Hive table • **Grant of Select**: permission to grant the **Select** permission to other users using Hive statements • **Grant of Delete**: permission to grant the **Delete** permission to other users using Hive statements • **Grant of Update**: permission to grant the **Update** permission to other users using Hive statements • **Grant of Insert**: permission to grant the **Insert** permission to other users using Hive statements

By default, permissions of a Hive resource type of each level are shared by resource types of sub-levels. However, the **Recursive** option is not selected by default. For example, if **Select** and **Insert** permissions are added to the **default** database, they are automatically added to the tables and columns in the database. If a child resource is set after the parent resource, the permission of the child resource is the union of the permissions of the parent resource and the current child resource.

Table 9-50 Yarn permission description

Resource Supporting Permission Management	Permission Setting
Cluster Admin Operations	Grants you Yarn administrator permissions.
root	Root queue of Yarn. It has the following permissions: • **Submit**: permission to submit jobs in the queue • **Admin**: permission to manage permissions of the current queue

Resource Supporting Permission Management	Permission Setting
Parent Queue	Yarn resource type, indicating a parent queue containing sub-queues. A root queue is a type of a parent queue. It has the following permissions: • **Submit**: permission to submit jobs in the queue • **Admin**: permission to manage permissions of the current queue
Leaf Queue	Yarn resource type, indicating a leaf queue. It has the following permissions: • **Submit**: permission to submit jobs in the queue • **Admin**: permission to manage permissions of the current queue

By default, permissions of a Yarn resource type of each level are shared by resource types of sub-levels. However, the **Recursive** option is not selected by default. For example, if the **Submit** permission is added to the **root** queue, it is automatically added to the sub-queue. Permissions inherited by sub-queues will not be displayed as selected in the **Permission** table. If a child resource is set after the parent resource, the permission of the child resource is the union of the permissions of the parent resource and the current child resource.

Table 9-51 Hue permission description

Resource Supporting Permission Management	Permission Setting
Storage Policy Admin	Grants you storage policy administrator permissions.

Step 5 Click **OK**. Return to **Manage Role**.

----End

Related Tasks

Modifying a role

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage Role**.

Step 3 In the row of the role to be modified, click **Modify** to modify role information.

NOTE

If you modify permissions assigned by the role, it takes 3 minutes to make new configurations take effect.

Step 4 Click **OK**. The modification is complete.

----End

Deleting a role

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage Role**.

Step 3 In the row of the role to be deleted, click **Delete**.

Step 4 Click **OK**. The role is deleted.

----End

9.14.4 Creating a User Group

Scenario

This section describes how to create user groups and specify their operation permissions on Manager. Management of single or multiple users can be unified in the user groups. After being added to a user group, users can obtain operation permissions owned by the user group.

Manager supports a maximum of 100 user groups.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Managing User Groups**.

Prerequisites

- Administrators have learned service requirements and created roles required by service scenarios.
- You have obtained a cluster with Kerberos authentication enabled or a common cluster with the EIP function enabled.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User Group**.

Step 4 Above the user group list, click **Create User Group**.

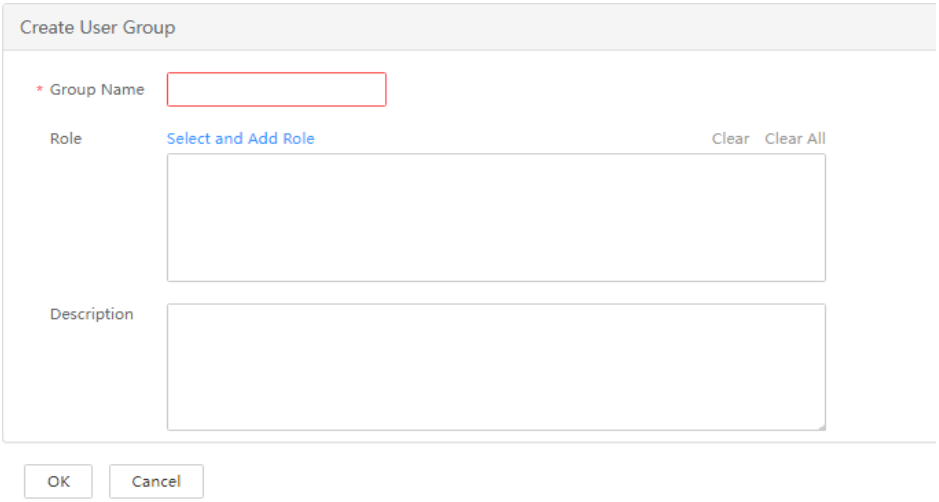

Step 5 Input **Group Name** and **Description**.

Group Name is mandatory and contains 3 to 20 characters. Only digits, letters, and underscores (_) are allowed. **Description** is optional.

Step 6 In **Role**, click **Select and Add Role** to select and add specified roles.

If you do not add the roles, the user group you are creating now does not have the permission to use MRS clusters.

Step 7 Click **OK**.

----End

Related Tasks

Modifying a user group

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User Group**.

Step 3 In the row of a user group to be modified, click **Modify**.

NOTE

If you change role permissions assigned to the user group, it takes 3 minutes to make new configurations take effect.

Step 4 Click **OK**. The modification is complete.

----End

Deleting a user group

Step 1 On MRS Manager, click **System**.

Step 2 In the **Permission** area, click **Manage User Group**.

Step 3 In the row of the user group to be deleted, click **Delete**.

Step 4 Click **OK**. The user group is deleted.

----End

9.14.5 Creating a User

Scenario

This section describes how to create users on Manager based on site requirements and specify their operation permissions to meet service requirements.

Up to 1,000 users can be created on Manager.

If a new password policy needs to be used for a new user's password, follow instructions in **Modifying a Password Policy** to modify the password policy and then perform the following operations to create a user.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Creating a User**.

Prerequisites

- Administrators have learned service requirements and created roles and role groups required by service scenarios.
- You have obtained a cluster with Kerberos authentication enabled or a common cluster with the EIP function enabled.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User**.

Step 4 Above the user list, click **Create User**.

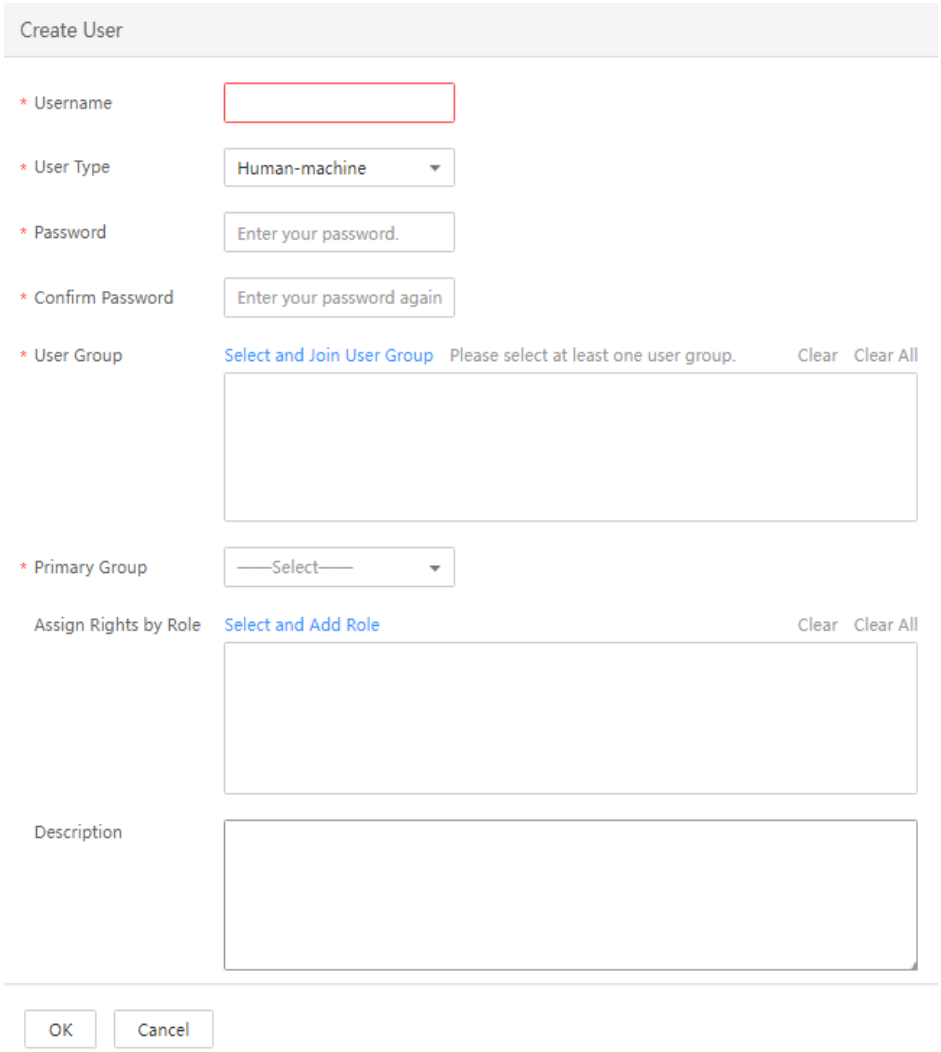

Step 5 Configure parameters as prompted and enter a username in **Username**.

NOTE

- A username that differs only in alphabetic case from an existing username is not allowed. For example, if **User1** has been created, you cannot create **user1**.
- When you use the user you created, enter the exactly correct username, which is case-sensitive.
- **Username** is mandatory and contains 3 to 20 characters. Only digits, letters, and underscores (_) are allowed.
- **root**, **omm**, and **ommdba** are reserved system user. Select another username.

Step 6 Set **User Type** to either **Human-machine** or **Machine-machine**.

- **Human-machine** user: used for MRS Manager O&M scenarios and component client operation scenarios. If you select this user type, you need to enter a password and confirm the password in **Password** and **Confirm Password** accordingly.
- **Machine-machine** users: used for MRS application development scenarios. If you select this user type, you do not need to enter a password, because the password is randomly generated.

Step 7 In **User Group**, click **Select and Join User Group** to select user groups and add users to them.

NOTE

- If roles have been added to user groups, the users can be granted with permissions of the roles.
- If you want to grant new users with Hive permissions, add the users to the Hive group.
- If a user needs to manage tenant resources, the user group must be assigned the **Manager_tenant** role and the role corresponding to the tenant.
- Users created on Manager cannot be added to the user group synchronized using the IAM user synchronization function.

Step 8 In **Primary Group**, select a group as the primary group for users to create directories and files. The drop-down list contains all groups selected in **User Group**.

Step 9 In **Assign Rights by Role**, click **Select and Add Role** to add roles for users based on onsite service requirements.

NOTE

- When you create a user, if permissions of a user group that is granted to the user cannot meet service requirements, you can assign other created roles to the user. It takes 3 minutes to make role permissions granted to the new user take effect.
- Adding a role when you create a user can specify the user rights.
- A new user can access web UIs of HDFS, HBase, Yarn, Spark, and Hue even when roles are not assigned to the user.

Step 10 In **Description**, provide description based on onsite service requirements.

Description is optional.

Step 11 Click **OK**.

If a new user is used in the MRS cluster for the first time, for example, used for logging in to MRS Manager or using the cluster client, the password must be changed. For details, see **Changing the Password of an Operation User**.

----End

9.14.6 Modifying User Information

Scenario

This section describes how to modify user information on Manager, including information about the user group, primary group, role, and description.

This operation is supported only in clusters with Kerberos authentication enabled or common clusters with the EIP function enabled.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Modifying User Information**.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User**.

Step 4 In the row of a user to be modified, click **Modify**.

NOTE

If you change user groups for a user or assign role permissions to a user, it takes 3 minutes to make new configurations take effect.

Step 5 Click **OK**. The modification is complete.

----End

9.14.7 Locking a User

This section describes how to lock users in MRS clusters. A locked user cannot log in to Manager or perform security authentication in the cluster. This operation is supported only in clusters with Kerberos authentication enabled or common clusters with the EIP function enabled.

A locked user can be unlocked by an administrator manually or until the lock duration expires. You can lock a user by using either of the following methods:

- Automatic lock: Set **Number of Password Retries** in **Configure Password Policy**. If user login attempts exceed the parameter value, the user is automatically locked. For details, see **Modifying a Password Policy**.
- Manual lock: The administrator manually locks a user.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Locking a User**.

The following describes how to manually lock a user. **Machine-machine** users cannot be locked.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User**.

Figure 9-19 Managing a user

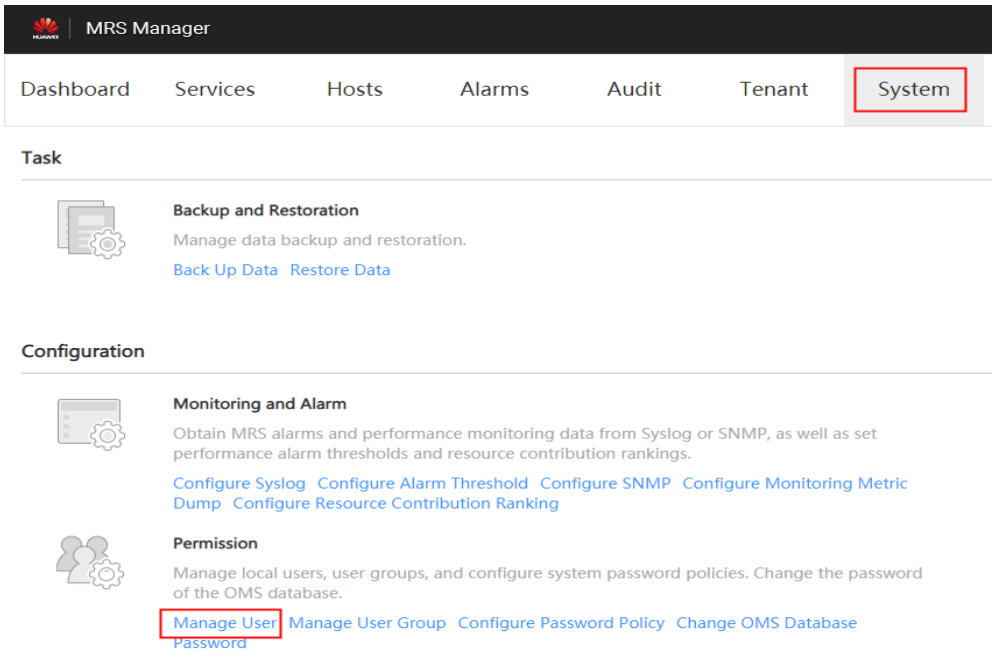

Step 4 In the row of a user you want to lock, click **Lock User**.

Figure 9-20 Locking a user

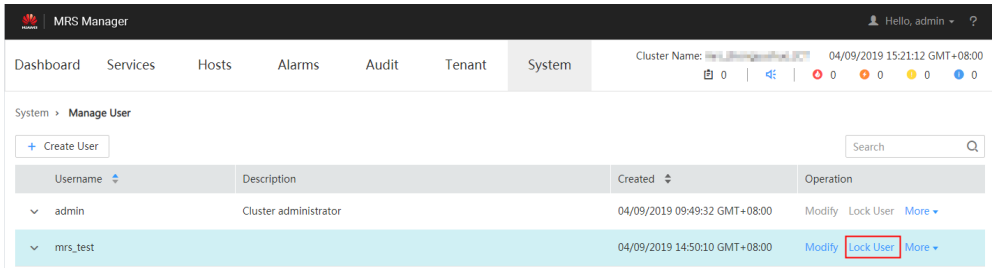

Step 5 In the window that is displayed, click **OK** to lock the user.

----End

9.14.8 Unlocking a User

If a user is locked because the number of login attempts exceeds the value of **Number of Password Retries**, or the user is manually locked by the administrator, the administrator can unlock the user on Manager. This operation is supported only in clusters with Kerberos authentication enabled or common clusters with the EIP function enabled.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Unlocking a User**.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User**.

Figure 9-21 Managing a user

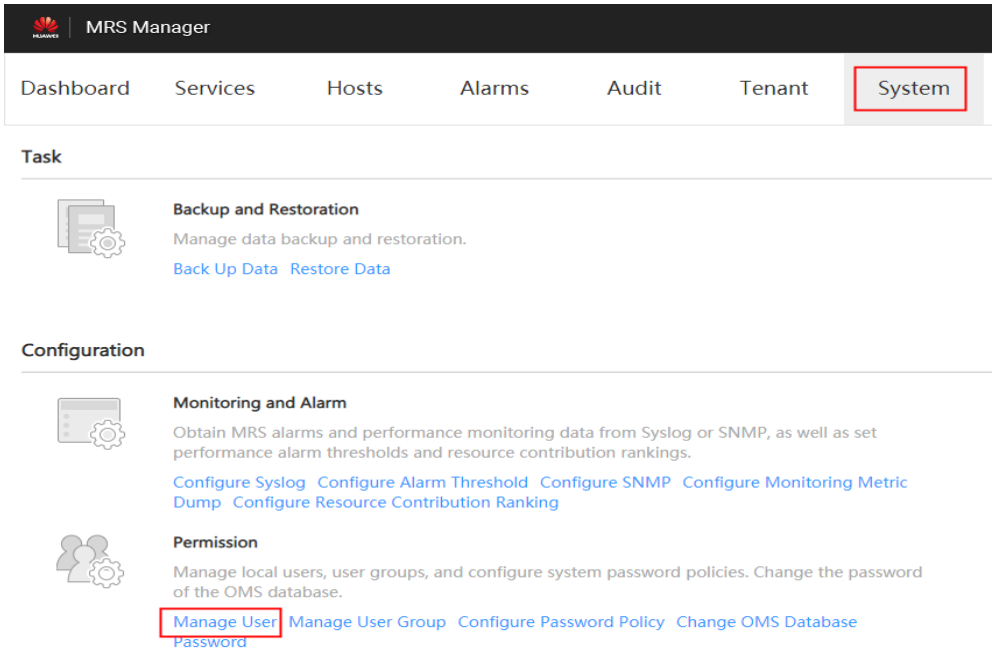

Step 4 In the row of a user to be unlocked, click **Unlock User**.

Figure 9-22 Unlocking a user

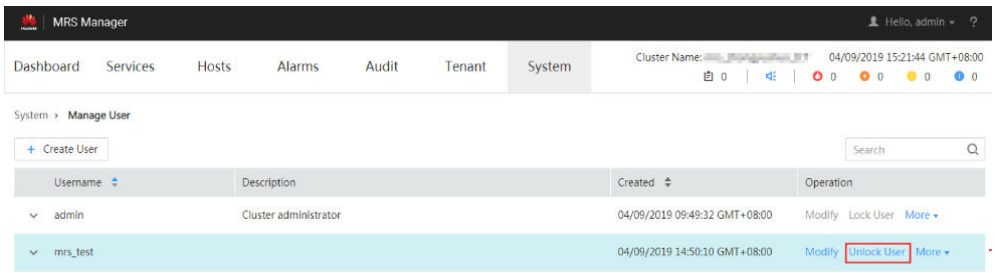

Step 5 In the window that is displayed, click **OK** to unlock the user.

----End

9.14.9 Deleting a User

The administrator can delete an MRS cluster user that is not required on MRS Manager. Deleting a user is allowed only in clusters with Kerberos authentication enabled or normal clusters with the EIP function enabled.

NOTE

If you want to create a new user with the same name as user A after deleting user A who has submitted a job on the client or MRS console, you need to delete user A's residual folders when deleting user A. Otherwise, the newly created user A may fail to submit a job.

To delete residual folders, log in to each Core node in the MRS cluster and run the following commands. In the following commands, **$user** indicates the folder named after the username.

cd /srv/BigData/hadoop/data1/nm/localdir/usercache/

rm -rf $user

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Deleting a User**.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User**.

Figure 9-23 User management

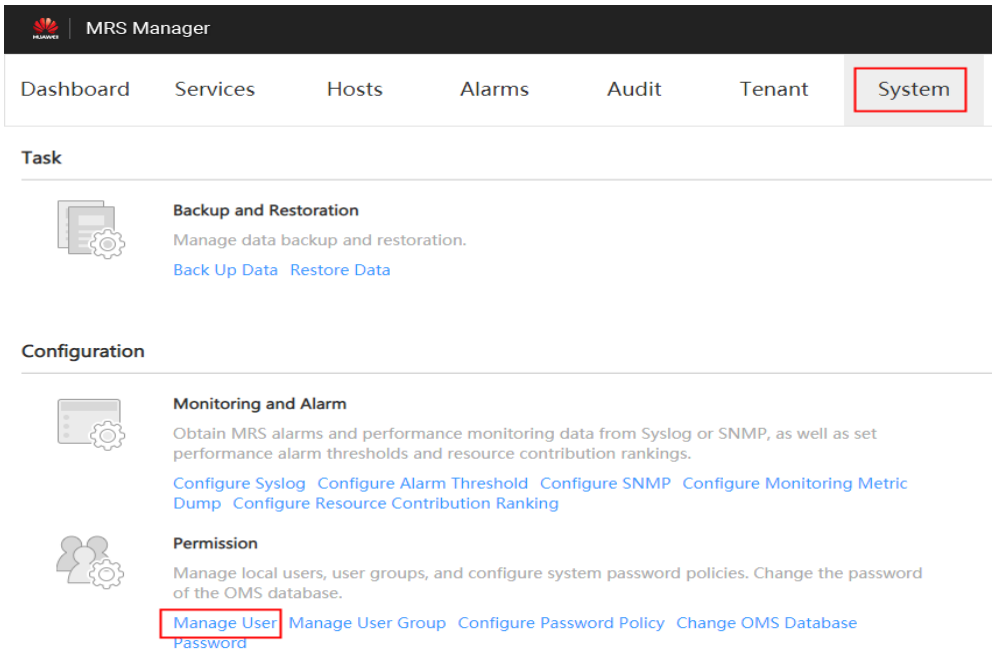

Step 4 In the row that contains the user to be deleted, choose **More** > **Delete**.

Figure 9-24 Deleting a user

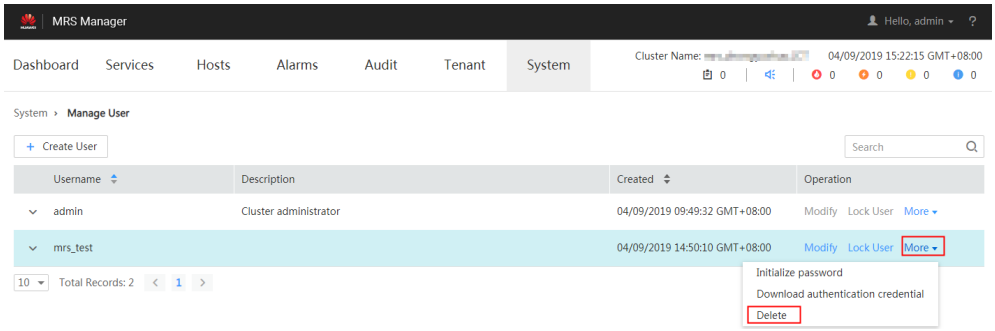

Step 5 Click **OK**.

----End

9.14.10 Changing the Password of an Operation User

Scenario

Passwords of **Human-machine** system users must be regularly changed to ensure MRS cluster security. This section describes how to change passwords on MRS Manager.

If a new password policy needs to be used for the password modified by the user, follow instructions in **Modifying a Password Policy** to modify the password policy and then perform the following operations to modify the password.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Changing a User Password**.

Impact on the System

If you have downloaded a user authentication file, download it again and obtain the keytab file after modifying the password of the MRS cluster user.

Prerequisites

- You have obtained the current password policy.
- You have obtained the URL for accessing MRS Manager.
- You have obtained a cluster with Kerberos authentication enabled or a common cluster with the EIP function enabled.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, move the mouse cursor to  in the upper right corner.

On the menu that is displayed, select **Change Password**.

Step 3 Fill in the **Old Password**, **New Password**, and **Confirm Password**. Click **OK**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

9.14.11 Initializing the Password of a System User

Scenario

This section describes how to initialize a password on Manager if a user forgets the password or the password of a public account needs to be changed regularly. After password initialization, the user must change the password upon the first login. This operation is supported only in clusters with Kerberos authentication enabled or common clusters with the EIP function enabled.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Initializing a Password**.

Impact on the System

If you have downloaded a user authentication file, download it again and obtain the keytab file after initializing the password of the MRS cluster user.

Initializing the Password of a Human-Machine User

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User**.

Step 4 Locate the row that contains the user whose password is to be initialized, choose **More** > **Initialize password**, and change the password as prompted.

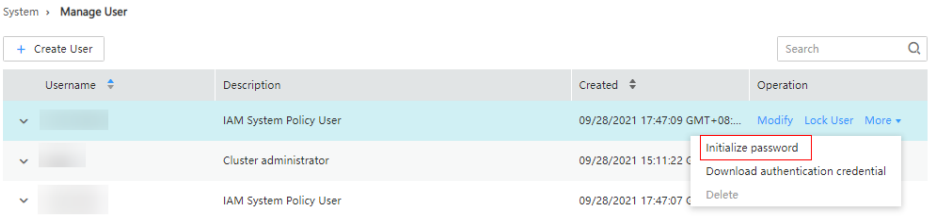

In the window that is displayed, enter the password of the current administrator account and click **OK**. Then in **Initialize password**, click **OK**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

Initializing the Password of a Machine-Machine User

Step 1 Prepare a client based on service conditions and log in to the node with the client installed.

Step 2 Run the following command to switch the user:

sudo su - omm

Step 3 Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

cd /opt/Bigdata/client

Step 4 Run the following command to configure environment variables:

source bigdata_env

Step 5 Run the following command to log in to the console as user **kadmin/admin**:

NOTE

The default password of user **kadmin/admin** is **KAdmin@123**, which will expire upon your first login. Change the password as prompted and keep the new password secure.

kadmin -p kadmin/admin

Step 6 Run the following command to reset the password of a component running user. This operation takes effect on all servers:

cpw *Component running user name*

For example, **cpw oms/manager**.

For the cluster, the default password complexity requirements are as follows:

- The password must contain 8 to 32 characters.
- The password must contain at least three types of the following: uppercase letters, lowercase letters, digits, spaces, and special characters ('~!@#$ %^&*()-_=+\|[{}];:'",<.>/?).
- The password cannot be the username or the reverse username.

----End

9.14.12 Downloading a User Authentication File

Scenario

When a user develops big data applications and runs them in an MRS cluster that supports Kerberos authentication, the user needs to prepare a **Machine-machine**

user authentication file for accessing the MRS cluster. The keytab file in the authentication file can be used for user authentication.

This section describes how to download a **Machine-machine** user authentication file and export the keytab file on Manager. This operation is supported only in clusters with Kerberos authentication enabled or common clusters with the EIP function enabled.

NOTE

Before downloading a **Human-machine** user authentication file, change the password for the user on MRS Manager to make the initial password set by the administrator invalid. Otherwise, the exported keytab file cannot be used. For details, see **Changing the Password of an Operation User**.

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Exporting an Authentication Credential File**.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 In the **Permission** area, click **Manage User**.

Step 4 In the row of a user for whom you want to export the keytab file, choose **More** > **Download authentication credential** to download the authentication file. After the file is automatically generated, save it to a specified path and keep it secure.

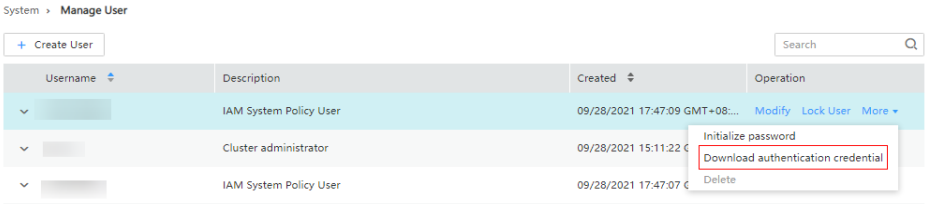

Step 5 Open the authentication file with a decompression program.

- **user.keytab** indicates a user keytab file used for user authentication.
- **krb5.conf** indicates the configuration file of the authentication server. The application connects to the authentication server according to this configuration file information when authenticating users.

----End

9.14.13 Modifying a Password Policy

Scenario

NOTICE

Because password policies are critical to the user management security, modify them based on service security requirements. Otherwise, security risks may be incurred.

This section describes how to set password and user login security rules as well as user lock rules. Password policies set on MRS Manager take effect for **Human-machine** users only, because the passwords of **Machine-machine** users are randomly generated. This operation is supported only in clusters with Kerberos authentication enabled or common clusters with the EIP function enabled.

If a new password policy needs to be used for a new user's password or the password modified by the user, perform the following operations to modify the password policy first, and then follow instructions in **Creating a User** or **Changing the Password of an Operation User**.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Configuring Password Policies**.

Procedure

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 On MRS Manager, click **System**.

Step 3 Click **Configure Password Policy**.

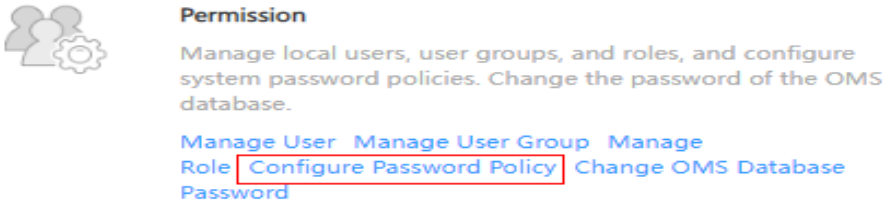

Step 4 Modify password policies as prompted. For parameter details, see **Table 9-52**.

Table 9-52 Password policy parameter description

Parameter	Description
Minimum Password Length	Indicates the minimum number of characters a password contains. The value ranges from 8 to 32. The default value is **8**.
Number of Character Types	Indicates the minimum number of character types a password contains. The character types include uppercase letters, lowercase letters, digits, spaces, and special characters (~`!?,.:;-_'(){}[]/<>@#$ %^&*+|\\=). The value can be **3** or **4**. The default value **3** indicates that the password must contain at least three types of the following characters: uppercase letters, lowercase letters, digits, special characters, and spaces.
Password Validity Period (days)	Indicates the validity period (days) of a password. The value ranges from 0 to 90. Value **0** means that the password is permanently valid. The default value is **90**.
Password Expiration Notification Days	Indicates the number of days to notify password expiration in advance. After the value is set, if the difference between the cluster time and the password expiration time is smaller than this value, the user receives password expiration notifications. When a user logs in to MRS Manager, a message is displayed, indicating that the password is about to expire and asking the user whether to change the password. The value ranges from **0** to *X* (*X* must be set to the half of the password validity period and rounded down). Value **0** indicates that no notification is sent. The default value is **5**.
Interval of Resetting Authentication Failure Count (min)	Indicates the interval (minutes) of retaining incorrect password attempts. The value ranges from 0 to 1440. Value **0** indicates that the number of incorrect password attempts are permanently retained and value **1440** indicates that the number of incorrect password attempts are retained for one day. The default value is **5**.
Number of Password Retries	Indicates the number of consecutive wrong passwords allowed before the system locks the user. The value ranges from 3 to 30. The default value is **5**.

Parameter	Description
Account Lock Duration (min)	Indicates the time period for which a user is locked when the user lockout conditions are met. The value ranges from 5 to 120. The default value is **5**.

----End

9.14.14 Configuring Cross-Cluster Mutual Trust Relationships

Scenario

If cluster A needs to access the resources of cluster B, the mutual trust relationship must be configured between these two clusters.

If no trust relationship is configured, resources of a cluster are available only for users in this cluster. MRS automatically assigns a unique **domain name** for each cluster to define the scope of resources for users.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Configuring Cross-Manager Mutual Trust Between Clusters**.

Impact on the System

- After cross-cluster mutual trust is configured, resources of a cluster become available for users in other cluster. User permission in the clusters must be regularly checked based on service and security requirements.
- After cross-cluster mutual trust is configured, the KrbServer service needs to be restarted and the cluster becomes unavailable during the restart.
- After cross-cluster mutual trust is configured, internal users **krbtgt/***Local cluster domain name*@*External cluster domain name* and **krbtgt/***External cluster domain name*@*Local cluster domain name* are added to the two clusters. The internal users cannot be deleted. The default password is **Crossrealm@123**.

Prerequisites

Both clusters are in the same VPC. If they are not, create a VPC peering connection between them. For details, see **VPC Peering Connection**.

Procedure

Step 1 On the MRS management console, query all security groups of the two clusters.

- If the security groups of the two clusters are the same, go to **Step 3**.
- If the security groups of the two clusters are different, go to **Step 2**.

Step 2 On the VPC management console, choose **Access Control** > **Security Groups**. On the **Security Groups** page, locate the row containing the target security group, click **Manage Rule** in the **Operation** column.

On the **Inbound Rules** tab page, click **Add Rule**. In the **Add Inbound Rule** dialog box that is displayed, configure related parameters.

- **Priority**: The value ranges from 1 to 100. The default value is **1**, which indicates the highest priority. A smaller value indicates a higher priority.
- **Action**: Select **Allow**.
- **Protocol & Port**: Choose **Protocols** > **All**.
- **Type**: Select **IPv4** or **IPv6**.
- **Source**: Select **Security group** and the security group of the peer cluster.
  - To add an inbound rule to the security group of cluster A, set **Source** to **Security group** and the security group of cluster B (peer cluster).
  - To add an inbound rule to the security group of cluster B, set **Source** to **Security group** and the security group of cluster A (peer cluster).

NOTE

For a common cluster with Kerberos authentication disabled, perform step **Step 1** to **Step 2** to configure cross-cluster mutual trust. For a security cluster with Kerberos authentication enabled, after completing the preceding steps, proceed to the following steps for configuration.

Step 3 Log in to MRS Manager of the two clusters separately. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**. Click **Service** and check whether the **Health Status** of all components is **Good**.

- If yes, go to **Step 4**.
- If no, contact technical support personnel for troubleshooting.

Step 4 Query configuration information.

1. On MRS Manager of the two clusters, choose **Services** > **KrbServer** > **Instance**. Query the **OM IP Address** of the two KerberosServer hosts.
2. Click **Service Configuration**. Set **Type** to **All**. Choose **KerberosServer** > **Port** in the navigation tree on the left. Query the value of **kdc_ports**. The default value is **21732**.
3. Click **Realm** and query the value of **default_realm**.

Step 5 On MRS Manager of either cluster, modify the **peer_realms** parameter.

Table 9-53 Parameter description

Parameter	Description
realm_name	Domain name of the mutual-trust cluster, that is, the value of **default_realm** obtained in step **4**.

Parameter	Description
ip_port	KDC address of the peer cluster. Format: *IP address of a KerberosServer node in the peer cluster:kdc_port* The addresses of the two KerberosServer nodes are separated by a comma. For example, if the IP addresses of the KerberosServer nodes are 10.0.0.1 and 10.0.0.2 respectively, the value of this parameter is **10.0.0.1:21732,10.0.0.2:21732**.

NOTE

- To deploy trust relationships with multiple clusters, click  to add items and specify relevant parameters. To delete an item, click  .
- A cluster can have trust relationships with a maximum of 16 clusters. By default, no trust relationship exists between different clusters that are trusted by a local cluster.

Step 6 Click **Save Configuration**. In the dialog box that is displayed, select **Restart the affected services or instances** and click **OK**. If you do not select **Restart the affected services or instances**, manually restart the affected services or instances.

After **Operation successful** is displayed, click **Finish**.

Step 7 Exit MRS Manager and log in to it again. If the login is successful, the configurations are valid.

Step 8 Log in to MRS Manager of the other cluster and repeat step **Step 5** to **Step 7**.

----End

Follow-up Operations

After cross-cluster mutual trust is configured, the service configuration parameters are modified on MRS Manager and the service is restarted. Therefore, you need to prepare the client configuration file again and update the client.

Scenario 1:

Cluster A and cluster B (peer cluster and mutually trusted cluster) are the same type, for example, analysis cluster or streaming cluster. Follow instructions in **Updating a Client (Versions Earlier Than 3.x)** to update the client configuration files of cluster A and B respectively.

- Update the client configuration file of cluster A.
- Update the client configuration file of cluster B.

Scenario 2:

Cluster A and cluster B (peer cluster and mutually trusted cluster) are the different type. Perform the following steps to update the configuration files.

- Update the client configuration file of cluster A to cluster B.
- Update the client configuration file of cluster B to cluster A.

- Update the client configuration file of cluster A.
- Update the client configuration file of cluster B.

Step 1 Log in to MRS Manager of cluster A.

Step 2 Click **Services**, and then **Download Client**.

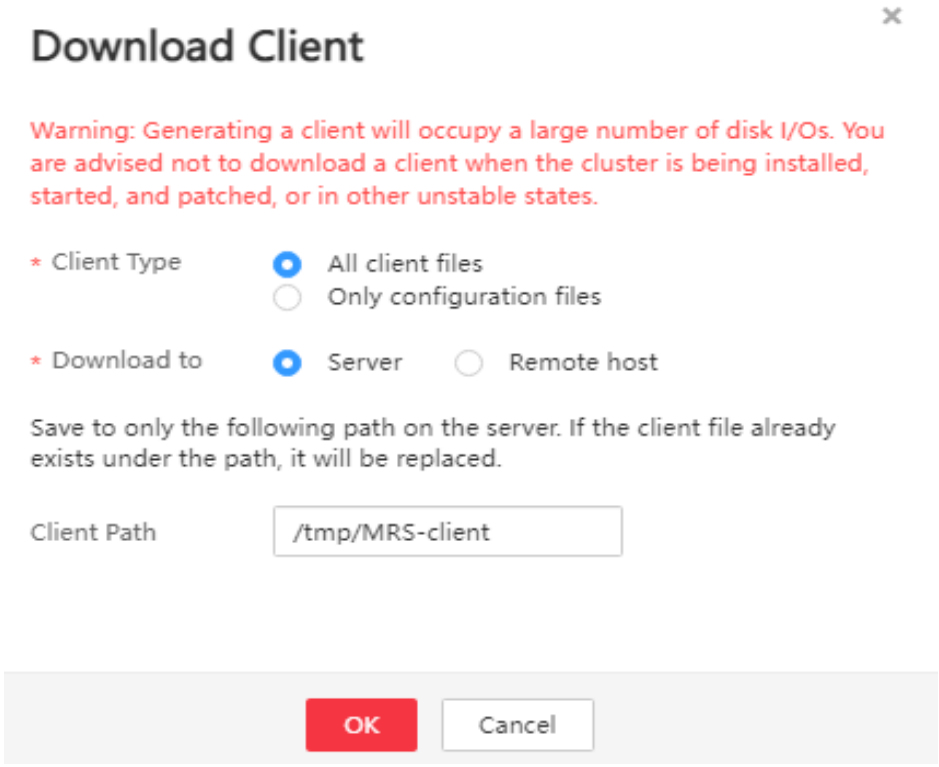

Step 3 Set **Client Type** to **Only configuration files**.

Step 4 Set **Download to** to **Remote host**.

Step 5 Set **Host IP Address** to the IP address of the active Master node of cluster B, **Host Port** to 22, and **Save Path** to **/tmp**.

- If the default port **22** for logging in to cluster B using SSH is changed, set **Host Port** to a new port.
- The value of **Save Path** contains a maximum of 256 characters.

Step 6 Set **Login User** to **root**.

If another user is used, ensure that the user has permissions to read, write, and execute the save path.

Step 7 Select **Password** or **SSH Private Key** for **Login Mode**.

- **Password**: Enter the password of user **root** set during cluster creation.
- **SSH Private Key**: Select and upload the key file used for creating the cluster.

Step 8 Click **OK** to generate a client file.

If the following information is displayed, the client file is saved. Click **Close**.

```
Client files downloaded to the remote host successfully.
```


If the following information is displayed, check the username, password, and security group configurations of the remote host. Ensure that the username and password are correct and an inbound rule of the SSH (22) port has been added to the security group of the remote host. And then, go to **Step 2** to download the client again.

```
Failed to connect to the server. Please check the network connection or parameter settings.
```

Step 9 Log in to the ECS of cluster B using VNC.

All images support Cloud-Init. The preset username for Cloud-Init is **root**, and the password is the one set during cluster creation.

Step 10 Run the following command to switch to the client directory, for example, **/opt/Bigdata/client**:

cd /opt/Bigdata/client

Step 11 Run the following command to update the client configuration of cluster A to cluster B:

sh refreshConfig.sh *Client installation directory Full path of the client configuration file package*

For example, run the following command:

sh refreshConfig.sh /opt/Bigdata/client /tmp/MRS_Services_Client.tar

If the following information is displayed, the configurations have been updated successfully.

```
ReFresh components client config is complete.
Succeed to refresh components client config.
```

NOTE

You can also refer to method 2 in **Updating a Client (Versions Earlier Than 3.x)** to perform operations in **Step 1** to **Step 11**.

Step 12 Repeat step **Step 1** to **Step 11** to update the client configuration file of cluster B to cluster A.

Step 13 Follow instructions in **Updating a Client (Versions Earlier Than 3.x)** to update the client configuration file of the local cluster.

- Update the client configuration file of cluster A.
- Update the client configuration file of cluster B.

----End

9.14.15 Configuring Users to Access Resources of a Trusted Cluster

Scenario

After cross-cluster mutual trust is configured, permission must be configured for users in the local cluster, so that the users can access the same resources in the peer cluster as the users in the peer cluster.

NOTE

The operations described in this section apply only to clusters of versions earlier than MRS 3.x.

For clusters of **MRS 3.*x*** or later, see **Assigning User Permissions After Cross-Cluster Mutual Trust Is Configured**.

Prerequisites

The mutual trust relationship has been configured between two clusters (clusters A and B). The clients of the clusters have been updated.

Procedure

Step 1 Log in to MRS Manager of cluster A and choose **System** > **Manage User**. Check whether cluster A has accounts that are the same as those of cluster B.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 Click on the left side of the username to unfold the detailed user information. Check whether the user group and role to which the user belongs meet the service requirements.

For example, user **admin** of cluster A has the permission to access and create files in the **/tmp** directory of cluster A. Then go to **Step 4**.

Step 3 Create the accounts in cluster A and bind the accounts to the user group and roles required by the services. Then go to **Step 4**.

Step 4 Choose **Service** > **HDFS** > **Instance**. Query the **OM IP Address** of **NameNode (Active)**.

Step 5 Log in to the client of cluster B.

For example, if you have updated the client on the Master2 node, log in to the Master2 node to use the client. For details, see **Using an MRS Client**.

Step 6 Run the following command to access the **/tmp** directory of cluster A.

hdfs dfs -ls hdfs://192.168.6.159:9820/tmp

In the preceding command, **192.168.6.159** is the IP address of the active NameNode of cluster A; **9820** is the default port for communication between the client and the NameNode.

Step 7 Run the following command to create a file in the **/tmp** directory of cluster A:

hdfs dfs -touchz hdfs://192.168.6.159:9820/tmp/mrstest.txt

If you can query the **mrstest.txt** file in the **/tmp** directory of cluster A, the cross-cluster mutual trust is configured successfully.

----End

9.15 Patch Operation Guide

9.15.1 Patch Operation Guide for Versions

If you obtain patch information from the following sources, upgrade the patch according to actual requirements.

- You obtain information about the patch released by MRS from a message pushed by the message center service.
- You obtain information about the patch by accessing the cluster and viewing patch information.

Preparing for Patch Installation

- Follow instructions in **Performing a Health Check** to check cluster status. If the cluster health status is normal, install a patch.
- You need to confirm the target patch to be installed according to the patch information in the patch content.

Installing a Patch

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters > Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's basic information.

Step 3 On the **Patch Information** page, click **Install** in the **Operation** column to install the target patch.

NOTE

- For details about rolling patch operations, see **Supporting Rolling Patches**.
- For the isolated host nodes in the cluster, follow instructions in **Restoring Patches for the Isolated Hosts** to restore the patch.

----End

Uninstalling a Patch

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters > Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's basic information.

Step 3 On the **Patch Information** page, click **Uninstall** in the **Operation** column to uninstall the target patch.

NOTE

- For details about rolling patch operations, see **Supporting Rolling Patches**.
- For the isolated host nodes in the cluster, follow instructions in **Restoring Patches for the Isolated Hosts** to restore the patch.

----End

9.15.2 Supporting Rolling Patches

The rolling patch function indicates that patches are installed or uninstalled for one or more services in a cluster by performing a rolling service restart (restarting

services or instances in batches), without interrupting the services or within a minimized service interruption interval. Services in a cluster are divided into the following three types based on whether they support rolling patch:

- Services supporting rolling patch installation or uninstallation: All businesses or part of them (varying depending on different services) of the services are not interrupted during patch installation or uninstallation.
- Services not supporting rolling patch installation or uninstallation: Businesses of the services are interrupted during patch installation or uninstallation.
- Services with some roles supporting rolling patch installation or uninstallation: Some businesses of the services are not interrupted during patch installation or uninstallation.

Table 9-54 provides services and instances that support or do not support rolling restart in the MRS cluster.

Table 9-54 Services and instances that support or do not support rolling restart

Service	Instance	Whether to Support Rolling Restart
HDFS	NameNode	Yes
	ZKFC	
	JournalNode	
	HttpFS	
	DataNode	
Yarn	ResourceManager	Yes
	NodeManager	
Hive	MetaStore	Yes
	WebHCat	
	HiveServer	
MapReduce	JobHistoryServer	Yes
HBase	HMaster	Yes
	RegionServer	
	ThriftServer	
	RESTServer	
Spark	JobHistory	Yes
	JDBCServer	
	SparkResource	No
Hue	Hue	No

Service	Instance	Whether to Support Rolling Restart
Tez	TezUI	No
Loader	Sqoop	No
ZooKeeper	QuorumPeer	Yes
Kafka	Broker	Yes
	MirrorMaker	No
Flume	Flume	Yes
	MonitorServer	
Storm	Nimbus	Yes
	UI	
	Supervisor	
	LogViewer	

Installing a Patch

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's basic information.

Step 3 On the **Patch Information** page, click **Install** in the **Operation** column.

Step 4 On the **Warning** page, enable or disable **Rolling Patch**.

Figure 9-25 Rolling patch installation

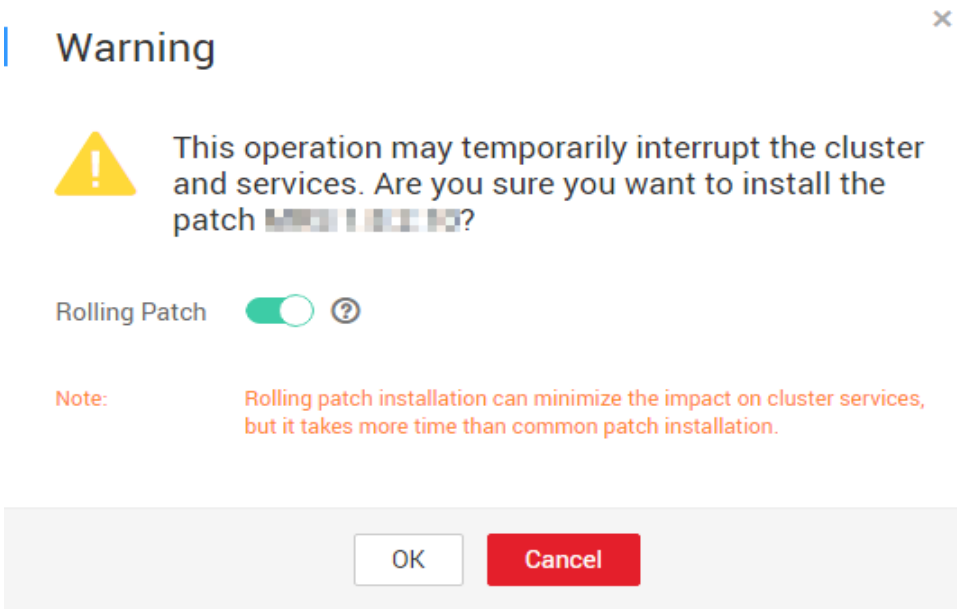

NOTE

- Enabling the rolling patch installation function: Services are not stopped before patch installation, and rolling service restart is performed after the patch installation. This minimizes the impact on cluster services but takes more time than common patch installation.
- Disabling the rolling patch uninstallation function: All services are stopped before patch uninstallation, and all services are restarted after the patch uninstallation. This temporarily interrupts the cluster and the services but takes less time than rolling patch uninstallation.
- The rolling patch installation function is not available in clusters with less than two Master nodes and three Core nodes.

Step 5 Click **OK** to install the target patch.

Step 6 View the patch installation progress.

1. Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.
2. Choose **System** > **Manage Patch**. On the **Manage Patch** page, you can view the patch installation progress.

NOTE

For the isolated host nodes in the cluster, follow instructions in **Restoring Patches for the Isolated Hosts** to restore the patch.

----End

Uninstalling a Patch

Step 1 Log in to the MRS management console.

Step 2 Choose **Clusters** > **Active Clusters** and click the name of the cluster to be queried to enter the page displaying the cluster's basic information.

Step 3 On the **Patch Information** page, click **Uninstall** in the **Operation** column.

Step 4 On the **Warning** page, enable or disable **Rolling Patch**.

Figure 9-26 Rolling patch uninstallation

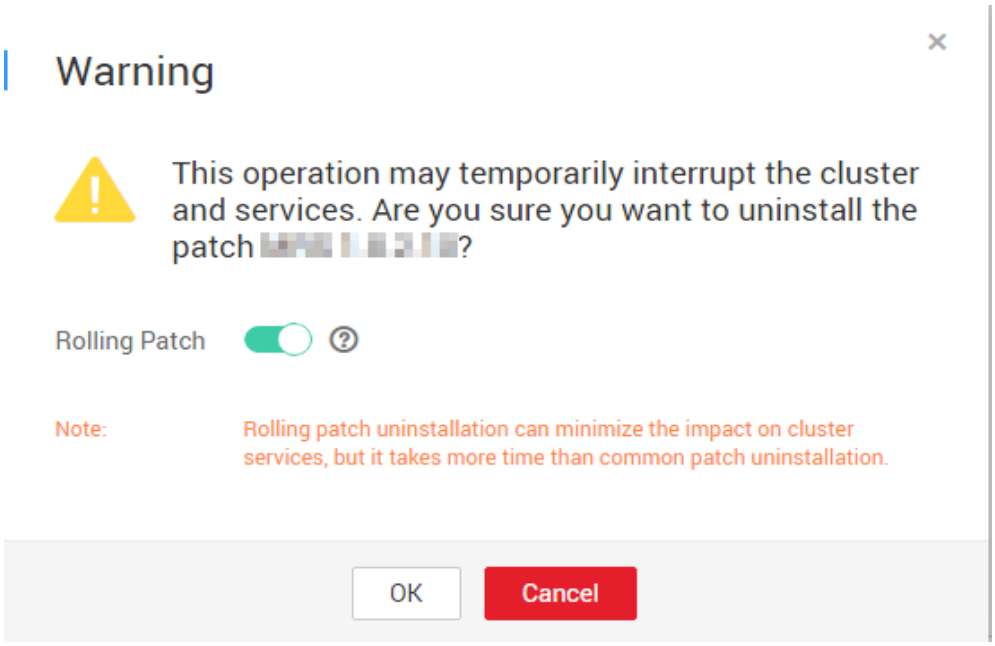

NOTE

- Enabling the rolling patch uninstallation function: Services are not stopped before patch uninstallation, and rolling service restart is performed after the patch uninstallation. This minimizes the impact on cluster services but takes more time than common patch uninstallation.
- Disabling the rolling patch uninstallation function: All services are stopped before patch uninstallation, and all services are restarted after the patch uninstallation. This temporarily interrupts the cluster and the services but takes less time than rolling patch uninstallation.
- The rolling patch uninstallation function is not available in clusters with less than two Master nodes and three Core nodes.

Step 5 Click **OK** to uninstall the target patch.

Step 6 View the patch uninstallation progress.

1. Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.
2. Choose **System** > **Manage Patch**. On the **Manage Patch** page, you can view the patch uninstallation progress.

NOTE

For the isolated host nodes in the cluster, follow instructions in **Restoring Patches for the Isolated Hosts** to restore the patch.

----End

9.16 Restoring Patches for the Isolated Hosts

If some hosts are isolated in a cluster, perform the following operations to restore patches for these isolated hosts after patch installation on other hosts in the

cluster. After patch restoration, versions of the isolated host nodes are consistent with those are not isolated.

Step 1 Access MRS Manager. For details, see **Accessing MRS Manager (MRS 2.x or Earlier)**.

Step 2 Choose **System** > **Manage Patch**. The **Manage Patch** page is displayed.

Step 3 In the **Operation** column, click **View Details**.

Step 4 On the patch details page, select host nodes whose **Status** is **Isolated**.

Step 5 Click **Select and Restore** to restore the isolated host nodes.

Figure 9-27 Restoring patches for the isolated hosts

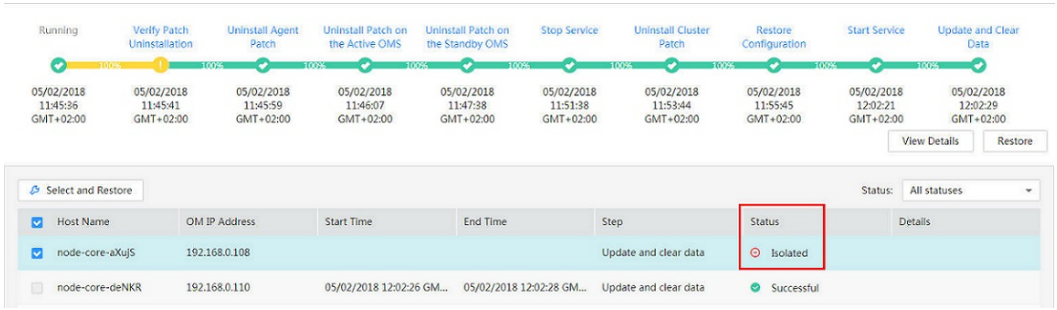

----End

9.17 Rolling Restart

After modifying the configuration items of a big data component, you need to restart the corresponding service to make new configurations take effect. If you use a normal restart mode, all services or instances are restarted concurrently, which may cause service interruption. To ensure that services are not affected during service restart, you can restart services or instances in batches by rolling restart. For instances in active/standby mode, a standby instance is restarted first and then an active instance is restarted. Rolling restart takes longer than normal restart.

Table 9-55 provides services and instances that support or do not support rolling restart in the MRS cluster.

Table 9-55 Services and instances that support or do not support rolling restart

Service	Instance	Whether to Support Rolling Restart
HDFS	NameNode	Yes
	ZKFC	
	JournalNode	
	HttpFS	
	DataNode	

Service	Instance	Whether to Support Rolling Restart
Yarn	ResourceManager	Yes
	NodeManager	
Hive	MetaStore	Yes
	WebHCat	
	HiveServer	
MapReduce	JobHistoryServer	Yes
HBase	HMaster	Yes
	RegionServer	
	ThriftServer	
	RESTServer	
Spark	JobHistory	Yes
	JDBCServer	
	SparkResource	No
Hue	Hue	No
Tez	TezUI	No
Loader	Sqoop	No
ZooKeeper	Quorumpeer	Yes
Kafka	Broker	Yes
	MirrorMaker	No
Flume	Flume	Yes
	MonitorServer	
Storm	Nimbus	Yes
	UI	
	Supervisor	
	Logviewer	

Restrictions

- Perform a rolling restart during off-peak hours.
  - Otherwise, a rolling restart failure may occur. For example, if the throughput of Kafka is high (over 100 MB/s) during the Kafka rolling restart, the Kafka rolling restart may fail.

  - For example, if the requests per second of each RegionServer on the native interface exceed 10,000 during the HBase rolling restart, you need to increase the number of handles to prevent a RegionServer restart failure caused by heavy loads during the restart.
- Before the restart, check the number of current requests of HBase. If requests of each RegionServer on the native interface exceed 10,000, increase the number of handles to prevent a failure.
- If the number of Core nodes in a cluster is less than six, services may be affected for a short period of time.
- Preferentially perform a rolling instance or service restart and select **Only restart instances whose configurations have expired**.

Performing a Rolling Service Restart

Step 1 On MRS Manager, click **Services** and select a service for which you want to perform a rolling restart.

Step 2 On the **Service Status** tab page, click **More** and select **Perform Rolling Service Restart**.

Figure 9-28 Service status

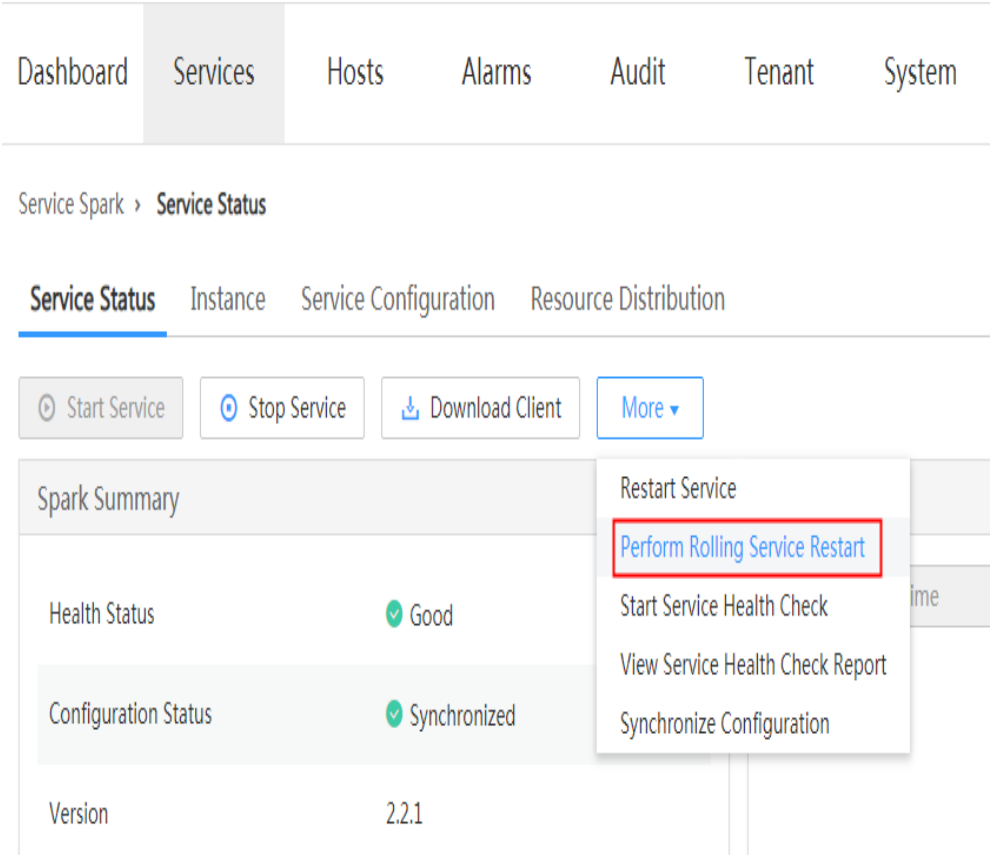

Step 3 After you enter the administrator password, the **Perform Rolling Service Restart** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the service.

Figure 9-29 Performing a rolling service restart

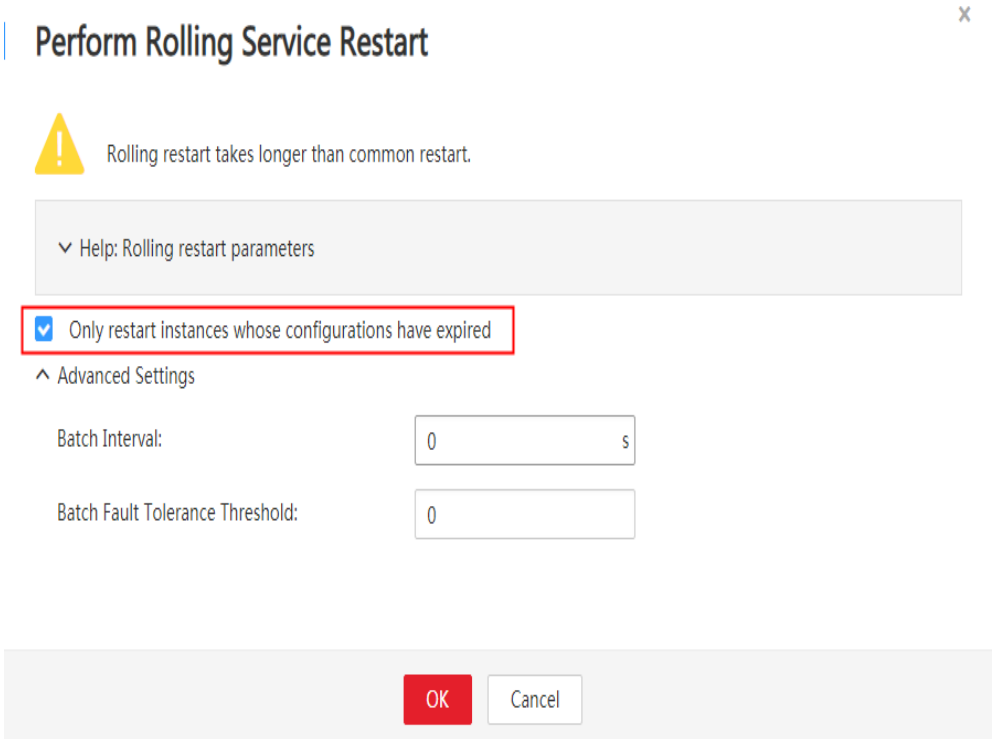

Step 4 After the rolling restart task is complete, click **Finish**.

Figure 9-30 Finishing the rolling service restart

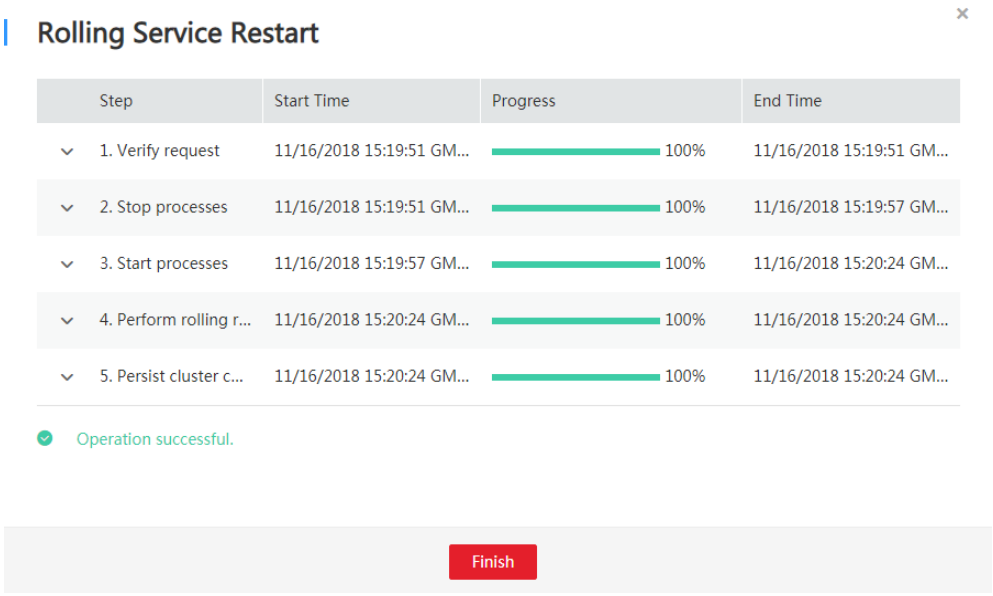

----End

Performing a Rolling Instance Restart

Step 1 On MRS Manager, click **Services** and select a service for which you want to perform a rolling restart.

Step 2 On the **Instance** tab page, select the instance to be restarted. Click **More** and select **Perform Rolling Instance Restart**.

Figure 9-31 Service instance

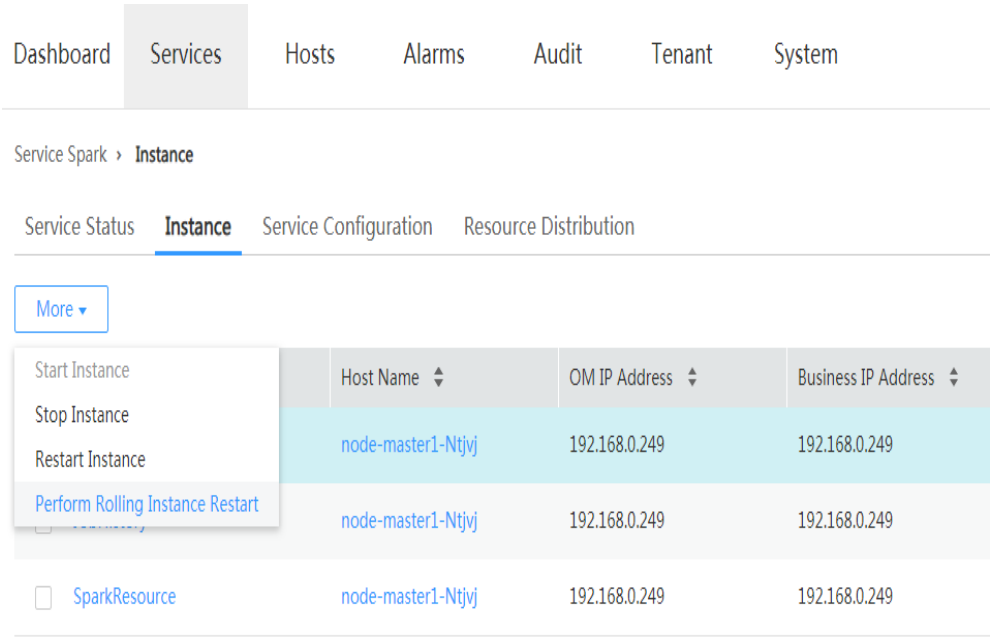

Step 3 After you enter the administrator password, the **Perform Rolling Instance Restart** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the instance.

Step 4 After the rolling restart task is complete, click **Finish**.

----End

Perform a Rolling Cluster Restart

Step 1 On MRS Manager, click **Services**. The **Services** page is displayed.

Step 2 Click **More** and select **Perform Rolling Cluster Restart**.

Figure 9-32 Services

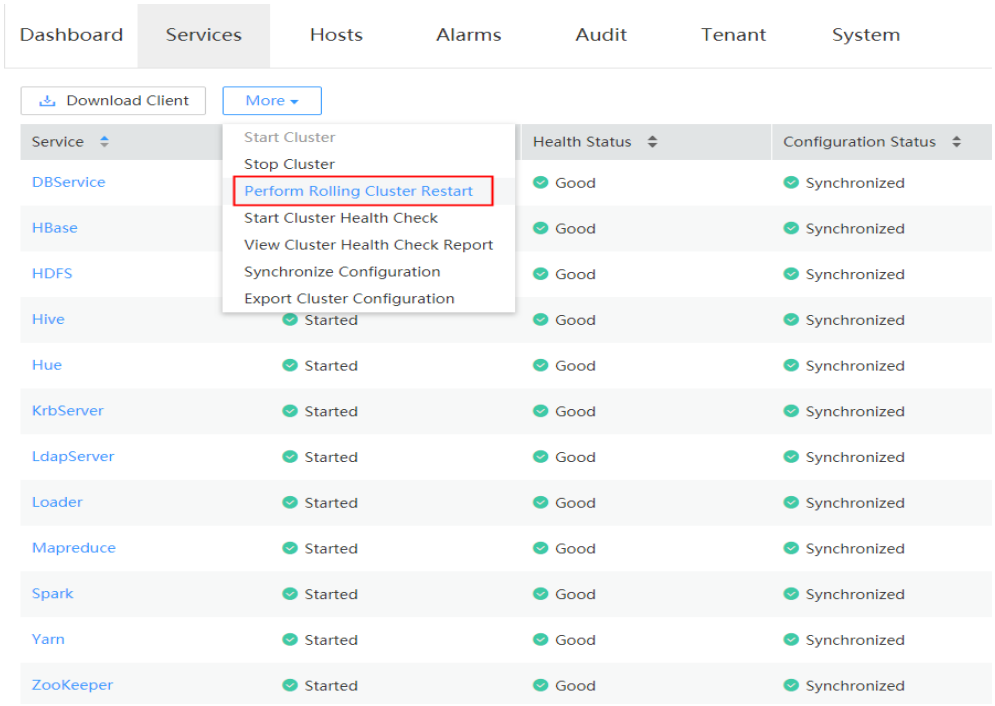

Step 3 After you enter the administrator password, the **Perform Rolling Cluster Restart** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the cluster.

Step 4 After the rolling restart task is complete, click **Finish**.

----End

Rolling Restart Parameter Description

Table 9-56 describes rolling restart parameters.

Table 9-56 Rolling restart parameter description

Parameter	Description
Only restart instances whose configurations have expired	Specifies whether to restart only the modified instances in a cluster.
Data Node Instances to Be Batch Restarted	Specifies the number of instances that are restarted in each batch when the batch rolling restart strategy is used. The default value is **1**. The value ranges from 1 to 20. This parameter is valid only for data nodes.

Parameter	Description
Batch Interval	Specifies the interval between two batches of instances for rolling restart. The default value is **0**. The value ranges from 0 to 2147483647. The unit is second. Note: Setting the batch interval parameter can increase the stability of the big data component process during the rolling restart. You are advised to set this parameter to a non-default value, for example, 10.
Batch Fault Tolerance Threshold	Specifies the tolerance times when the rolling restart of instances fails to be executed in batches. The default value is **0**, which indicates that the rolling restart task ends after any batch of instances fails to be restarted. The value ranges from 0 to 214748364.

Procedure in a Typical Scenario

Step 1 On MRS Manager, click **Services** and select HBase. The HBase service page is displayed.

Step 2 Click the **Service Configuration** tab, and modify an HBase parameter. After the following dialog box is displayed, click **OK** to save the configurations.

Figure 9-33 Saving configurations

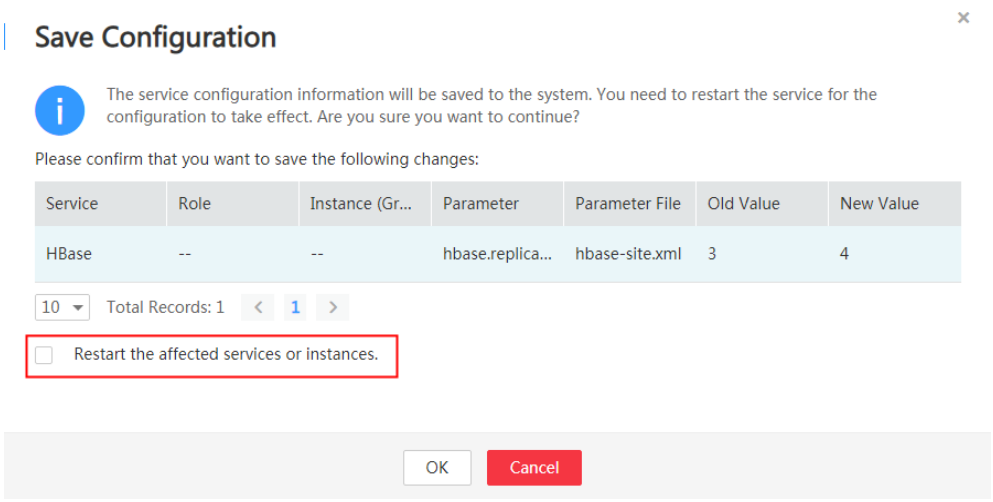

NOTE

Do not select **Restart the affected services or instances**. This option indicates a normal restart. If you select this option, all services or instances will be restarted, which may cause service interruption.

Step 3 After saving the configurations, click **Finish**.

Step 4 Click the **Service Status** tab.

Step 5 On the **Service Status** tab page, click **More** and select **Perform Rolling Service Restart**.

Figure 9-34 Service status-rolling restart

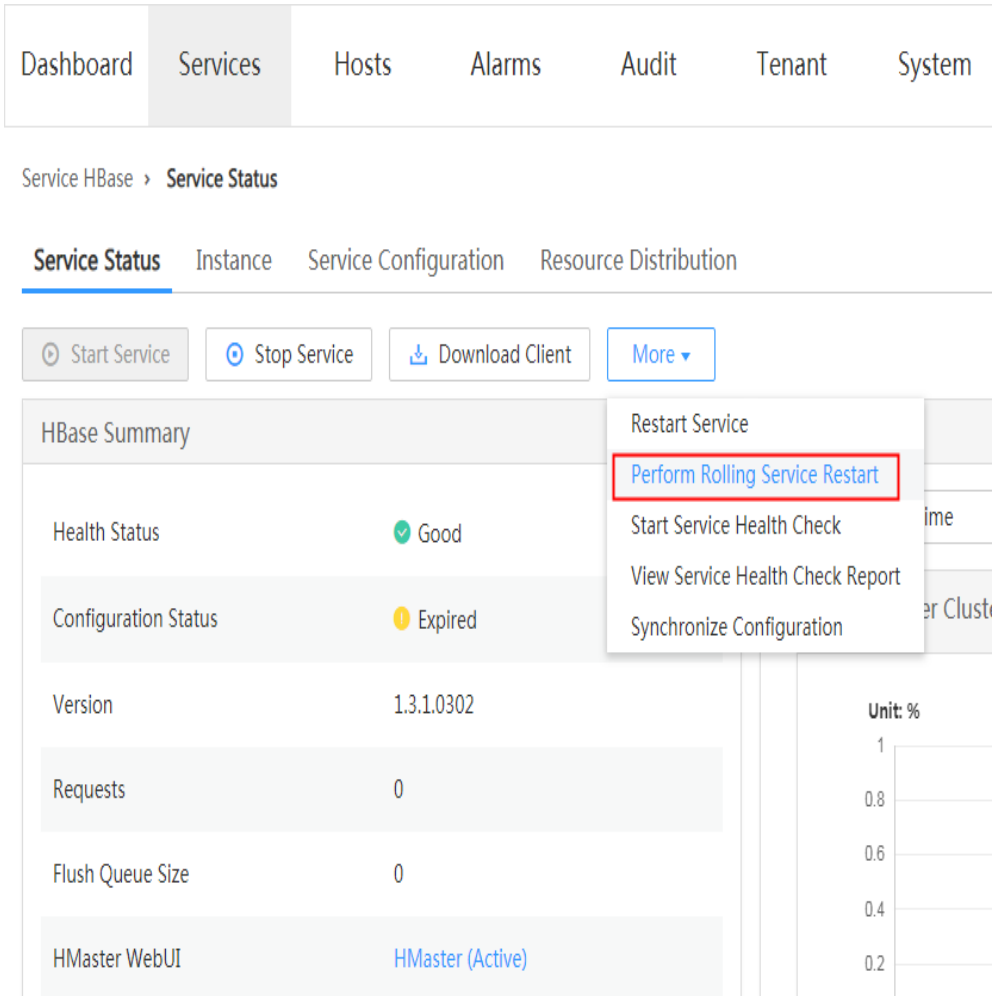

Step 6 After you enter the administrator password, the **Perform Rolling Service Restart** page is displayed. Select **Only restart instances whose configurations have expired** and click **OK** to perform rolling restart for the service.

Figure 9-35 Configuring the rolling service restart

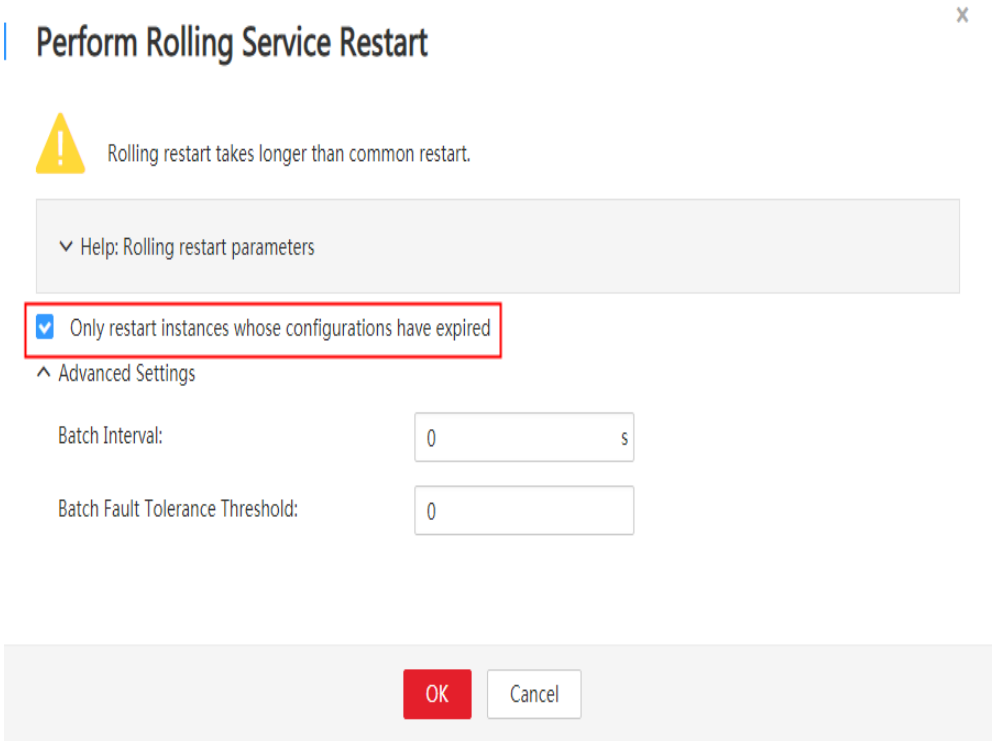

Step 7 After the rolling restart task is complete, click **Finish**.

Figure 9-36 Finishing the rolling service restart

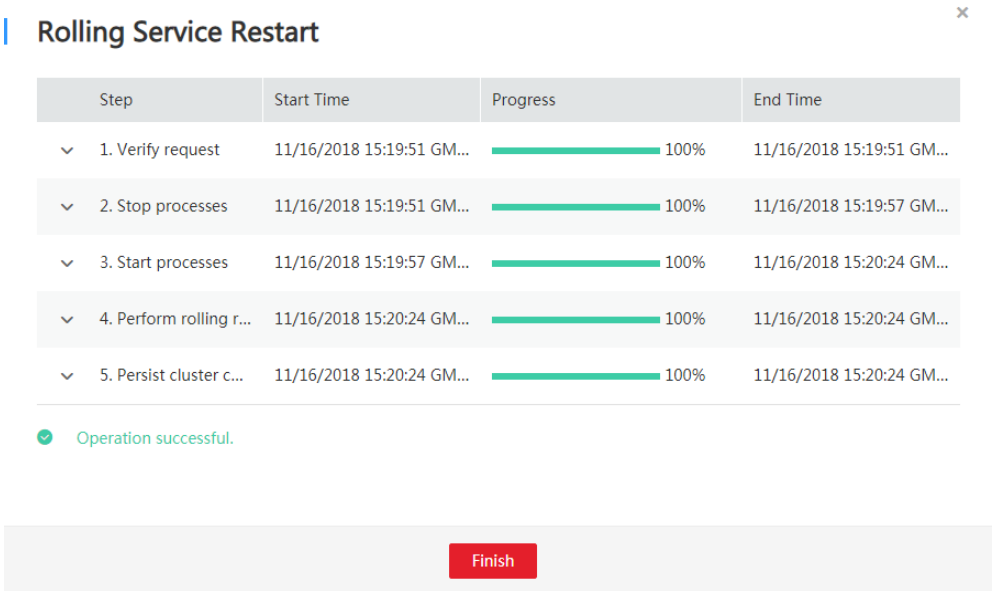

----End

10 Alarm Reference (Applicable to MRS 3.x)

10.1 ALM-12001 Audit Log Dumping Failure

Description

Cluster audit logs need to be dumped on a third-party server due to the local historical data backup policy. The system starts to check the dump server at 3 a.m. every day. If the dump server meets the configuration conditions, audit logs can be successfully dumped. This alarm is generated when the audit log dump fails if the disk space of the dump directory on the third-party server is insufficient or a user changes the username, password, or dump directory of the dump server.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12001	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

System can store a maximum of only 50 dump files locally. If the fault persists on the dump server, the local audit logs may be lost.

Possible Causes

- The network connection is abnormal.
- The username, password, or dump directory of the dump server does not meet the configuration conditions.
- The disk space of the dump directory is insufficient.

Procedure

Check whether the network connection is normal.

Step 1 On the FusionInsight Manager home page, choose **Audit > Configurations**.

Step 2 Check whether the SFTP IP on the dump configuration page is valid.

Log in to the node where Manager is located as user **root** and run the **ping** command to check whether the network connection between the SFTP server and the cluster is normal.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Repair the network connection, reset the SFTP password, and click **OK**.

Step 4 Wait for 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the username, password, or dump directory are correct.

Step 5 On the dump configuration page, check whether the username, password, and dump directory of the third-party server are correct.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 Change the username, password, or dump directory, reset the SFTP password and click **OK**.

Step 7 Wait for 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check whether the disk space of the dump directory is sufficient.

Step 8 Log in to the third-party server as user **root** and run the **df** command to check whether the disk space of the dump directory of the third-party server exceeds 100 MB.

- If yes, go to **Step 11**.
- If no, go to **Step 9**.

Step 9 Expand disk space capacity for the third-party server, Reset the SFTP password and click **OK**

Step 10 Wait for 2 minutes, view real-time alarms and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Reset the dump rule.

Step 11 On the FusionInsight Manager home page, choose **Audit > Configurations**.

Step 12 Reset dump rules, set the parameters properly, and click **OK**.

Step 13 Wait for 2 minutes, view real-time alarms and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Collect fault information.

Step 14 On the FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 15 Select **OmmServer** from the **Service** and click **OK**.

Step 16 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 17 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.2 ALM-12004 OLdap Resource Abnormal

Description

The system checks LDAP resources every 60 seconds. This alarm is generated when the system detects that the LDAP resources in Manager are abnormal for six consecutive times.

This alarm is cleared when the Ldap resource in the Manager recovers and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12004	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The Manager and component WebUI authentication services are unavailable and cannot provide security authentication and user management functions for web upper-layer services. Users may be unable to log in to the WebUIs of Manager and components.

Possible Causes

The LdapServer process in the Manager is abnormal.

Procedure

Check whether the LdapServer process in the Manager is normal.

Step 1 Log in the Manager node in the cluster as user **omm**.

Log in to FusionInsight Manager using the floating IP address, and run the **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** command to check the information about the current Manager two-node cluster.

Step 2 Run **ps -ef | grep slapd** command to check whether the LdapServer resource process in the **${BIGDATA_HOME}/om-server/om/** in the process configuration file is running properly.

NOTE

You can determine that the resource is normal by checking the following information:

1. After the **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** command runs, **ResHAStatus** of the OLdap is **Normal**.
2. After the **ps -ef | grep slapd** command runs, the slapd process of port 21750 can be viewed.
   - If yes, go to **Step 3**.
   - If no, go to **Step 4**.

Step 3 Run the **kill -2** *ldap pid* command to restart the LdapServer process and wait for 20 seconds. The HA starts the OLdap process automatically. Check whether the current OLdap resource is in normal state.

- If yes, the operation is complete.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On the FusionInsight Manager home page, choose **O&M** > **Log > Download**.

Step 5 Select **OmsLdapServer** and **OmmServer** from the **Service** and click **OK**.

Step 6 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.3 ALM-12005 OKerberos Resource Abnormal

Description

The alarm module checks the status of the Kerberos resource in Manager every 80 seconds. This alarm is generated when the alarm module detects that the Kerberos resources are abnormal for six consecutive times.

This alarm is cleared when the Kerberos resource recovers and the alarm handling is complete.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12005	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The component WebUI authentication services are unavailable and cannot provide security authentication functions for web upper-layer services. Users may be unable to log in to FusionInsight Manager and the WebUIs of components.

Possible Causes

The OLdap resource on which the Okerberos depends is abnormal.

Procedure

Check whether the OLdap resource on which the Okerberos depends is abnormal in the Manager.

Step 1 Log in the Manager node in the cluster as user **omm**.

Log in to FusionInsight Manager using the floating IP address, and run the **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** command to check the information about the current Manager two-node cluster.

Step 2 Run the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the OLdap resource status managed by HA is normal. (In single-node mode, the OLdap resource is in the Active_normal state; in the two-node mode, the OLdap resource is in the Active_normal state on the active node and in the Standby_normal state on the standby node.)

- If yes, go to **Step 4**.

- If no, go to **Step 3**.

Step 3 See the procedure in **ALM-12004 OLdap Resource Abnormal** to resolve the problem. After the OLdap resource status recovers, check whether the OKerberos resource status is normal.

- If yes, the operation is complete.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On the FusionInsight Manager home page, choose **O&M** > **Log > Download**.

Step 5 Select **OmsKerberos** and **OmmServer** from the **Service** and click **OK**.

Step 6 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.4 ALM-12006 Node Fault

Alarm Description

Controller checks the NodeAgent heartbeat every 30 seconds. If Controller does not receive heartbeat messages from a NodeAgent, it attempts to restart the NodeAgent process. This alarm is generated if the NodeAgent fails to be restarted for three consecutive times.

This alarm is cleared when Controller can properly receive the status report of the NodeAgent.

NOTE

In MRS 3.3.0 and later versions, the alarm name is changed to **NodeAgent Process Is Abnormal**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12006	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Services on the node are unavailable.

Possible Causes

- The network is disconnected, the hardware is faulty, or the operating system runs slowly.
- The memory of the NodeAgent process is insufficient.
- The NodeAgent process is faulty.

Handling Procedure

Check whether the network is disconnected, whether the hardware is faulty, or whether the operating system runs commands slowly.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click ⌄ in the row containing the alarm, click the host name, and view the IP address of the host for which the alarm is generated.

Step 2 Log in to the active management node as user **root**.

NOTE

If the faulty node is the active management node and fails login, the network of the active management node may be faulty. In this case, go to **Step 4**.

Step 3 Run the **ping** *IP address of the faulty host* command to check whether the faulty node is reachable.

- If yes, go to **Step 12**.
- If no, go to **Step 4**.

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 6**.

Step 6 Contact the hardware administrator to check whether the hardware (CPU or memory) of the node is faulty.

- If yes, go to **Step 7**.
- If no, go to **Step 12**.

Step 7 Repair or replace faulty components and restart the node. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Step 8 If a large number of node faults are reported in the cluster, the floating IP addresses may be abnormal. As a result, Controller cannot detect the NodeAgent heartbeat.

Log in to any management node and view the **/var/log/Bigdata/omm/oms/ha/scriptlog/floatip.log** log to check whether the logs generated one to two minutes before and after the faults occur are complete.

For example, a complete log is in the following format:

```
2017-12-09 04:10:51,000 INFO (floatip) Read from ${BIGDATA_HOME}/om-server_*/om/etc/om/
routeSetConf.ini,value is : yes
2017-12-09 04:10:51,000 INFO (floatip) check wsNetExport : eth0 is up.
2017-12-09 04:10:51,000 INFO (floatip) check omNetExport : eth0 is up.
2017-12-09 04:10:51,000 INFO (floatip) check wsInterface : eRth0:oms, wsFloatIp: XXX.XXX.XXX.XXX.
2017-12-09 04:10:51,000 INFO (floatip) check omInterface : eth0:oms, omFloatIp: XXX.XXX.XXX.XXX.
2017-12-09 04:10:51,000 INFO (floatip) check  wsFloatIp : XXX.XXX.XXX.XXX is reachable.
2017-12-09 04:10:52,000 INFO (floatip) check  omFloatIp : XXX.XXX.XXX.XXX is reachable.
```

- If yes, go to **Step 12**.
- If no, go to **Step 9**.

Step 9 Check whether the omNetExport log is printed after the wsNetExport is detected or whether the interval for printing two logs exceeds 10 seconds or longer.

- If yes, go to **Step 10**.
- If no, go to **Step 12**.

Step 10 View the **/var/log/message** file of the OS to check whether sssd frequently restarts or nscd exception information is displayed when the fault occurs.

sssd restart example

```
Feb  7 11:38:16 10-132-190-105 sssd[pam]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd[be[default]]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd: Starting up
Feb  7 11:38:16 10-132-190-105 sssd[be[default]]: Starting up
Feb  7 11:38:16 10-132-190-105 sssd[nss]: Starting up
Feb  7 11:38:16 10-132-190-105 sssd[pam]: Starting up
```

Example nscd exception information

```
Feb 11 11:44:42 10-120-205-33 nscd: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780:
Can't contact LDAP server
Feb 11 11:44:43 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780:
Can't contact LDAP server
```



```
Feb 11 11:44:44 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.92:21780: Can't contact LDAP server
```

- If yes, go to **Step 11**.
- If no, go to **Step 12**.

Step 11 Check whether the LdapServer node is faulty, for example, the service IP address is unreachable or the network latency is too high. If the fault occurs periodically, locate and eliminate it and run the **top** command to check whether abnormal software exists.

Check whether the memory of the NodeAgent process is insufficient.

Step 12 Log in to the faulty node as user **root** and run the following command to view the NodeAgent process logs:

vi /var/log/Bigdata/nodeagent/scriptlog/agent_gc.log.*.current

Step 13 Check whether the log file contains an error indicating that the metaspace size or heap memory size is insufficient.

- If yes, go to **Step 14**.
- If no, go to **Step 17**.

Step 14 Run the **su - omm** command to switch to user **omm**, edit the corresponding file based on the cluster version, increase the values of **nodeagent.Xms** (initial heap memory) and **nodeagent.Xmx** (maximum heap memory), and save the modification.

The path of the file containing the parameters is as follows:

- Versions earlier than MRS 3.2.1: **/opt/Bigdata/om-agent/nodeagent/bin/nodeagent_ctl.sh**
- MRS 3.2.1 or later: **$NODE_AGENT_HOME/etc/agent/nodeagent.properties**

Step 15 Run the following commands to restart the NodeAgent service:

sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/stop-agent.sh

sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/start-agent.sh

Step 16 Wait a moment and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 17**.

Check whether the NodeAgent process is faulty.

Step 17 Log in to the faulty node as user **omm** and run the following command:

ps -ef | grep "Dprocess.name=nodeagent" | grep -v grep

Step 18 Check whether the command output is empty.

- If yes, go to **Step 19**.
- If no, go to **Step 21**.

Step 19 View the NodeAgent startup and run logs to locate the fault. After the fault is rectified, go to **Step 20**.

- NodeAgent run logs: **/var/log/Bigdata/nodeagent/agentlog/agent.log**

- NodeAgent start and stop logs: **/var/log/Bigdata/nodeagent/scriptlog/nodeagent_ctl.log**

Step 20 Run the following commands to restart the NodeAgent service:

sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/stop-agent.sh

sh ${BIGDATA_HOME}/om-agent/nodeagent/bin/start-agent.sh

Collect fault information.

Step 21 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 22 Select the following nodes from **Services** and click **OK**.

- NodeAgent
- Controller
- OS

Step 23 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 24 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.5 ALM-12007 Process Fault

Description

This alarm is generated when the process health check module detects that the process connection status is **Bad** for three consecutive times. The process health check module checks the process status every 5 seconds.

This alarm is cleared when the process can be connected.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12007	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The service provided by the process is unavailable.

Possible Causes

- The instance process is abnormal.
- The disk space is insufficient.

NOTE

If a large number of process fault alarms exist in a time segment, files in the installation directory may be deleted mistakenly or permission on the directory may be modified.

Procedure

Check whether the instance process is abnormal.

Step 1 In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**, click ⌄ in the row where the alarm is located , and click the host name to view the host address for which the alarm is generated

Step 2 On the **Alarms** page, check whether the **ALM-12006 Node Fault** is generated.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Handle the alarm according to **ALM-12006 Node Fault**.

Step 4 Log in to the host for which the alarm is generated as user **root**. Check whether the installation directory user, user group, and permission of the alarm role are correct. The user, user group, and the permission must be **omm:ficommon 750**.

For example, the NameNode installation directory is *${BIGDATA_HOME}*/**FusionInsight_Current/***1_8_NameNode*/**etc**.

- If yes, go to **Step 6**.
- If no, go to **Step 5**.

Step 5 Run the following command to set the permission to **750** and **User:Group** to **omm:ficommon**:

chmod 750 *<folder_name>*

chown omm:ficommon *<folder_name>*

Step 6 Wait for 5 minutes. In the alarm list, check whether **ALM-12007 Process Fault** is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Step 7 Log in to the active OMS node as user **root** and run the following command to view the **configurations.xml** file. In the preceding command, "Service name" is the service name queried in **Step 1**.

vi $BIGDATA_HOME/components/current/*Service name***/configurations.xml**

Search for the keyword **healthMonitor.properties**, find the health check configuration item corresponding to the alarm reporting instance, and record the interface or script path specified by **monitor.info**, as shown in the following figure.

Check the logs recorded in the interface or script and rectify the fault.

```
<config category="healthMonitor.properties" format="propertyfileconfigurer">
    <property type="hidden" scope="setup" classification="System">
        <name>monitor.type</name>
        <value>SCRIPT</value>
    </property>
    <property type="hidden" scope="setup" classification="System">
        <name>monitor.preInitDelay</name>
        <value>120000</value>
        <!-- 2min -->
    </property>
    <property type="hidden" scope="setup" classification="System">
        <name>monitor.recheckTimes</name>
        <value>90</value>
        </property>
    <property type="hidden" scope="setup" classification="System">
        <name>monitor.checkIntervals</name>
        <value>10000</value>
    </property>
    <property type="hidden" scope="setup" classification="System">
        <name>monitor.info</name>
        <value>#{install_home}/sbin/dbservice_mon.sh</value>
    </property>
    <property type="hidden" scope="setup" classification="System">
        <name>monitor.metric.connector</name>
        <value>#{install_home}/sbin/dbservice_ha.sh</value>
    </property>
</config>
```

Step 8 Wait for 5 minutes. In the alarm list, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether disk space is sufficient.

Step 9 On the FusionInsight Manager, check whether the alarm list contains **ALM-12017 Insufficient Disk Capacity**.

- If yes, go to **Step 10**.
- If no, go to **Step 13**.

Step 10 Rectify the fault by following the steps provided in **ALM-12017 Insufficient Disk Capacity**.

Step 11 Wait for 5 minutes. In the alarm list, check whether **ALM-12017 Insufficient Disk Capacity** is cleared.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 Wait for 5 minutes. In the alarm list, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collect fault information.

Step 13 On the FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 14 According to the service name obtained in **Step 1**, select the component and **NodeAgent** from the **Service** and click **OK**.

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.6 ALM-12010 Manager Heartbeat Interruption Between the Active and Standby Nodes

Description

This alarm is generated when the active Mager does not receive the heartbeat signal from the standby Manager within 7 seconds.

This alarm is cleared when the active Manager receives heartbeat signals from the standby Manager.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12010	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

When the active Manager process is abnormal, an active/standby failover cannot be performed, and services are affected.

Possible Causes

- The link between the active and standby Manager is abnormal.
- The node name configuration is incorrect.
- The port is disabled by the firewall.

Procedure

Check whether the network between the active and standby Manager server is normal.

Step 1 In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**, click ⌄ in the row containing the alarm and view the IP address of the standby Manager (Peer Manager) server in the alarm details.

Step 2 Log in to the active Manager server as user **root**.

Step 3 Run the **ping** *standby Manager heartbeat IP address* command to check whether the standby Manager server is reachable.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.

- If no, go to **Step 6**.

Step 6 Run the following command to go to the software installation directory:

cd /opt

Step 7 Run the following command to find the configuration file directory of the active and standby nodes.

find -name hacom_local.xml

Step 8 Run the following command to go to the **workspace** directory:

cd${BIGDATA_HOME}/om-server/OMS/workspace0/ha/local/hacom/conf/

Step 9 Run the **vim** command to open the **hacom_local.xml** file. Check whether the local and peer nodes are correctly configured. The local node is configured as the active node, and the peer node is configured as the standby node.

- If yes, go to **Step 12**.
- If no, go to **Step 10**.

Step 10 Modify the configuration of the active and standby nodes in the **hacom_local.xml** file and press **Esc** to return to the command mode. Run the **:wq** command to save the modification and exit.

Step 11 Check whether the alarm is cleared automatically.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check whether the port is disabled by the firewall.

Step 12 Run the **lsof -i :20012** command to check whether the heartbeat ports of the active and standby nodes are enabled. If the command output is displayed, the ports are enabled. Otherwise, the ports are disabled by the firewall.

- If yes, go to **Step 13**.
- If no, go to **Step 16**.

Step 13 Run the **iptables -P INPUT ACCEPT** command to avoid the server disconnection.

Step 14 Run the following command to clear the firewall:

iptables -F

Step 15 Check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 16**.

Collect fault information.

Step 16 On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 17 Select the following nodes from the **Service** and click **OK**:

- OmmServer
- Controller
- NodeAgent

Step 18 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.7 ALM-12011 Manager Data Synchronization Exception Between the Active and Standby Nodes

Description

The system checks data synchronization between the active and standby Manager nodes every 60 seconds. This alarm is generated when the standby Manager fails to synchronize files with the active Manager.

This alarm is cleared when the standby Manager synchronizes files with the active Manager.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12011	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Some configurations will be lost after an active/standby switchover because the configuration files on the standby Manager are not updated. Maybe Manager and some components cannot run properly.

Possible Causes

- The link between the active and standby Managers is interrupted or The storage space of the **/srv/BigData/LocalBackup** directory is full.
- The synchronization file does not exist or the file permission is incorrect.

Procedure

Check whether the network between the active Manager server and the standby Manager server is normal.

Step 1 In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**, click ⌄ in the row where the alarm is located and obtain the standby Manager server IP address (Peer Manager IP address) in the alarm details.

Step 2 Log in to the active Manager server as user **root**.

Step 3 Run the **ping** *standby Manager IP address* command to check whether the standby Manager server is reachable.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the storage space of the /srv/BigData/LocalBackup directory is full.

Step 6 Run the following command to check whether the storage space of the **/srv/BigData/LocalBackup** directory is full:

df -hl /srv/BigData/LocalBackup

- If yes, go to **Step 7**.
- If no, go to **Step 10**.

Step 7 Run the following command to clear unnecessary backup files:

rm -rf *Directory to be cleared*

Example:

rm -rf /srv/BigData/LocalBackup/0/default-oms_20191211143443

Step 8 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

In the **Operation** column of the backup task to be performed, click **Configure** and change the value of **Maximum Number of Backup Copies** to reduce the number of backup file sets.

Step 9 Wait about 1 minute and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check whether the synchronization file exists and whether the file permission is normal.

Step 10 Run the following command to check whether the synchronization file exists.

find /srv/BigData/ -name "sed*"

find /opt -name "sed*"

- If yes, go to **Step 11**.
- If no, go to **Step 12**.

Step 11 Run the following command to view the synchronization file information and permission obtained in **Step 10**.

ll *path of the file to be found*

- If the size of the file is 0 and the permission column is **-**, the file is a junk file. Run the following command to delete it.

  rm -rf *files to be deleted*

  Wait for several minutes and check whether the alarm is cleared. If the alarm persists, go to **Step 12**.
- If the file size is not 0, go to **Step 12**.

Step 12 View the log files generated when the alarm is generated.

1. Run the following command to switch to the HA run log file path.

   cd /var/log/Bigdata/omm/oms/ha/runlog/

2. Decompress and view the log files generated when the alarm is generated.

   For example, if the name of the file to be viewed is **ha.log.2021-03-22_12-00-07.gz**, run the following command:

   gunzip *ha.log.2021-03-22_12-00-07.gz*

   vi *ha.log.2021-03-22_12-00-07*

   Check whether error information is reported before and after the alarm generation time.

   – If yes, rectify the fault based on the error information. Then go to **Step 13**.

   For example, if the following error information is displayed, the directory permission is insufficient. In this case, change the directory permission to be the same as that on the normal node.

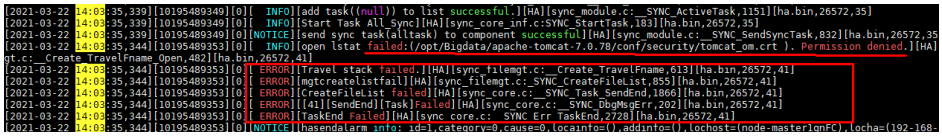

- If no, go to **Step 14**.

Step 13 Wait about 10 minute and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Collect fault information.

Step 14 On the FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 15 Select the following nodes from the **Service** and click **OK**:

- OmmServer
- Controller
- NodeAgent

Step 16 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 17 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.8 ALM-12012 NTP Service Is Abnormal

Alarm Description

The system checks whether the NTP service on a node synchronizes time with the NTP service on the active OMS node every 60 seconds. This alarm is generated when the NTP service fails to synchronize time for two consecutive times.

This alarm is generated when the time difference between the NTP service on a node and the NTP service on the active OMS node is greater than or equal to 20s for two consecutive times. This alarm is cleared when the time difference is less than 20s.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12012	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The time on the node is inconsistent with that on other nodes in the cluster. Therefore, some FusionInsight applications on the node may not run properly.

Possible Causes

- The NTP service on the current node cannot start properly.
- The current node fails to synchronize time with the NTP service on the active OMS node.
- The key authenticated by the NTP service on the current node is inconsistent with that on the active OMS node.
- The time offset between the node and the NTP service on the active OMS node is large.

Handling Procedure

Check the NTP service mode of the node.

Step 1 Log in to the active management node as user **root**, run the **su - omm** command to switch to user **omm**, and run the following command to check the resource status on the active and standby nodes:

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

- If "chrony" is displayed in the **ResName** column of the command output, go to **Step 2**.
- If "ntp" is displayed in the **ResName** column, go to **Step 20**.

NOTE

If both "chrony" and "ntp" are displayed in the **ResName** column of the command output, the NTP service mode is being switched. Wait for 10 minutes and go to **Step 1** again. If both "chrony" and "ntp" persist, contact O&M personnel personnel.

Check whether the chrony service on the node is started properly.

Step 2 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click ⌄ in the row containing the alarm, and view the name of the host for which the alarm is generated in **Location**.

Step 3 Check whether the chronyd process is running on the node where the alarm is generated. Log in to the node for which the alarm is generated as user **root** and run the **ps -ef | grep chronyd | grep -v grep** command to check whether the command output contains the chronyd process.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Start the NTP service.

Step 5 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the current node can synchronize time properly with the chrony service on the active OMS node.

Step 6 Check whether the node can synchronize time with the NTP service on the active OMS node based on additional information of the alarm.

- If yes, go to **Step 7**.
- If no, go to **Step 17**.

Step 7 Check whether the synchronization with the chrony service on the active OMS node is faulty.

Log in to the node for which the alarm is generated as user **root** and run the **chronyc sources** command.

In the command output, if there is an asterisk (*) before the IP address of the chrony service on the active OMS node, the synchronization is normal. The command output is as follows:

```
MS Name/IP address         Stratum Poll Reach LastRx Last sample
===============================================================================
^* 10.10.10.162                10  10   377   626    +16us[  +15us] +/-  308us
```

In the command output, if there is no asterisk (*) before the IP address of the NTP service on the active OMS node, and the value of **Reach** is **0**, the synchronization is abnormal.

```
MS Name/IP address         Stratum Poll Reach LastRx Last sample
===============================================================================
^? 10.1.1.1                     0  10     0     -     +0ns[   +0ns] +/-    0ns
```

- If yes, go to **Step 8**.
- If no, go to **Step 38**.

Step 8 The chrony synchronization failure is typically caused by the system firewall. If the firewall can be disabled, disable it. If the firewall cannot be disabled, check the firewall configuration policy and ensure that UDP ports 123 and 323 are not disabled. (For details, see the firewall configuration policy of each system.)

Step 9 Check whether the alarm is cleared 10 minutes later.

- If yes, no further action is required.

- If no, go to **Step 10**.

Step 10 Log in to the active OMS node as user **root** and run the following command to view the authentication code whose key index is **1M**:

In Red Hat Enterprise Linux, run the **cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys** command.

Step 11 Run the following command to check whether the key is the same as that queried in **Step 10**:

In Red Hat Enterprise Linux, run the **diff ${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys /etc/chrony.keys** command.

NOTE

If the keys are the same, no result is returned after the command is executed. For example:

```
host01:~ # cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys
1 M sdYbq;o^CzEAWo<U=Tw5
host01:~ # diff ${BIGDATA_HOME}/om-server/OMS/workspace/conf/chrony.keys /etc/chrony.keys
host01:~ #
```

- If yes, go to **Step 12**.
- If no, go to **Step 38**.

Step 12 Run the **cat ${BIGDATA_HOME}/om-server/om/packaged-distributables/ntpKeyFile** command to check whether the key is the same as that queried in **Step 10**. (Compare the key with that of the authentication key index **1M** queried in **Step 10**.)

- If yes, go to **Step 13**.
- If no, go to **Step 15**.

Step 13 Log in to the faulty node as user **root** and run the **cat /etc/chrony.keys** command in Red Hat Enterprise Linux to check whether the key is the same as that queried in **Step 12** (compare it with that of the authentication key index **1M** for comparison).

- If yes, go to **Step 38**.
- If no, go to **Step 14**.

Step 14 Run the **su - omm** command to switch to user **omm**, change the key of the authentication key index **1M** in **${NODE_AGENT_HOME}/chrony.keys** to the key of **ntpKeyFile** in **Step 12**, and go to **Step 16**.

Step 15 Run the following commands as user **root** or **omm** to change the NTP key of the active OMS node (change **ntp.keys** to **ntpkeys** in Red Hat Enterprise Linux):

cd ${BIGDATA_HOME}/om-server/OMS/workspace/conf

sed -i "`cat chrony.keys | grep -n '1 M'|awk -F ':' '{print $1}'`d" chrony.keys

echo "1 M `cat ${BIGDATA_HOME}/om-server/om/packaged-distributables/ntpKeyFile`" >> chrony.keys

Check whether the key of the authentication key index **1M** in **chrony.keys** is the same as that of **ntpKeyFile**.

- If yes, go to **Step 16**.

- If no, change the key of the authentication key index **1M** in **chrony.keys** to the key of **ntpKeyFile** and go to **Step 16**.

Step 16 After 5 minutes, run the **systemctl restart chronyd** command to restart the chrony service on the active OMS node. After 15 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 38**.

Check whether the time deviation between the node and the chrony service on the active OMS node is large.

Step 17 Check whether the time deviation is large in additional information of the alarm.

- If yes, go to **Step 18**.
- If no, go to **Step 38**.

Step 18 On the **Hosts** tab page, select the host for which the alarm is generated, and choose **More** > **Stop All Instances** to stop all the services on the node.

If the time on the alarm node is later than that on the chrony service of the active OMS node, adjust the time of the alarm node. After adjusting the time, choose **More** > **Start All Instances** to start the services on the node.

If the time on the alarm node is earlier than that on the chrony service of the active OMS node, wait until the time deviation is due and adjust the time of the alarm node. After adjusting the time, choose **More** > **Start All Instances** to start the services on the node.

NOTE

If you do not wait, data loss may occur.

Step 19 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 38**.

Check whether the NTP service on the node is started properly.

Step 20 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click ⌄ in the row containing the alarm, and view the name of the host for which the alarm is generated in **Location**.

Step 21 Check whether the ntpd process is running on the node using the following method. Log in to the alarm node as user **root** and run the **ps -ef | grep ntpd | grep -v grep** command to check whether the command output contains the ntpd process.

- If yes, go to **Step 24**.
- If no, go to **Step 22**.

Step 22 Run the **service ntp start** command (or the **service ntpd start** command in Red Hat Enterprise Linux) to start the NTP service.

Step 23 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 24**.

Check whether the node can synchronize time properly with the NTP service on the active OMS node.

Step 24 Check whether the node can synchronize time with the NTP service on the active OMS node based on additional information of the alarm.

- If yes, go to **Step 25**.
- If no, go to **Step 35**.

Step 25 Check whether the synchronization with the NTP service on the active OMS node is faulty.

Log in to the alarm node as user **root** and run the **ntpq -np** command.

If an asterisk (*) exists before the IP address of the NTP service on the active OMS node in the command output, the synchronization is in normal state. The command output is as follows:

```
remote refid st t when poll reach delay offset jitter
==============================================================================
*10.10.10.162 .LOCL. 1 u 1 16 377 0.270 -1.562 0.014
```

If there is no asterisk (*) before the IP address of the NTP service on the active OMS node, as shown in the following command output, and the value of **refid** is **.INIT.**, the synchronization is abnormal.

```
remote refid st t when poll reach delay offset jitter
==============================================================================
10.10.10.162 .INIT. 1 u 1 16 377 0.270 -1.562 0.014
```

- If yes, go to **Step 26**.
- If no, go to **Step 38**.

Step 26 The NTP synchronization failure is typically caused by the system firewall. If the firewall can be disabled, run the **iptables -F** command to disable it. If the firewall cannot be disabled, run the **iptables -L** command to check the firewall configuration policy and ensure that the UDP port 123 is not disabled. (For details, see the firewall configuration policy of each system.)

Step 27 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 28**.

Step 28 Log in to the active OMS node as user **root** and run the following command to view the authentication key index **1M**:

In SUSE Linux, run the **cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys** command.

In Red Hat Enterprise Linux or EulerOS, run the **cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntpkeys** command.

Step 29 Run the following command to check whether the key is the same as that queried in **Step 28**:

In SUSE Linux, run the **diff ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys /etc/ntp.keys** command.

In Red Hat Enterprise Linux or EulerOS, run the **diff ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntpkeys /etc/ntp/ntpkeys** command.

NOTE

If the keys are the same, no result is returned after the command is executed. For example:

```
host01:~ # cat ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys
1 M sdYbq;o^CzEAWo<U=Tw5
host01:~ # diff ${BIGDATA_HOME}/om-server/OMS/workspace/conf/ntp.keys /etc/ntp.keys
host01:~ #
```

- If yes, go to **Step 30**.
- If no, go to **Step 38**.

Step 30 Run the **cat ${BIGDATA_HOME}/om-server/om/packaged-distributables/ntpKeyFile** command to check whether the key is the same as that queried in **Step 28**. (Compare the key with that of the authentication key index **1M** queried in **Step 28**.)

- If yes, go to **Step 31**.
- If no, go to **Step 33**.

Step 31 Log in to the faulty node as user **root** and run the **cat /etc/ntp.keys** command in SUSE Linux (or the **cat /etc/ntp/ntpkeys** command in Red Hat Enterprise Linux) to check whether the key is the same as the value queried in **Step 30** (use the key of the authentication key index **1M** for comparison).

- If yes, go to **Step 38**.
- If no, go to **Step 32**.

Step 32 Run the **su - omm** command to switch to user **omm**, change the key of the authentication key index **1M** in **${NODE_AGENT_HOME}/ntp.keys** (**${NODE_AGENT_HOME}/ntpkeys** in Red Hat Enterprise Linux) to the key of **ntpKeyFile** in **Step 30**, and go to **Step 34**.

Step 33 Run the following commands as user **root** or **omm** to change the NTP key of the active OMS node (change **ntp.keys** to **ntpkeys** in Red Hat Enterprise Linux):

cd ${BIGDATA_HOME}/om-server/OMS/workspace/conf

sed -i "`cat ntp.keys | grep -n '1 M'|awk -F ':' '{print $1}'`d" ntp.keys

echo "1 M `cat ${BIGDATA_HOME}/om-server/om/packaged-distributables/ntpKeyFile`" >>ntp.keys

Check whether the key of the authentication key index **1M** in **ntp.keys** is the same as that of **ntpKeyFile**.

- If yes, go to **Step 34**.
- If no, change the key of the authentication key index **1M** in **ntp.keys** to the key of **ntpKeyFile** and go to **Step 34**.

Step 34 After 5 minutes, run the **service ntp restart** command to restart the NTP service on the active OMS node. After 15 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 38**.

Check whether the time deviation between the node and the NTP service on the active OMS node is large.

Step 35 Check whether the time deviation is large in additional information of the alarm.

- If yes, go to **Step 36**.
- If no, go to **Step 38**.

Step 36 On the **Hosts** tab page, select the host for which the alarm is generated, and choose **More** > **Stop All Instances** to stop all the services on the node.

If the time on the alarm node is later than that on the NTP service of the active OMS node, adjust the time of the alarm node. After adjusting the time, choose **More** > **Start All Instances** to start the services on the node.

If the time on the alarm node is earlier than that on the NTP service of the active OMS node, wait until the time deviation is due and adjust the time of the alarm node. After adjusting the time, choose **More** > **Start All Instances** to start the services on the node.

NOTE

If you do not wait, data loss may occur.

Step 37 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 38**.

Collect fault information.

Step 38 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 39 Expand the **Service** drop-down list, select **NodeAgent** and **OmmServer** for the target cluster, and click **OK**. Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

Step 40 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 41 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.9 ALM-12014 Partition Lost

Description

The system checks the partition status every 60 seconds. This alarm is generated when the system detects that a partition to which service directories are mounted

is lost (because the device is removed or goes offline, or the partition is deleted). The system checks the partition status periodically.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12014	Major	• Yes: MRS 3.3.0 and later versions • No: Versions earlier than MRS 3.3.0

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DirName	Specifies the directory for which the alarm is generated.
PartitionName	Specifies the device partition for which the alarm is generated.

Impact on the System

Service data fails to be written into the partition, and the service system runs abnormally.

Possible Causes

- The hard disk is removed.
- The hard disk is offline, or a bad sector exists on the hard disk.

Procedure

Step 1 On FusionInsight Manager, click **O&M > Alarm > Alarms**, and click ⌄ in the row where the alarm is located.

Step 2 Obtain **HostName**, **PartitionName** and **DirName** from **Location**.

Step 3 Check whether the disk of **PartitionName** on **HostName** is inserted to the correct server slot.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Contact hardware engineers to remove the faulty disk.

Step 5 Log in to the **HostName** node where an alarm is reported and check whether there is a line containing **DirName** in the **/etc/fstab** file as user **root**.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Run the **vi /etc/fstab** command to edit the file and delete the line containing **DirName**.

Step 7 Contact hardware engineers to insert a new disk. For details, see the hardware product document of the relevant model. If the faulty disk is in a RAID group, configure the RAID group. For details, see the configuration methods of the relevant RAID controller card.

Step 8 Wait 20 to 30 minutes (The disk size determines the waiting time), and run the **mount** command to check whether the disk has been mounted to the **DirName** directory.

- If yes, go to **Step 9** for MRS 3.3.0 and later, MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions. For clusters earlier than MRS 3.3.0, manually clear the alarm. No further action is required.
- If no, go to **Step 10**.

Step 9 Wait about 2 minute and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On the FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 11 Select the **OmmServer** from the Services drop-down list and click **OK**.

Step 12 Set Start Date for log collection to 10 minutes ahead of the alarm generation time and End Date to 10 minutes behind the alarm generation time and click **Download**.

Step 13 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

MRS 3.3.0 and later, MRS 3.1.0.0.10/3.1.5.0.3 and later patch versions: After the fault is rectified, the system automatically clears this alarm.

Versions earlier than MRS 3.3.0: After the fault is rectified, the system does not automatically clear this alarm, and you need to manually clear the alarm.

Related Information

None

10.10 ALM-12015 Partition Filesystem Readonly

Description

The system checks the partition status every 60 seconds. This alarm is generated when the system detects that a partition to which service directories are mounted enters the read-only mode (due to a bad sector or a faulty file system). The system checks the partition status periodically.

This alarm is cleared when the system detects that the partition to which service directories are mounted exits from the read-only mode (because the file system is restored to read/write mode, the device is removed, or the device is formatted).

Attribute

Alarm ID	Alarm Severity	Auto Clear
12015	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DirName	Specifies the directory for which the alarm is generated.
PartitionName	Specifies the device partition for which the alarm is generated.

Impact on the System

Service data fails to be written into the partition, and the service system runs abnormally.

Possible Causes

The hard disk is faulty, for example, a bad sector exists.

Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, click ⌄ in the row where the alarm is located.

Step 2 Obtain **HostName** and **PartitionName** from **Location**. **HostName** is the node where the alarm is reported, and **PartitionName** is the partition of the faulty disk.

Step 3 Contact hardware engineers to check whether the disk is faulty. If the disk is faulty, remove it from the server.

Step 4 After the disk is removed, alarm **ALM-12014 Partition Lost** is reported. Handle the alarm. For details, see **ALM-12014 Partition Lost**. After the alarm **ALM-12014 Partition Lost** is cleared, alarm **ALM-12015 Partition Filesystem Readonly** is automatically cleared.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.11 ALM-12016 CPU Usage Exceeds the Threshold

Description

The system checks the CPU usage every 30 seconds and compares the actual CPU usage with the threshold. The CPU usage has a default threshold. This alarm is generated when the CPU usage exceeds the threshold for several times (configurable, 10 times by default) consecutively.

The alarm is cleared in the following two scenarios: The value of **Trigger Count** is 1 and the CPU usage is smaller than or equal to the threshold; the value of **Trigger Count** is greater than 1 and the CPU usage is smaller than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12016	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Service processes respond slowly or become unavailable.

Possible Causes

- The alarm threshold or alarm smoothing times are incorrect.
- CPU configuration cannot meet service requirements. The CPU usage reaches the upper limit. Or the service is in peak hours. As a result, the CPU usage reaches the upper limit in a short period of time.

Procedure

Check whether the alarm threshold or alarm Trigger Count are correct.

Step 1 Change the alarm threshold and alarm **Trigger Count** based on CPU usage.

On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **CPU** > **Host CPU Usage** and change the alarm smoothing times based on CPU usage, as shown in **Figure 10-1**.

NOTE

This option defines the alarm check phase. **Trigger Count** indicates the alarm check threshold. An alarm is generated when the number of check times exceeds the threshold.

Figure 10-1 Setting alarm smoothing times

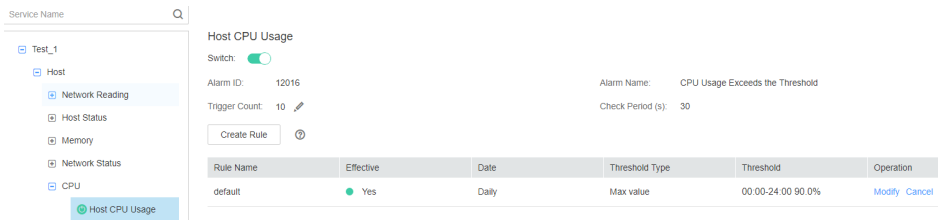

On **Host CPU Usage** page and click **Modify** in the **Operation** column to change the alarm threshold, as shown in **Figure 10-2**.

Figure 10-2 Setting an alarm threshold

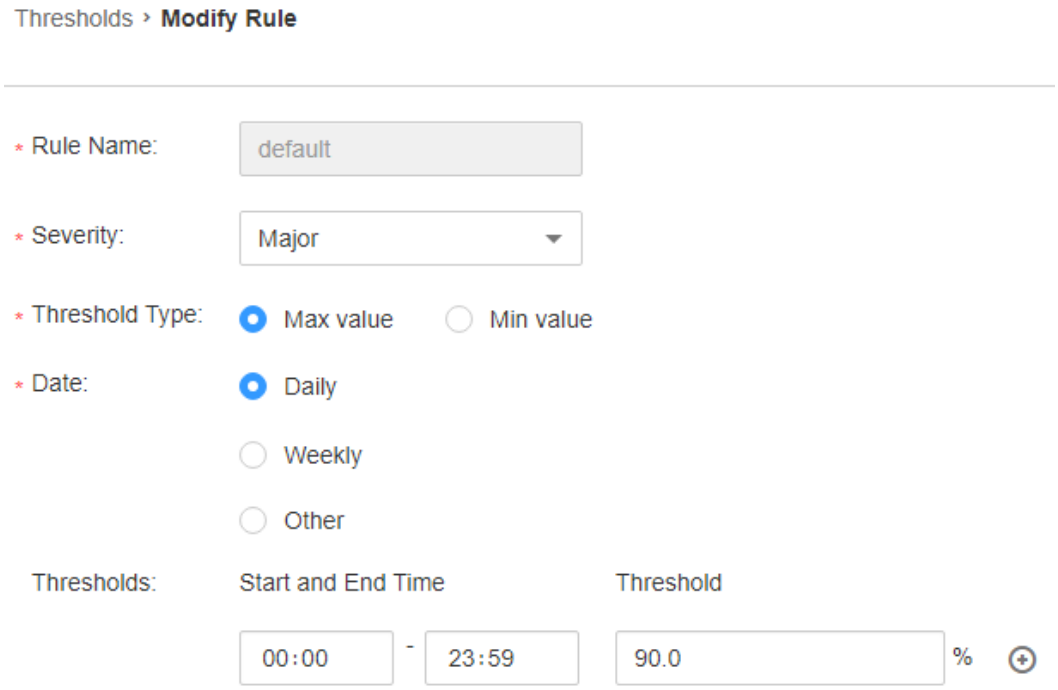

Step 2 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Check whether the CPU usage reaches the upper limit.

Step 3 In the alarm list on FusionInsight Manager, click ∨ in the row where the alarm is located to view the alarm host address in the alarm details.

Step 4 On the **Hosts** page, click the node on which the alarm is reported.

Step 5 View the CPU usage for 5 minutes. If the CPU usage exceeds the threshold for multiple times, contact the system administrator to add more CPUs.

Step 6 Check whether the current traffic is in peak hours. If the alarm is generated during peak hours, you are advised to expand the capacity of the node or contact the system administrator to add more CPUs.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager in the active cluster, choose **O&M** > **Log** > **Download**.

Step 9 Select **OmmServer** from the **Service** and click **OK**.

Step 10 Set **Start Date** for log collection to 10 minutes ahead of the alarm generation time and **End Date** to 10 minutes behind the alarm generation time in **Time Range** and click **Download**.

Step 11 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.12 ALM-12017 Insufficient Disk Capacity

Description

The system checks the host disk usage of the system every 30 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold, this alarm is generated when the host disk usage exceeds the specified threshold.

When the **Trigger Count** is 1, this alarm is cleared when the usage of a host disk partition is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the usage of a host disk partition is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12017	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
PartitionName	Specifies the device partition for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Service processes become unavailable.

Possible Causes

- The alarm threshold is incorrect.
- Disk configuration of the server cannot meet service requirements.

Procedure

Check whether the alarm threshold is appropriate.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm > Thresholds >** *Name of the desired cluster* > **Host** > **Disk** > **Disk Usage** and check whether the threshold (configurable, 90% by default) is appropriate.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Choose **O&M** > **Alarm > Thresholds >** *Name of the desired cluster* > **Host** > **Disk** > **Disk Usage** and click **Modify** in the **Operation** column to change the alarm threshold based on site requirements. As shown in **Figure 10-3**:

Figure 10-3 Setting an alarm threshold

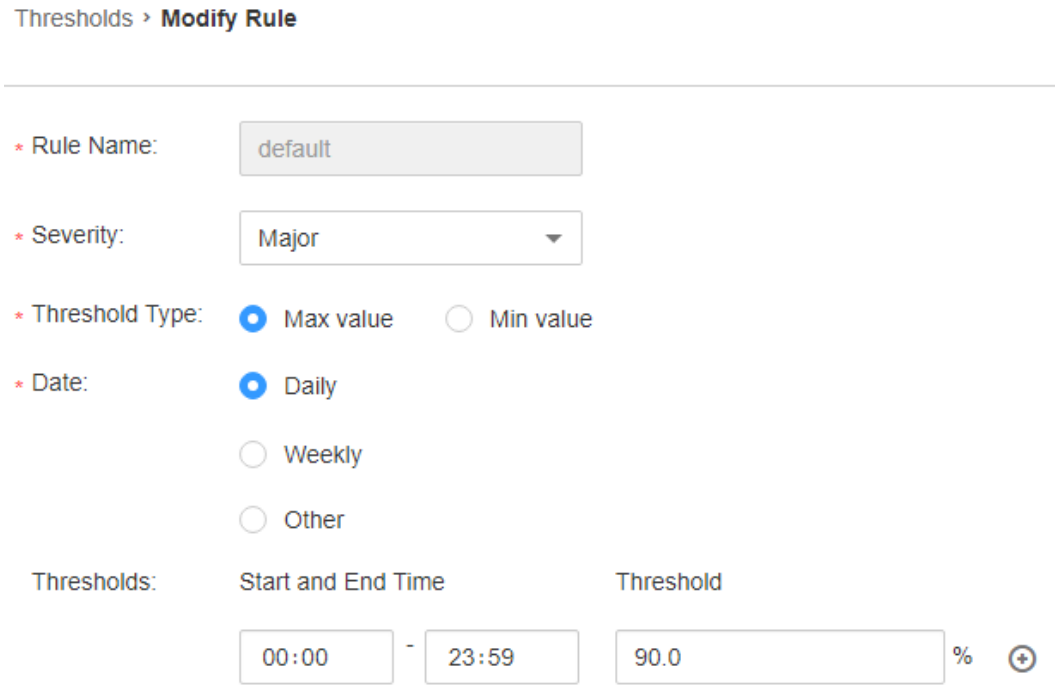

Step 3 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the disk usage reaches the upper limit.

Step 4 In the alarm list on FusionInsight Manager, click ⌄ in the row where the alarm is located to view the alarm host name and disk partition information in the alarm details.

Step 5 Log in to the node where the alarm is generated as user **root**.

Step 6 Run the **df -lmPT | awk '$2 != "iso9660"' | grep '^/dev/' | awk '{"readlink -m "$1 | getline real }{$1=real; print $0}' | sort -u -k 1,1** command to check the system disk partition usage. Check whether the disk is mounted to the following directories based on the disk partition name obtained in **Step 4**: **/**, **/opt**, **/tmp**, **/var**, **/var/log**, and **/srv/BigData**(can be customized).

- If yes, the disk is a system disk. Then go to **Step 10**.
- If no, the disk is not a system disk. Then go to **Step 7**.

Step 7 Run the **df -lmPT | awk '$2 != "iso9660"' | grep '^/dev/' | awk '{"readlink -m "$1 | getline real }{$1=real; print $0}' | sort -u -k 1,1** command to check the system disk partition usage. Determine the role of the disk based on the disk partition name obtained in **Step 4**.

Step 8 Check the disk service.

In MRS, check whether the disk service is HDFS, Yarn, Kafka, Supervisor.

- If yes, adjust the capacity. Then go to **Step 9**.
- If no, go to **Step 12**.

Step 9 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Step 10 Run the **find / -xdev -size +500M -execls -l {} \;** command to check whether a file larger than 500 MB exists on the node and disk.

- If yes, go to **Step 11**.
- If no, go to **Step 12**.

Step 11 Handle the large file and check whether the alarm is cleared 2 minutes later.

- If yes, no further action is required.
- If no, go to **Step 12**.

Step 12 Contact the system administrator to expand the disk capacity.

Step 13 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Collect fault information.

Step 14 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 15 Select **OMS** from the **Service** and click **OK**.

Step 16 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 17 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.13 ALM-12018 Memory Usage Exceeds the Threshold

Description

The system checks the memory usage of the system every 30 seconds and compares the actual memory usage with the threshold. The memory usage has a default threshold, this alarm is generated when the value of the memory usage exceeds the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the host memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1,

this alarm is cleared when the host memory usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12018	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Service processes respond slowly or become unavailable.

Possible Causes

Memory configuration cannot meet service requirements. The memory usage reaches the upper limit.

Procedure

Expand the system.

Step 1 In the alarm list on FusionInsight Manager, click ⌄ in the row where the alarm is located to view the alarm host address in the alarm details.

Step 2 Log in to the host where the alarm is generated as user **root**.

Step 3 If the memory usage exceeds the threshold, perform memory capacity expansion.

Step 4 Run the command **free -m | grep Mem\: | awk '{printf("%s,", $3 * 100 / $2)}'** to check the system memory usage.

Step 5 Wait for 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager in the active cluster, choose **O&M** > **Log** > **Download**.

Step 7 Select **OmmServer** from the **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.14 ALM-12027 Host PID Usage Exceeds the Threshold

Description

The system checks the PID usage every 30 seconds and compares the actual PID usage with the default PID usage threshold. This alarm is generated when the system detects that the PID usage exceeds the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the PID usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the PID usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12027	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

No PID is available for new processes and service processes are unavailable.

Possible Causes

Too many processes are running on the node. You need to increase the value of **pid_max**.

Procedure

Increase the value of pid_max.

Step 1 In the alarm list on FusionInsight Manager, click ⌄ in the row where the alarm is located to view the alarm host address in the alarm details.

Step 2 Log in to the host where the alarm is generated as user **root**.

Step 3 Run the **cat /proc/sys/kernel/pid_max**command to check the value of **pid_max**.

Step 4 If the PID usage exceeds the threshold, edit the **/etc/sysctl.conf** file and increase the value of **kernel.pid_max** to twice of the value of **pid_max** queried in **Step 3**. If this parameter does not exist, add it to the end of the file.

For example, change the parameter to **kernel.pid_max=65536** and run the following command to make the parameter take effect immediately:

sysctl -p

NOTE

The maximum value of **kernel.pid_max** is as follows:

- On 32-bit systems: 32768
- On 64-bit systems: 4194304 (2^22)

Step 5 Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log** > **Download**.

Step 7 Select all services from the **Service** and click **OK**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.15 ALM-12028 Number of Processes in the D State and Z State on a Host Exceeds the Threshold

Alarm Description

The system checks the number of processes in the D state and Z state of user **omm** on the host every 30 seconds and compares the actual number with the threshold. The number of processes in the D state and Z state on the host has a default threshold range. This alarm is generated when the number of processes exceeds the threshold.

This alarm is cleared when **Trigger Count** is **1** and the total number of processes in the D state and Z state of user **omm** on the host does not exceed the threshold. This alarm is cleared when **Trigger Count** is greater than **1** and the total number of processes in the D state and Z state of user **omm** on the host is less than or equal to 90% of the threshold.

NOTE

The function of checking the number of processes in the Z state on the host applies to MRS 3.2.0 or later.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12028	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Excessive system resources are used and service processes respond slowly.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state.

Handling Procedure

Check the processes in the D state and Z state.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the IP address of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**. () Then run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command as user **omm** to view the PID of the process that is in the D state and Z state:

ps -elf | grep -v "\[thread_checkio\]" | awk 'NR!=1 {print $2, $3, $4}' | grep omm | awk -F' ' '{print $1, $3}' | grep -E "Z|D" | awk '{print $2}'

Step 4 Check whether the command output is empty.

- If yes, the service process is running properly. Then go to **Step 6**.

- If no, go to **Step 5**.

Step 5 Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

Step 6 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Select **OMS** for **Service** and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.16 ALM-12033 Slow Disk Fault

Alarm Description

For MRS 3.3.0 and its later versions:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 150 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 1000 ms within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 20 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when neither of the preceding conditions is met for three consecutive detection periods (30 or 300 seconds).

For versions earlier than MRS 3.3.0:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - The system runs the **iostat** command every 3 seconds by default, and detects that the **svctm** value exceeds 1000 ms in at least seven consecutive periods within 30 seconds.
  - The system runs the **iostat** command every 3 seconds, and detects that more than 50% of I/Os take more than 150 ms within 300s.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - The system runs the **iostat** command every 3 seconds by default, and detects that the **svctm** value exceeds 1000 ms for at least 10 periods within 30 seconds.
  - The system runs the **iostat** command every 3 seconds by default, and detects that more than 60% of I/Os take more than 20 ms within 300s.

This alarm is automatically cleared when the preceding conditions have not been met for 15 minutes.

NOTE

For details about how to obtain the **svctm** value, see **Related Information**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12033	• Minor: MRS 3.3.0 and its later versions • Major: versions earlier than MRS 3.3.0	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
DiskName	Specifies the disk for which the alarm was generated.

Impact on the System

Service performance deteriorates, service processing capabilities become poor, and services may be unavailable.

Possible Causes

The disk is aged or has bad sectors.

Handling Procedure

Check the disk status.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**.

Step 2 View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is generated.

Step 3 Check whether the node for which the alarm is generated is in a virtualization environment.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 Check whether the storage performance provided by the virtualization environment meets the hardware requirements. Then, go to **Step 5**.

Step 5 Log in to the alarm node as user **root**, run the **df -h** command, and check whether the command output contains the value of the **DiskName** field.

- If yes, go to **Step 7**.
- If no, go to **Step 6**.

Step 6 Run the **lsblk** command to check whether the mapping between the value of **DiskName** and the disk has been created.

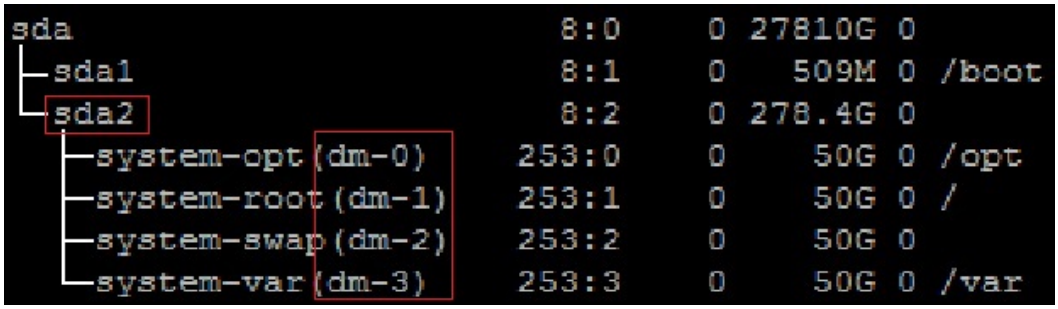

- If yes, go to **Step 7**.
- If no, go to **Step 22**.

Step 7 Log in to the alarm node as user **root**, run the **lsscsi | grep "/dev/sd[x]"** command to view the disk information, and check whether RAID has been set up.

NOTE

In the command, **/dev/sd[x]** indicates the disk name obtained in **Step 2**.

Example:

lsscsi | grep "/dev/sda"

In the command output, if **ATA**, **SATA**, or **SAS** is displayed in the third line, the disk has not been organized into a RAID group. If other information is displayed, RAID has been set up.

- If yes, go to **Step 12**.
- If no, go to **Step 8**.

Step 8 Run the **smartctl -i /dev/sd[x]** command to check whether the hardware supports the SMART tool.

Example:

smartctl -i /dev/sda

In the command output, if "SMART support is: Enabled" is displayed, the hardware supports SMART. If "Device does not support SMART" or other information is displayed, the hardware does not support SMART.

- If yes, go to **Step 9**.
- If no, go to **Step 16**.

Step 9 Run the **smartctl -H --all /dev/sd[x]** command to check basic SMART information and determine whether the disk is working properly.

Example:

smartctl -H --all /dev/sda

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or **Elements in grown defect list**. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to **Step 10**.
- If no, go to **Step 18**.

Step 10 Run the **smartctl -l error -H /dev/sd[x]** command to check the Glist of the disk and determine whether the disk is normal.

Example:

smartctl -l error -H /dev/sda

Check the **Command/Feature_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can trigger the disk SMART self-check.

- If yes, go to **Step 11**.
- If no, go to **Step 18**.

Step 11 Run the **smartctl -t long /dev/sd[x]** command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be

completed is displayed. After the self-check is completed, repeat **Step 9** and **Step 10** to check whether the disk is working properly.

Example:

smartctl -t long /dev/sda

- If yes, go to **Step 17**.
- If no, go to **Step 18**.

Step 12 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]** command to check whether the hardware supports SMART.

NOTE

- In the command, **[sat|scsi]** indicates the disk type. Both types need to be used.
- **[DID]** indicates the slot information. Slots 0 to 15 need to be used.

For example, run the following commands in sequence:

smartctl -d sat+megaraid,0 -H --all /dev/sda

smartctl -d sat+megaraid,1 -H --all /dev/sda

smartctl -d sat+megaraid,2 -H --all /dev/sda

...

Try the command combinations of different disk types and slot information. If "SMART support is: Enabled" is displayed in the command output, the disk supports SMART. Record the parameters of the disk type and slot information when a command is successfully executed. If "SMART support is: Enabled" is not displayed in the command output, the disk does not support SMART.

- If yes, go to **Step 13**.
- If no, go to **Step 16**.

Step 13 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -H --all /dev/sd[x]** command recorded in **Step 12** to check basic SMART information and determine whether the disk is normal.

Example:

smartctl -d sat+megaraid,2 -H --all /dev/sda

Check the value of **SMART overall-health self-assessment test result** in the command output. If the value is **FAILED**, the disk is faulty and needs to be replaced. If the value is **PASSED**, check the value of **Reallocated_Sector_Ct** or **Elements in grown defect list**. If the value is greater than 100, the disk is faulty and needs to be replaced.

- If yes, go to **Step 14**.
- If no, go to **Step 18**.

Step 14 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -l error -H /dev/sd[x]** command to check the Glist of the disk and determine whether the hard disk is working properly.

Example:

smartctl -d sat+megaraid,2 -l error -H /dev/sda

Check the **Command/Feature_name** column in the command output. If **READ SECTOR(S)** or **WRITE SECTOR(S)** is displayed, the disk has bad sectors. If other errors occur, the disk circuit board is faulty. Both errors indicate that the disk is abnormal and needs to be replaced.

If "No Errors Logged" is displayed, no error log exists. You can trigger the disk SMART self-check.

- If yes, go to **Step 15**.
- If no, go to **Step 18**.

Step 15 Run the **smartctl -d [sat|scsi]+megaraid,[DID] -t long /dev/sd[x]** command to trigger the disk SMART self-check. After the command is executed, the time when the self-check is to be completed is displayed. After the self-check is completed, repeat **Step 13** and **Step 14** to check whether the disk is working properly.

Example:

smartctl -d sat+megaraid,2 -t long /dev/sda

- If yes, go to **Step 17**.
- If no, go to **Step 18**.

Step 16 If the configured RAID controller card does not support SMART, the disk does not support SMART. In this case, use the check tool provided by the corresponding RAID controller card vendor to rectify the fault. Then go to **Step 17**.

For example, LSI is a MegaCLI tool.

Step 17 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, click **Clear** in the **Operation** column of the alarm, and check whether the alarm is reported on the same disk again.

If the alarm is reported for three times, replace the disk.

- If yes, go to **Step 18**.
- If no, no further action is required.

Replace the disk.

Step 18 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**.

Step 19 View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

Step 20 Replace the disk.

Step 21 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 22**.

Collect the fault information.

Step 22 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 23 Select **OMS** for **Service** and click **OK**.

Step 24 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 25 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

The **svctm** value can be obtained as follows:

- MRS 3.1.0:

  Run the **iostat -x -t** command in the OS.

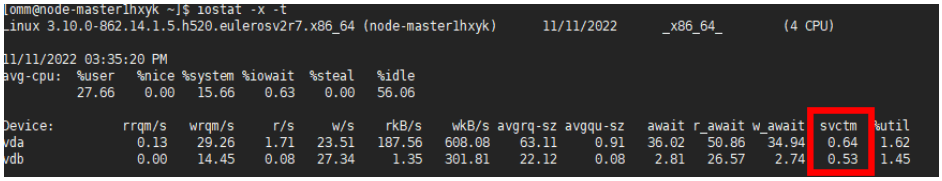

- Versions later than MRS 3.1.0:

  svctm = (tot_ticks_new - tot_ticks_old) / (rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old)

- Versions earlier than MRS 3.3.0: If **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, then **svctm = 0**.
- MRS 3.3.0 and its later versions:

  When the detection period is 30 seconds, if **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, then **svctm = 0**.

  When the detection period is 300 seconds and **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, if **tot_ticks_new - tot_ticks_old = 0**, then **svctm = 0**; otherwise, the value of **svctm** is infinite.

The parameters can be obtained as follows:

The system runs the **cat /proc/diskstats** command every 3 seconds to collect data. For example:

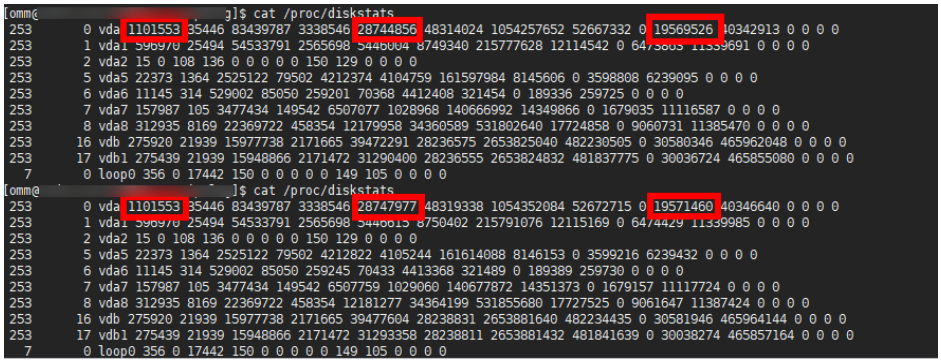

In the two commands:

In the data collected for the first time, the number in the fourth column is the **rd_ios_old** value, the number in the eighth column is the **wr_ios_old** value, and the number in the thirteenth column is the **tot_ticks_old** value.

In the data collected for the second time, the number in the fourth column is the **rd_ios_new** value, the number in the eighth column is the **wr_ios_new** value, and the number in the thirteenth column is the **tot_ticks_new** value.

In this case, the value of **svctm** is as follows:

(19571460 - 19569526) / (1101553 + 28747977 - 1101553 - 28744856) = 0.6197

10.17 ALM-12034 Periodical Backup Failure

Description

The system executes the periodic backup task every 60 minutes. This alarm is generated when a periodical backup task fails to be executed. This alarm is cleared when the next backup task is executed successfully.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12034	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TaskName	Specifies the task.

Impact on the System

There are not available backup packages for a long time, so the system cannot be restored in case of exceptions.

Possible Causes

The alarm cause depends on the task details. Handle the alarm according to the logs and alarm details.

Procedure

Check whether the disk space is sufficient.

Step 1 In the FusionInsight Manager portal, click **O&M > Alarm > Alarms**.

Step 2 In the alarm list, click ⌄ in the row where the alarm is located and obtain **TaskName** from **Location**.

Step 3 Log in to the active node of the cluster as the **root** user and check whether information similar to the following is recorded in backup and restoration logs in **/var/log/Bigdata/controller/backup/**.

```
Upload backup files to *** file failed, error info: ***
```

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On FusionInsight Manager, choose **O&M** > **Backup and Restoration > Backup Management**. Locate the backup task based on **TaskName**, click **Configure** in the **Operation** column, and check whether all configuration items are correctly configured.

- If yes, go to **Step 7**.
- If no, modify the configuration, save the modification, and go to **Step 5**.

Step 5 Choose **More > Back Up Now** to start the backup task and check whether the backup task is successfully executed.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Step 7 Choose **More > View History** in the **Operation** column. In the displayed dialog box view the task details.

Figure 10-4 View History

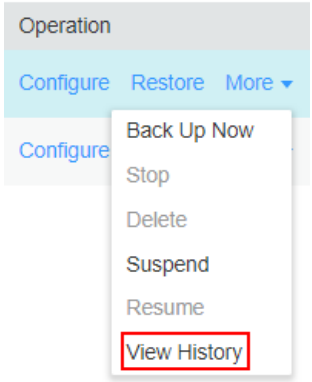

Step 8 In the displayed dialog box and click ⌄ to check whether the following message is displayed: Failed to backup xx due to insufficient disk space, move the data in the xx directory to other directories.

- If yes, go to **Step 9**.
- If no, go to **Step 16**.

Step 9 Choose **Backup Path** > **View** and obtain the **Backup Path**.

Step 10 Log in to the node as user **root** and run the following command to check the node mounting details:

df -h

Step 11 Check whether the available space of the node to which the backup path is mounted is less than 20 GB.

- If yes, go to **9**.
- If no, go to **Step 16**.

Step 12 Check whether there are many backup packages in the backup directory.

- If yes, go to **Step 13**.
- If no, go to **Step 16**.

Step 13 Enable the available space of the node to which the backup directory is mounted to be greater than 20 GB by moving backup packages out of the backup directory or delete the backup packages.

Step 14 After the problem is resolved, perform the backup task again and check whether the backup task execution is successful.

- If yes, go to **Step 15**.
- If no, go to **Step 16**.

Step 15 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Collect fault information.

Step 16 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 17 Select **Controller** from the **Service** and click **OK**.

Step 18 Click ✎ in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.18 ALM-12035 Unknown Data Status After Recovery Task Failure

Description

After the recovery task fails, the system automatically rolls back every 60 minutes. If the rollback fails, data may be lost. If this occurs, an alarm is reported. This alarm is cleared when the next recovery task execution is successful.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12035	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TaskName	Specifies the task.

Impact on the System

After the recovery task fails, the system automatically rolls back. If the rollback fails, data may be lost or the data status may be unknown, which may affect services.

Possible Causes

The alarm cause depends on the task details. Handle the alarm according to the logs and alarm details.

Procedure

Collect fault information.

Step 1 In the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the running status of the component meets the requirements. (The OMS and DBService must be in the normal state, and other components must be stopped.)

- If yes, go to **Step 9**.
- If no, go to **Step 2**.

Step 2 Restore the component status as required and start the recovery task again.

Step 3 Log in to the FusionInsight Manager portal and click **O&M** > **Alarm** > **Alarms**.

Step 4 In the alarm list, click ⌄ in the row where the alarm is located to obtain **TaskName** from **Location**.

Step 5 Choose **O&M** > **Backup and Restoration** > **Restoration Management**.

Step 6 Find the restoration task by **Task Name** and view the task details.

Step 7 Perform the recovery task again and check whether the recovery task execution is successful.

- If yes, go to **8**.
- If no, go to **9**.

Step 8 After 2 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **9**.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 10 Select **Controller** from the **Service** and click **OK**.

Step 11 Click ✎ in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.19 ALM-12037 NTP Server Abnormal

Description

The system checks the NTP server status every 60 seconds. This alarm is generated when the system detects that the NTP server is abnormal for 10 consecutive times.

This alarm is cleared when the NTP server recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12037	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the IP address of the NTP server for which the alarm is generated.

Impact on the System

The NTP server configured on the active OMS node is abnormal. In this case, the active OMS node cannot synchronize time with the NTP server and a time offset may be generated in the cluster.

Possible Causes

- The NTP server network is abnormal.
- The NTP server authentication fails.
- The NTP server time cannot be obtained.
- The time obtained from the NTP server is not continuously updated.

Procedure

Check the NTP server network.

Step 1 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and click ⌄ in the row where the alarm is located.

Step 2 View the alarm additional information to check whether the NTP server fails to be pinged.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Contact the network administrator to check the network configuration and ensure that the network between the NTP server and the active OMS node is normal. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the NTP server authentication fails.

Step 4 Log in to the active OMS node as user **root**.

Step 5 Run the following command to check the status of the resources on the active and standby nodes:

su - omm

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

- If "chrony" is displayed in the **ResName** column of the command output, go to **Step 6**.
- If "ntp" is displayed in the **ResName** column, go to **Step 7**.

NOTE

If both "chrony" and "ntp" are displayed in the **ResName** column of the command output, the NTP service mode is being switched. Wait for 10 minutes and perform **Step 5** again. If both "chrony" and "ntp" still exist in the **ResName** column, contact O&M personnel.

Step 6 Run the command **chronyc sources** to check whether the NTP server authentication fails.

If the value of **Reach** for chrony is **0**, the connection or authentication fails.

- If yes, go to **Step 12**.
- If no, go to **Step 8**.

Step 7 Run the command **ntpq -np** to check whether the NTP server authentication fails.

If **refid** of the NTP server is **.AUTH.**, the authentication fails.

- If yes, go to **Step 12**.
- If no, go to **Step 8**.

Check whether the time can be obtained from the NTP server.

Step 8 View the alarm additional information to check whether the time can be obtained from the NTP server.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 Contact the provider of the NTP server to rectify the NTP server fault. After the NTP server is normal, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check whether the time obtained from the NTP server is not continuously updated.

Step 10 View the alarm additional information to check whether the time obtained from the NTP server is not continuously updated.

- If yes, go to **Step 11**.
- If no, go to **Step 12**.

Step 11 Contact the provider of the NTP server to rectify the NTP server fault. After the NTP server is normal, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Collect fault information.

Step 12 On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 13 Select **NodeAgent** and **OmmServer** from the **Service** and click **OK**.

Step 14 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 15 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.20 ALM-12038 Monitoring Indicator Dumping Failure

Description

After monitoring indicator dumping is configured on FusionInsight Manager, the system checks the monitoring indicator dumping result at the dumping interval (60 seconds by default). This alarm is generated when the dumping fails.

This alarm is cleared when dumping is successful.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12038	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The upper-layer management system cannot obtain monitoring indicators from the FusionInsight Manager system.

Possible Causes

- The server cannot be connected.
- The save path on the server cannot be accessed.
- The monitoring indicator file fails to be uploaded.

Procedure

Check whether the server connection is normal.

Step 1 Check whether the network between the FusionInsight Manager system and the server is normal.

- If yes, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 Contact the network administrator to recover the network and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Choose **System** > **Interconnection** > **Upload Performance Data** and check whether the FTP username, password, port, dump mode, and public key

configured on the upload performance data page are consistent with the configuration on the server.

- If yes, go to **Step 5**.
- If no, go to **Step 4**.

Step 4 Enter the correct configuration information, click **OK**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the permission of the save path on the server is correct.

Step 5 Choose **System** > **Interconnection > Upload Performance Data** and check the configuration items **FTP Username**, **Save Path**, and **Dump Mode**.

- If the dump mode is FTP, go to **Step 6**.
- If the dump mode is SFTP, go to **Step 7**.

Step 6 Log in to the server in FTP mode. In the default path, check whether **FTP Username** has the read and write permission of the relative path **Save Path**.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

Step 7 Log in to the server in SFTP mode and check whether **FTP Username** has the read and write permission of the absolute path **Save Path**.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

Step 8 Add the read and write permission and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether the save path on the server has sufficient disk space.

Step 9 Log in to the server and check whether the save path has sufficient disk space.

- If yes, go to **Step 11**.
- If no, go to **Step 10**.

Step 10 Delete unnecessary files or go to the monitoring indicator dumping configuration page to change the save path. Then, check whether the save path has sufficient disk space.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 12 Select **OMS** from the **Service** and click **OK**.

Step 13 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.21 ALM-12039 Active/Standby OMS Databases Not Synchronized

Description

The system checks the data synchronization status between the active and standby OMS Databases every 10 seconds. This alarm is generated when the synchronization status cannot be queried for 30 consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the data synchronization status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12039	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local GaussDB HA IP	Specifies the HA IP address of the local GaussDB.
Peer GaussDB HA IP	Specifies the HA IP address of the peer GaussDB.

Name	Meaning
SYNC_PERCENT	Specifies the synchronization percentage.

Impact on the System

When data is not synchronized between the active and standby OMS Databases, data may be lost or abnormal if the active instance becomes abnormal.

Possible Causes

- The network between the active and standby nodes is unstable.
- The standby OMS Database is abnormal.
- The standby node disk space is full.

Procedure

Check whether the network between the active and standby nodes is normal.

Step 1 Log in to FusionInsight Manager, click **O&M > Alarm > Alarms**, click ⌄ in the row where the alarm is located, and query the standby OMS Database IP address.

Step 2 Log in to the active OMS Database node as user **root**.

Step 3 Run the **ping** *Standby OMS Database heartbeat IP address* command to check whether the standby OMS Database node is reachable.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the standby OMS Database is normal. (Skip this check for versions later than MRS 3.1.2.)

Step 6 Log in to the standby OMS Database node as user **root**.

Step 7 Run the **su - omm** command to switch to user **omm**.

Step 8 Go to the **${BIGDATA_HOME}/om-server/om/sbin/** directory and run the **./status-oms.sh** command to check whether the OMS Database resource status of the standby DBService is normal. In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:

For example:

```
10_10_10_231 gaussDB Standby_normal Normal Active_standby
```


- If yes, go to **Step 9**.
- If no, go to **Step 16**.

Check whether the standby node disk space is full.

Step 9 Log in to the standby OMS Database node as user **root**.

Step 10 Run the **su - omm** command to switch to user **omm**.

Step 11 Run the **echo ${BIGDATA_DATA_HOME}/dbdata_om** command to obtain the OMS Database data directory.

Step 12 Run the **df -h** command to view the system disk partition usage information.

Step 13 Check whether the disk where the OMS Database data directory is mounted is full.

- If yes, go to **Step 14**.
- If no, go to **Step 16**.

Step 14 Expand the disk capacity.

Step 15 After the disk capacity is expanded, wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Collect fault information.

Step 16 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 17 Select **OMMServer** from the **Service** and click **OK**.

Step 18 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.22 ALM-12040 Insufficient System Entropy

Alarm Description

MRS 3.2.0 or later:

The system checks whether the rng-tools or haveged tool has been enabled and correctly configured every 5 minutes. If neither tool is configured, this alarm is

generated. If either is configured, the system continues to check the entropy. If the entropy is less than 100 for five consecutive times, this alarm is generated.

This alarm is cleared when rng-tools or haveged has been installed and enabled on the target node and the entropy of the OS is greater than or equal to 100 in at least one of five entropy checks.

MRS 3.1.2 or earlier:

The system checks the entropy for five consecutive times at 00:00 every day. Specifically, the system checks whether rng-tools or haveged has been enabled and correctly configured. If neither is configured, the system continues to check the entropy. If the entropy is less than 100 for five consecutive times, this alarm is reported.

This alarm is cleared when the system detects that the true random number mode has been configured, the random number parameters have been configured in the pseudo-random number mode, or neither mode is configured but the entropy of the OS is greater than or equal to 100 in at least one of five entropy checks.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12040	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The system is not running properly.

Possible Causes

- rng-tools or haveged has not been installed or started.
- The entropy of the OS is smaller than 100 for multiple consecutive times.

Handling Procedure

Check whether haveged or rng-tools has been installed or started.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**.

Step 2 Check the value of **HostName** in the **Location** area to obtain the name of the host for which the alarm is generated.

Step 3 Log in to the node for which the alarm is generated as user **root**.

Step 4 Run the **/bin/rpm -qa | grep -w "haveged"** command to check the haveged installation status and check whether the command output is empty.

- If yes, go to **Step 6**.
- If no, go to **Step 5**.

Step 5 Run the **/sbin/service haveged status |grep "running"** command and check the command output.

- If the command is executed successfully, haveged has been installed and configured correctly and is running properly. Go to **Step 8**.
- If the command fails to execute, haveged is not running properly. Run the following command to manually restart haveged and go to **Step 9**:

  systemctl restart haveged.service

Step 6 Run the **/bin/rpm -qa | grep -w "rng-tools"** command to check the rng-tools installation and check whether the command output is empty.

- If yes, contact the OS vendor to install and start haveged or rng-tools. Then go to **Step 9**.
- If no, go to **Step 7**.

Step 7 Run the **ps -ef | grep -v "grep" | grep rngd | tr -d " " | grep "\-r/dev/urandom"** command and check the command output.

- If the command is executed successfully, rngd has been installed and configured correctly and is running properly. Go to **Step 8**.
- If the command fails to execute, rngd is not running properly. Run the following command to manually restart rngd and go to **Step 9**:

  systemctl restart rngd.service

Check the entropy of the OS.

Step 8 Manually check the entropy of the OS.

Log in to the target node as user **root** and run the **cat /proc/sys/kernel/random/entropy_avail** command to check whether the entropy of the OS meets cluster installation requirements (no less than 100).

- If yes, the entropy of the OS is not less than 100. Go to **Step 9**.
- If no, the entropy of the OS is less than 100. Use either of the following methods and go to **Step 9**.
  - Method 1: Use haveged (true random number mode). Contact the OS vendor to install and start haveged.
  - Method 2: Use rng-tools (pseudo-random number mode). Contact the OS vendor to install and start rng-tools and configure it based on the OS type.

Step 9 Wait until the system to check the entropy at 00:00 on the following day and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Select **NodeAgent** for **Service** and click **OK**.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

10.23 ALM-12041 Incorrect Permission on Key Files

Description

The system checks whether the permission, user, and user group information about critical directories or files is normal every 5 minutes. This alarm is generated when the information is abnormal.

This alarm is cleared when the information becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12041	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
PathName	Specifies the path or name of the abnormal file.

Impact on the System

System functions are unavailable.

Possible Causes

The file permission is abnormal or the file is lost due to a user manually modified information such as the file permission, user, and user group, or the system is powered off unexpectedly.

Procedure

Check whether the abnormal file exists and whether the permission on the abnormal file is correct.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**.

Step 2 Check the value of **HostName** to obtain the host name involved in this alarm. Check the value of **PathName** to obtain the path or name of the abnormal file.

Step 3 Log in to the node for which the alarm is generated as user **root**.

Step 4 Run the **ll** *pathName* command, where *pathName* indicates the name of the abnormal file to obtain the user, permission, and user group information about the file or directory.

Step 5 Go to **${BIGDATA_HOME}/om-agent/nodeagent/etc/agent/autocheck** directory. Then run the **vi keyfile** command and search for the name of the abnormal file and check the due permission of the file.

NOTE

To ensure proper configuration synchronization between the active and standby OMS servers, files, directories, and files and sub-directories in the directories configured in **$OMS_RUN_PATH/workspace/ha/module/hasync/plugin/conf/filesync.xml** will also be monitored except files and directories in **keyfile**. User **omm** must have read and write permissions of files and read and execute permissions of directories.

Step 6 Compare the real-world permission of the file with the due permission obtained in **Step 5** and correct the permission, user, and user group information for the file.

Step 7 Wait a hour and check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 8**.

NOTE

If the disk partition where the cluster installation directory resides is used up, some temporary files will be generated in the program installation directory when running the **sed** command fails. Users do not have the read, write, and execute permissions of these temporary files. The system reports an alarm indicating that permissions of temporary files are abnormal if these files are within the monitoring range of the alarm. Perform the preceding alarm handling processes to clear the alarm. Alternatively, you can directly delete the temporary files after confirming that files with abnormal permissions are temporary. The temporary file generated after a **sed** command execution failure is similar to the following.

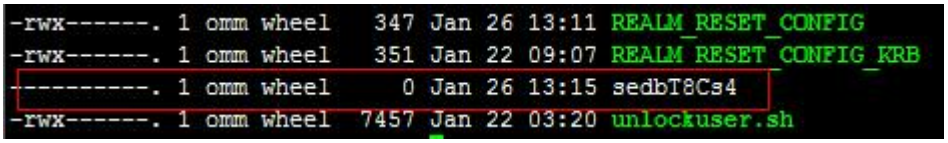

Collect fault information.

Step 8 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 9 Select **NodeAgent** from the **Service** and click **OK**.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.24 ALM-12042 Incorrect Configuration of Key Files

Description

The system checks whether critical configurations are correct every 5 minutes. This alarm is generated when the configurations are abnormal.

This alarm is cleared when the configurations become normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12042	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
PathName	Specifies the path or name of the abnormal file.

Impact on the System

Functions related to the file are abnormal.

Possible Causes

The file configuration is modified manually or the system is powered off unexpectedly.

Procedure

Check abnormal file configuration.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**.

Step 2 Check the value of **HostName** to obtain the host name involved in this alarm. Check the value of **PathName** to obtain the path or name of the abnormal file.

Step 3 Log in to the node for which the alarm is generated as user **root**.

Step 4 View the $BIGDATA_LOG_HOME/nodeagent/scriptlog/checkfileconfig.log file and analyze the cause based on the error log. Locate the check standards of the file in the **Related Information** and manually check and modify the file based on the standards.

Run the **vi** *file name* command to enter the editing mode, and then press **Insert** to start editing.

After the modification is complete, press **Esc** to exit the editing mode and enter **:wq** to save the settings and exit.

For example:

vi /etc/ssh/sshd_config

Step 5 Wait a hour and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 7 Select **NodeAgent** from the **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

- **Check standards of /etc/fstab**

  Check whether the partitions configured in the **/etc/fstab** file can be found in **/proc/mounts**.

  Check whether the swap partitions configured in fstab correspond to those in /proc/swaps.

- **Check the /etc/hosts configuration file.**

  Run **cat /ect/hosts**. If any of the following situations occurs, the **/etc/hosts** configuration file is abnormal:

  a. The **/etc/hosts** file does not exist.

  b. The host name is not configured in the file.

  c. The host name maps to multiple IP addresses in the file.

  d. The IP address corresponding to the host name does not exist in the command output of the **ifconfig** command.

  e. One IP address maps to multiple host names in the file.

- **Check standards of /etc/ssh/sshd_config**

  Run the **vi /etc/ssh/sshd_config** command to check whether configuration items are configured as follows:

  a. The value of **UseDNS** must be set to **no**.

  b. The value of **MaxStartups** must be greater than or equal to 1000.

  c. At least one of the **PasswordAuthentication** and **ChallengeResponseAuthentication** parameters must be left blank or at least one of the parameters be set to **yes**.

10.25 ALM-12045 Read Packet Dropped Rate Exceeds the Threshold

Alarm Description

The system checks the read packet dropped rate every 30 seconds. This alarm is generated when the read packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is **5**).

To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Packet Dropped Rate**.

This alarm is cleared when **Trigger Count** is 1 and the read packet dropped rate is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the read packet dropped rate is less than or equal to 90% of the threshold.

The alarm detection is disabled by default. If you want to enable this function, check whether this function can be enabled based on Checking System Environments.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12045	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
PortName	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The service performance deteriorates or some services time out.

Risk warning: In SUSE kernel 3.0 or later or Red Hat 7.2, the system kernel modifies the mechanism for counting the number of dropped read packets. In this case, this alarm may be generated even if the network is running properly, but services are not affected. You are advised to check the system environment first.

Possible Causes

- The NICs are bonded in active/standby mode.
- The alarm threshold is improperly configured.
- The network quality is poor.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click ⌄ in the row containing the alarm, and view the name of the host for which the alarm is generated and the NIC name.

Check whether the NICs are bonded in active/standby mode.

Step 2 Log in to the alarm node as user **omm** and run the **ls -l /proc/net/bonding** command to check whether the **/proc/net/bonding** directory exists on the node.

- If yes, the bond mode is configured for the node. Go to **Step 3**.
  ```
  # ls -l /proc/net/bonding/
  total 0
  -r--r--r-- 1 root root 0 Oct 11 17:35 bond0
  ```
- If no, the bond mode is not configured for the node. Go to **Step 5**.
  ```
  # ls -l /proc/net/bonding/
  ls: cannot access /proc/net/bonding/: No such file or directory
  ```

Step 3 Run the **cat /proc/net/bonding/***bond0* command to check whether the value of **Bonding Mode** in the configuration file is **fault-tolerance**.

NOTE

In the command, **bond0** indicates the name of the bond configuration file. Use the file name obtained in **Step 2**.

```
# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)

Bonding Mode: fault-tolerance (active-backup)
Primary Slave: eth1 (primary_reselect always)
Currently Active Slave: eth1
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0

Slave Interface: eth0
MII Status: up
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0

Slave Interface: eth1
MII Status: up
```



```
Speed: 1000 Mbps
Duplex: full
Link Failure Count: 1
Slave queue ID: 0
```

- If yes, the NICs are bonded in active/standby mode. Go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Check whether the NIC specified by **NetworkCardName** in the alarm is the standby NIC.

- If yes, the alarm of the standby NIC cannot be automatically cleared. Manually clear the alarm on the alarm management page. No further action is required.
- If no, go to **Step 5**.

NOTE

To determine the standby NIC, check the **/proc/net/bonding/bond0** configuration file. If the NIC name corresponding to **NetworkCardName** is **Slave Interface** but not **Currently Active Slave** (the current active NIC), the NIC is the standby one.

Check whether the threshold is set properly.

Step 5 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Packet Dropped Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 Choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Packet Dropped Rate**. Click **Modify** in the **Operation** column to change the threshold. See **Figure 10-5**.

Figure 10-5 Configuring the alarm threshold

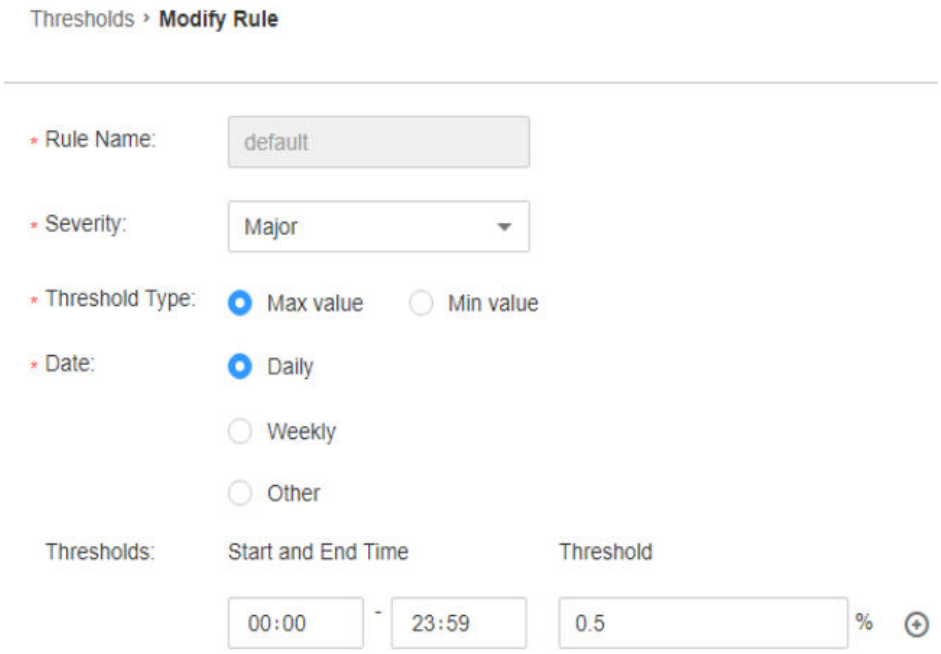

Step 7 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check whether the network connection is normal.

Step 8 Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to **Step 9**.
- If no, go to **Step 10**.

Step 9 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect the fault information.

Step 10 On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Select **OMS** for **Service** and click **OK**.

Step 12 Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.26 ALM-12046 Write Packet Dropped Rate Exceeds the Threshold

Alarm Description

The system checks the write packet dropped rate every 30 seconds. This alarm is generated when the write packet dropped rate exceeds the threshold (the default threshold is 0.5%) for multiple times (the default value is **5**).

To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Packet Dropped Rate**.

If **Trigger Count** is **1**, this alarm is cleared when the network write packet dropped rate is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this

alarm is cleared when the network write packet dropped rate is less than or equal to 90% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12046	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Port Name	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The service performance deteriorates or some services time out.

Possible Causes

- The alarm threshold is improperly configured.
- The network quality is poor.

Handling Procedure

Check whether the threshold is set properly.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Packet Dropped Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Packet Dropped Rate**. Click **Modify** in the **Operation** column to change the threshold.

See **Figure 10-6**.

Figure 10-6 Configuring the alarm threshold

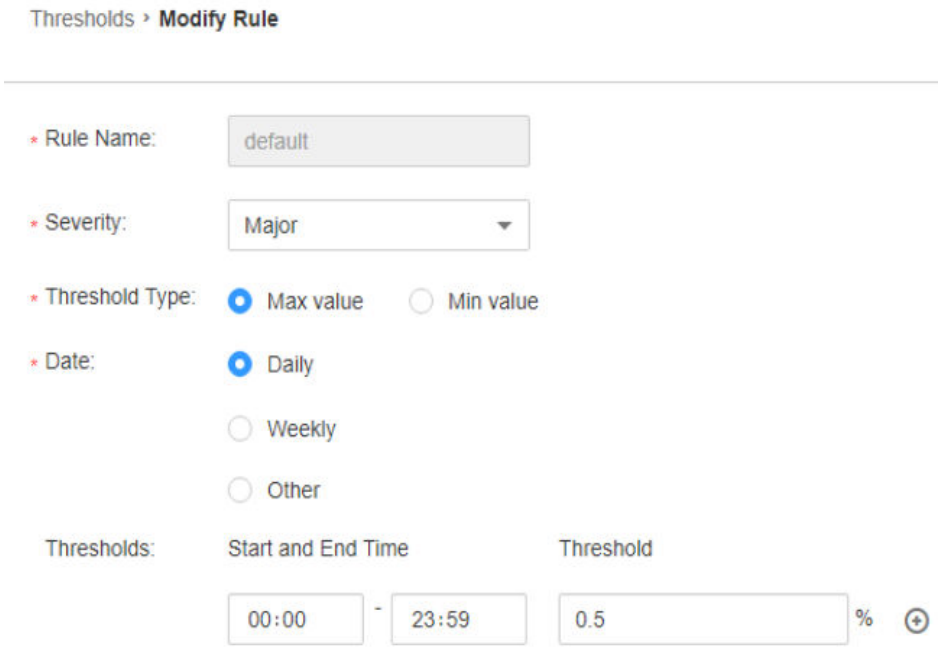

Step 3 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network connection is normal.

Step 4 Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to **Step 5**.
- If no, go to **Step 6**.

Step 5 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Select **OMS** for **Service** and click **OK**.

Step 8 Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.27 ALM-12047 Read Packet Error Rate Exceeds the Threshold

Alarm Description

The system checks the read packet error rate every 30 seconds. This alarm is generated when the read packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Packet Error Rate**.

If **Trigger Count** is **1**, this alarm is cleared when the read packet error rate is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the read packet error rate is less than or equal to 90% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12047	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Parameter	Description
Port Name	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The communication is intermittently interrupted, and services time out.

Possible Causes

- The alarm threshold is improperly configured.
- The network quality is poor.

Handling Procedure

Check whether the threshold is set properly.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Packet Error Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Packet Error Rate**. Click **Modify** in the **Operation** column to change the threshold.

See **Figure 10-7**.

Figure 10-7 Configuring the alarm threshold

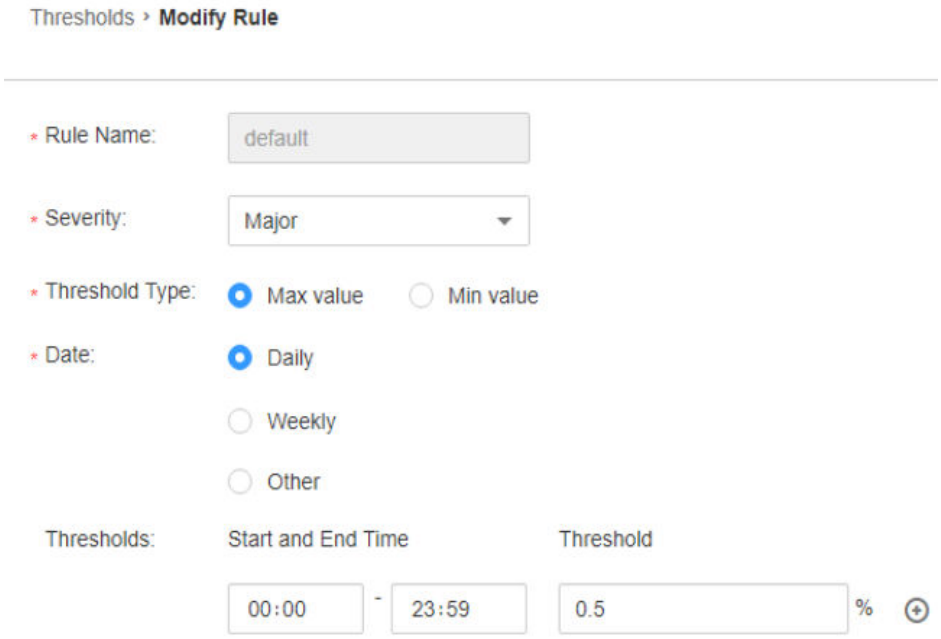

Step 3 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network connection is normal.

Step 4 Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to **Step 5**.
- If no, go to **Step 6**.

Step 5 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Select **OMS** for **Service** and click **OK**.

Step 8 Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.28 ALM-12048 Write Packet Error Rate Exceeds the Threshold

Alarm Description

The system checks the write packet error rate every 30 seconds. This alarm is generated when the write packet error rate exceeds the threshold (the default threshold is **0.5%**) for multiple times (the default value is **5**).

To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Packet Error Rate**.

If **Trigger Count** is **1**, this alarm is cleared when the write packet error rate is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is

cleared when the write packet error rate is less than or equal to 90% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12048	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Port Name	Specifies the network port for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The communication is intermittently interrupted, and services time out.

Possible Causes

- The alarm threshold is improperly configured.
- The network quality is poor.

Handling Procedure

Check whether the threshold is set properly.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Packet Error Rate**, and check whether the alarm threshold is configured properly. The default value is **0.5%**. You can adjust the threshold as needed.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Packet Error Rate**. Click **Modify** in the **Operation** column to change the threshold.

See **Figure 10-8**.

Figure 10-8 Configuring the alarm threshold

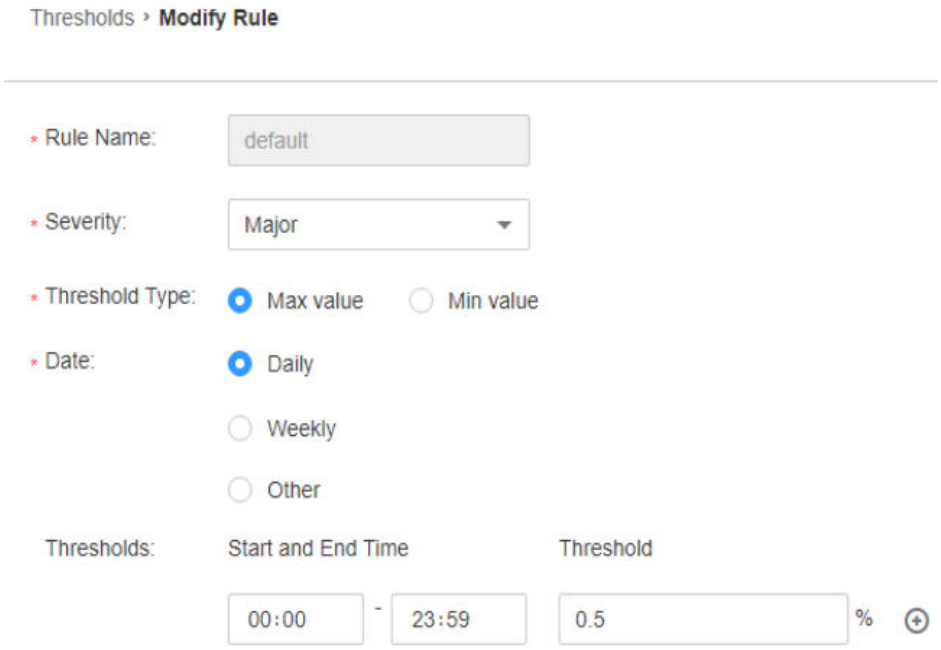

Step 3 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network connection is normal.

Step 4 Contact the network administrator to check whether the network is normal.

- If yes, rectify the fault and go to **Step 5**.
- If no, go to **Step 6**.

Step 5 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager of the active cluster, choose **O&M** > **Log** > **Download**.

Step 7 Select **OMS** for **Service** and click **OK**.

Step 8 Expand the **Hosts** dialog box and select the alarm node and the active OMS node.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.29 ALM-12049 Network Read Throughput Rate Exceeds the Threshold

Description

The system checks the network read throughput rate every 30 seconds and compares the actual throughput rate with the threshold (the default threshold is 80%). This alarm is generated when the system detects that the network read throughput rate exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Throughput Rate**.

When the **Trigger Count** is 1, this alarm is cleared when the network read throughput rate is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the network read throughput rate is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12049	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The service system runs improperly or is unavailable.

Possible Causes

- The alarm threshold is set improperly.
- The network port rate cannot meet the current service requirements.

Procedure

Check whether the threshold is set properly.

Step 1 On the FusionInsight Manager, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Throughput Rate** and check whether the alarm threshold is set properly. (By default, 80% is a proper value. However, users can configure the value as required.)

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Based on actual usage condition, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Reading** > **Read Throughput Rate** and click **Modify** in the **Operation** column to modify the alarm threshold.

For details, see **Figure 10-9**.

Figure 10-9 Setting alarm thresholds

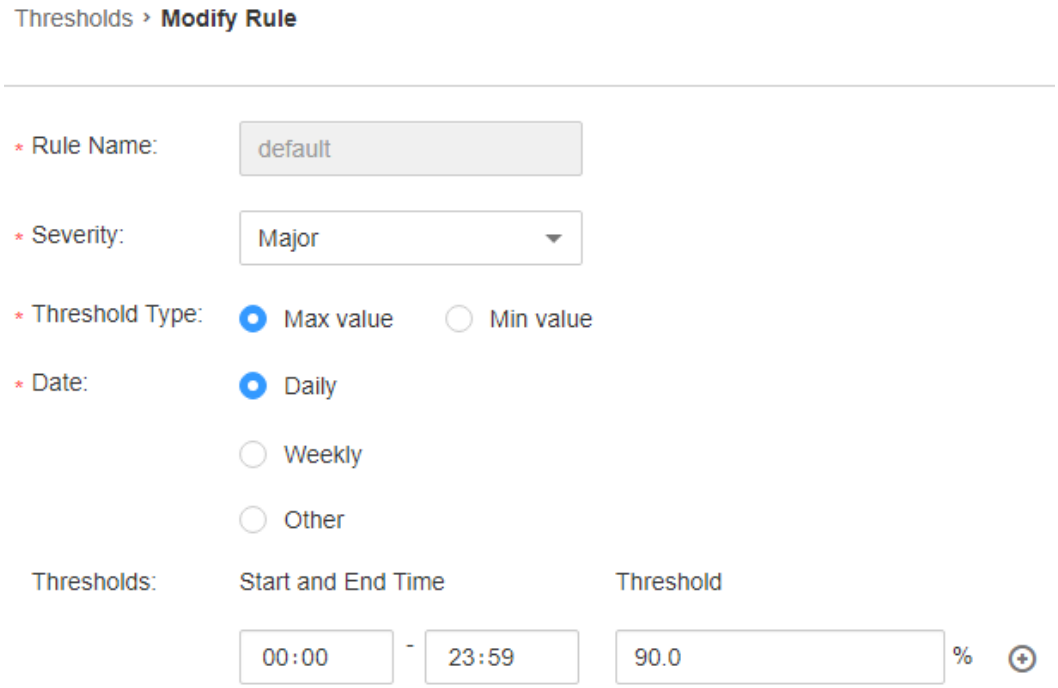

Step 3 Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network port rate can meet the service requirements.

Step 4 On FusionInsight Manager, click ⌄ in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host and the network port name for which the alarm is generated.

Step 5 Log in to the host for which the alarm is generated as user **root**.

Step 6 Run the **ethtool** *network port name* command to check the maximum speed of the current network port.

NOTE

In the VM environment, you cannot run a command to query the network port rate. It is recommended that you contact the system administrator to confirm whether the network port rate meets the requirements.

Step 7 If the network read throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log** > **Download**.

Step 10 Select **OMS** from the **Service** and click **OK**.

Step 11 Set **Host** to the node for which the alarm is generated and the active OMS node.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.30 ALM-12050 Network Write Throughput Rate Exceeds the Threshold

Description

The system checks the network write throughput rate every 30 seconds and compares the actual throughput rate with the threshold (the default threshold is 80%). This alarm is generated when the system detects that the network write throughput rate exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Throughput Rate**.

When the **Trigger Count** is 1, this alarm is cleared when the network write throughput rate is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the network write throughput rate is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12050	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NetworkCardName	Specifies the network port for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The service system runs improperly or is unavailable.

Possible Causes

- The alarm threshold is set improperly.
- The network port rate cannot meet the current service requirements.

Procedure

Check whether the threshold is set properly.

Step 1 On the FusionInsight Manager, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Throughput Rate** and check whether the alarm threshold is set properly. (By default, 80% is a proper value. However, users can configure the value as required.)

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Based on actual usage condition, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Writing** > **Write Throughput Rate** and click **Modify** in the **Operation** column to modify the alarm threshold.

For details, see **Figure 10-10**.

Figure 10-10 Setting alarm thresholds

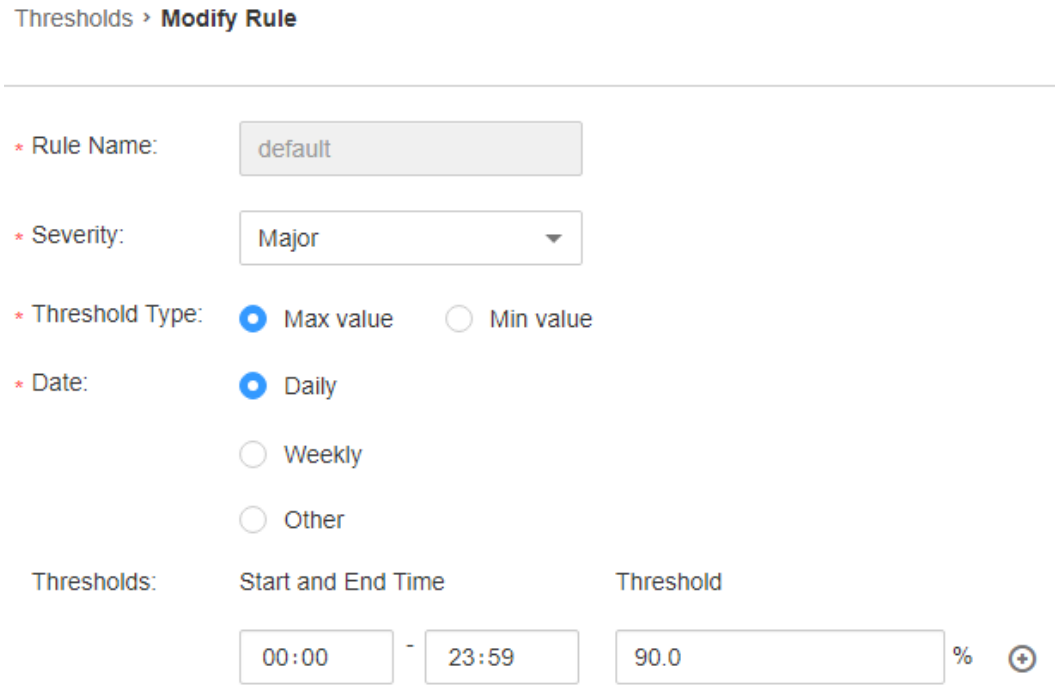

Step 3 Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the network port rate can meet the service requirements.

Step 4 On FusionInsight Manager, click ⌄ in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host and the network port name for which the alarm is generated.

Step 5 Log in to the host for which the alarm is generated as user **root**.

Step 6 Run the **ethtool***network port name* command to check the maximum speed of the current network port.

NOTE

In the VM environment, you cannot run a command to query the network port rate. It is recommended that you contact the system administrator to confirm whether the network port rate meets the requirements.

Step 7 If the network write throughput rate exceeds the threshold, contact the system administrator to increase the network port rate.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log** > **Download**.

Step 10 Select **OMS** from the **Service** and click **OK**.

Step 11 Set **Host** to the node for which the alarm is generated and the active OMS node.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.31 ALM-12051 Disk Inode Usage Exceeds the Threshold

Description

The system checks the disk Inode usage every 30 seconds and compares the actual Inode usage with the threshold (the default threshold is 80%). This alarm is generated when the Inode usage exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Disk** > **Disk Inode Usage**.

When the **Trigger Count** is 1, this alarm is cleared when the disk Inode usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the disk Inode usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12051	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Data cannot be properly written to the file system.

Possible Causes

Massive small files are stored in the disk.

Procedure

Massive small files are stored in the disk.

Step 1 On FusionInsight Manager, choose **O&M > Alarm > Alarms** and click ⌄ in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host and the disk partition for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **df -i | grep -iE "***partition name***|**Filesystem" command to check the current disk Inode usage.

```
# df -i | grep -iE "xvda2|Filesystem"
Filesystem            Inodes   IUsed   IFree IUse% Mounted on
/dev/xvda2            2359296  207420 2151876    9% /
```

Step 4 If the Inode usage exceeds the threshold, manually check small files stored in the disk partition and confirm whether these small files can be deleted.

NOTE

Run the **for i in /*; do echo $i; find $i|wc -l; done** command to query the number of files in a partition. Replace **/*** with the specified partition.

```
# for i in /srv/*; do echo $i; find $i|wc -l; done
/srv/BigData
4284
/srv/ftp
1
/srv/www
13
```


- If yes, run the **rm -rf** *Path of the file or folder* to be deleted command to delete the file or folder and go to **Step 5**.

  NOTE

  Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.
- If no, expand the capacity. Then, perform **Step 5**.

Step 5 Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log** > **Download**.

Step 7 Select **OMS** from the **Service** and click **OK**.

Step 8 Set **Host** to the node for which the alarm is generated and the active OMS node.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.32 ALM-12052 TCP Temporary Port Usage Exceeds the Threshold

Description

The system checks the TCP temporary port usage every 30 seconds and compares the actual usage with the threshold (the default threshold is 80%). This alarm is generated when the TCP temporary port usage exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Network Status** > **TCP Ephemeral Port Usage**.

When the **Trigger Count** is 1, this alarm is cleared when the TCP temporary port usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the TCP temporary port usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12052	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Services on the host cannot establish external connections, and therefore they are interrupted.

Possible Causes

- The temporary port cannot meet the current service requirements.
- The system is abnormal.

Procedure

Expand the temporary port number range.

Step 1 On FusionInsight Manager, click ⌄ in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **omm**.

Step 3 Run the **cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 1** command to obtain the value of the start port and run the **cat /proc/sys/net/ipv4/ip_local_port_range |cut -f 2** command to obtain the value of the end port. The total number of temporary ports is the value of the end port minus the value of the start port. If the total number of temporary ports is smaller than 28,232, the

random port range of the OS is narrow. Contact the system administrator to increase the port range.

Step 4 Run the **ss -ant 2>/dev/null | grep -v LISTEN | awk 'NR > 2 {print $4}'| awk -F':' '{print $NF}' | awk '$1 >"***Value of the start port***" {print $1}' | sort -u | wc -l** command to calculate the number of used temporary ports.

Step 5 The formula for calculating the usage of the temporary ports is: Usage of the temporary ports = (Number of used temporary ports/Total number of temporary ports) x 100%. Check whether the temporary port usage exceeds the threshold.

- If yes, go to **Step 7**.
- If no, go to **Step 6**.

Step 6 Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check whether the system environment is abnormal.

Step 7 Run the following command to import the temporary file and view the frequently used ports in the **port_result.txt file**:

netstat -tnp|sort > $BIGDATA_HOME/tmp/port_result.txt

```
netstat -tnp|sort

Active Internet connections (w/o servers)

Proto Recv Send LocalAddress ForeignAddress State PID/ProgramName tcp   0   0 10-120-85-154:45433
10-120-85-154:9866 CLOSE_WAIT 94237/java
tcp   0   0 10-120-85-154:45434  10-120-85-154:9866 CLOSE_WAIT 94237/java
tcp   0   0 10-120-85-154:45435  10-120-85-154:9866 CLOSE_WAIT 94237/java
...
```

Step 8 Run the following command to view the processes that occupy a large number of ports:

ps -ef |grep *PID*

NOTE

- PID is the processes ID queried in **Step 7**.
- Run the following command to collect information about all processes and check the processes that occupy a large number of ports:

  ps -ef > $BIGDATA_HOME/tmp/ps_result.txt

Step 9 After obtaining the administrator's approval, clear the processes that occupy a large number of ports. Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log** > **Download**.

Step 11 Select **OMS** from the **Service** and click **OK**.

Step 12 Set **Host** to the node for which the alarm is generated and the active OMS node.

Step 13 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected log information and files **port_result.txt** and **ps_result.txt**. Then, delete the two residual temporary files from the environment.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.33 ALM-12053 Host File Handle Usage Exceeds the Threshold

Description

The system checks the file handle usage every 30 seconds and compares the actual usage with the threshold (the default threshold is 80%). This alarm is generated when the host file handle usage exceeds the threshold for several times (5 times by default) consecutively.

To change the threshold, choose **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Host** > **Host Status** > **Host File Handle Usage**.

When the **Trigger Count** is 1, this alarm is cleared when the host file handle usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the host file handle usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12053	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The I/O operations, such as opening a file or connecting to network, cannot be performed and programs are abnormal.

Possible Causes

- The application process is abnormal. For example, the opened file or socket is not closed.
- The number of file handles cannot meet the current service requirements.
- The system is abnormal.

Procedure

Check information about files opened in processes.

Step 1 On FusionInsight Manager, click ⌄ in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **lsof -n|awk '{print $2}'|sort|uniq -c|sort -nr|more** command to check the process that occupies excessive file handles.

Step 4 Check whether the processes in which a large number of files are opened are normal. For example, check whether there are files or sockets not closed.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Release the abnormal processes that occupy too many file handles.

Step 6 Five minutes later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Increase the number of file handles.

Step 7 On FusionInsight Manager, click ⌄ in the row where the alarm is located in the real-time alarm list and obtain the IP address of the host for which the alarm is generated.

Step 8 Log in to the host for which the alarm is generated as user **root**.

Step 9 Contact the system administrator to increase the number of system file handles.

Step 10 Run the **cat /proc/sys/fs/file-nr** command to view the used handles and the maximum number of file handles. The first value is the number of used handles, the third value is the maximum number. Please check whether the usage exceeds the threshold.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

```
# cat /proc/sys/fs/file-nr
12704 0 640000
```

Step 11 Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check whether the system environment is abnormal.

Step 12 Contact the system administrator to check whether the operating system is abnormal.

- If yes, go to **Step 13** to rectify the fault.
- If no, go to **Step 14**.

Step 13 Wait for 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Collect fault information.

Step 14 On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log** > **Download**.

Step 15 Select **OMS** from the **Service** and click **OK**.

Step 16 Set **Host** to the node for which the alarm is generated and the active OMS node.

Step 17 Click ✎ in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.34 ALM-12054 Invalid Certificate File

Alarm Description

The system checks whether the certificate file is invalid (has expired or is not valid yet) on 23:00 every day. This alarm is generated when the certificate file is invalid.

This alarm is cleared when a valid certificate is imported and the alarm detection mechanism is triggered on the next hour.

NOTE

For MRS 3.2.0 or later, the certificate file is checked at the beginning of each hour.
For versions earlier than MRS 3.2.0, the certificate file is checked on 23:00 every day.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12054	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Some functions are unavailable.

Possible Causes

No certificate (CA certificate, HA root certificate, HA user certificate, Gaussdb root certificate, or Gaussdb user certificate) is imported to the system, the certificate fails to be imported, or the certificate file is invalid.

Handling Procedure

Check the alarm cause.

Step 1 On FusionInsight Manager, locate the target alarm in the real-time alarm list and click ⌄.

View **Additional Information** to obtain the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, log in to the active OMS management node as user **omm** and go to **Step 2**.
- If **HA root Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to **Step 3**.
- If **HA server Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to **Step 4**.
- If **Certificate has expired** is displayed in the additional information, view **Location** to obtain the name of the host for which the alarm is generated. Then, log in to the host as user **omm** and perform **Step 2** to **Step 4** in sequence to check whether the certificates have expired. If these certificates have not expired, check whether other certificates have been imported. If yes, import the certificate files again.

Check the validity period of the certificate files in the system.

Step 2 Check whether the current system time is in the validity period of the CA certificate.

Run the **bash ${CONTROLLER_HOME}/security/cert/conf/querycertvalidity.sh** command to check the effective time and due time of the CA root certificate.

- If yes, go to **Step 7**.
- If no, go to **Step 5**.

Step 3 Check whether the current system time is in the validity period of the HA root certificate.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/root-ca.crt** command to check the effective time and due time of the HA root certificate.

- If yes, go to **Step 7**.
- If no, go to **Step 6**.

Step 4 Check whether the current system time is in the validity period of the HA user certificate.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/server.crt** command to check the effective time and due time of the HA user certificate.

- If yes, go to **Step 7**.
- If no, go to **Step 6**.

The following is an example of the effective time and due time of a CA or HA certificate:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            97:d5:0e:84:af:ec:34:d8
        Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=CN, ST=xxx, L=yyy, O=zzz, OU=IT, CN=HADOOP.COM
        Validity
            Not Before: Dec 13 06:38:26 2016 GMT            // Effective time
            Not After : Dec 11 06:38:26 2026 GMT            // Due time
```

Import certificate files.

Step 5 Import a new CA certificate file.

Apply for or generate a new CA certificate file and import it to the system. For details, see **Replacing the CA Certificate**. The alarm is automatically cleared after the CA certificate is imported. Check whether this alarm is reported again during periodic check.

- If yes, go to **Step 7**.
- If no, no further action is required.

Step 6 Import a new HA certificate file.

Apply for or generate a new HA certificate file and import it to the system. For details, see **Replacing HA Certificates**. The alarm is automatically cleared after the CA certificate is imported. Check whether this alarm is reported again during periodic check.

- If yes, go to **Step 7**.
- If no, no further action is required.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 In the **Services** area, select **Controller**, **OmmServer**, **OmmCore**, and **Tomcat**, and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.35 ALM-12055 Certificate File Is About to Expire

Alarm Description

The system checks the certificate file on 23:00 every day. This alarm is generated if the certificate file is about to expire within 30 days.

This alarm is cleared when a certificate that is not about to expire is imported and the alarm detection mechanism is triggered on the next hour.

NOTE

For MRS 3.2.0 or later, the certificate file is checked at the beginning of each hour.
For versions earlier than MRS 3.2.0, the certificate file is checked on 23:00 every day.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12055	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Some functions are unavailable.

Possible Causes

The remaining validity period of a system certificate (CA certificate, HA root certificate, HA user certificate, Gaussdb root certificate, or Gaussdb user certificate) is less than 30 days.

Handling Procedure

Check the alarm cause.

Step 1 On FusionInsight Manager, locate the target alarm in the real-time alarm list and click ⌄.

View **Additional Information** to obtain the additional information about the alarm.

- If **CA Certificate** is displayed in the additional alarm information, log in to the active OMS management node as user **omm** and go to **Step 2**.
- If **HA root Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to **Step 3**.
- If **HA server Certificate** is displayed in the additional information, view **Location** to obtain the name of the host involved in this alarm. Then, log in to the host as user **omm** and go to **Step 4**.

Check the validity period of the certificate files in the system.

Step 2 Check whether the remaining validity period of the CA certificate is smaller than the alarm threshold.

Run the **bash ${CONTROLLER_HOME}/security/cert/conf/querycertvalidity.sh** command to check the effective time and due time of the CA root certificate.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 3 Check whether the remaining validity period of the HA root certificate is smaller than the alarm threshold.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/root-ca.crt** command to check the effective time and due time of the HA root certificate.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 4 Check whether the remaining validity period of the HA user certificate is smaller than the alarm threshold.

Run the **openssl x509 -noout -text -in ${CONTROLLER_HOME}/security/certHA/server.crt** command to check the effective time and due time of the HA user certificate.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

The following is an example of the effective time and due time of a CA or HA certificate:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            97:d5:0e:84:af:ec:34:d8
        Signature Algorithm: sha256WithRSAEncryption
```



```
Issuer: C=CN, ST=xxx, L=yyy, O=zzz, OU=IT, CN=HADOOP.COM
Validity
   Not Before: Dec 13 06:38:26 2016 GMT          // Effective time
   Not After : Dec 11 06:38:26 2026 GMT          // Due time
```

Import certificate files.

Step 5 Import a new CA certificate file.

Apply for or generate a new CA certificate file and import it to the system. For details, see **Replacing the CA Certificate**. Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to **Step 7**.
- If no, no further action is required.

Step 6 Import a new HA certificate file.

Apply for or generate a new HA certificate file and import it to the system. For details, see **Replacing HA Certificates**. Manually clear the alarm and check whether this alarm is generated again during periodic check.

- If yes, go to **Step 7**.
- If no, no further action is required.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 In the **Services** area, select **Controller**, **OmmServer**, **OmmCore**, and **Tomcat**, and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.36 ALM-12057 Metadata Not Configured with the Task to Periodically Back Up Data to a Third-Party Server

Description

After the system is installed, it checks whether the task for periodically backing up metadata to the third-party server, and then performs the check hourly. If the task for periodically backing up metadata to a third-party server is not configured, a critical alarm is generated.

This alarm is cleared when a user creates such a backup task.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12057	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If metadata is not backed up to a third-party server, metadata cannot be restored if both the active and standby management nodes of the cluster are faulty and local backup data is lost.

Possible Causes

Metadata is not configured with the task to periodically back up data to a third-party server.

Procedure

Step 1 On the FusionInsight Manager portal choose **O&M > Alarm > Alarms**.

Step 2 In the alarm list, click ⌄ in the row where the alarm is located and identify the data module from which the alarm is generated based on **Additional Information**.

Step 3 Choose **O&M** > **Backup and Restoration > Backup Management** > **Create**.

Step 4 Configure a backup task. The backup data to be configured is consistent with the data in Additional Information of the alarm.

Back up data to a third-party server, for example, remote HDFS (RemoteHDFS), NAS (NFS/CIFS), Object Storage Service (OBS), and SFTP server (SFTP). For details, see **Backing Up Data**.

Step 5 After the backup task is created successfully, wait for two minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information

Step 6 On FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 7 In the **Service** area, select **Controller** and click **OK**.

Step 8 Click ✎ in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.37 ALM-12061 Process Usage Exceeds the Threshold

Description

The system checks the usage of the omm process every 30 seconds. Users can run the **ps -o nlwp, pid, args, -u omm | awk '{sum+=$1} END {print "", sum}'** command to obtain the number of concurrent processes of user **omm**. Run the **ulimit -u**command to obtain the maximum number of processes that can be simultaneously opened by user **omm**. Divide the number of concurrent processes by the maximum number to obtain the process usage of user **omm**. The process

usage has a default threshold. This alarm is generated when the process usage exceeds the threshold.

If **Trigger Count** is **3** and the process usage is less than or equal to the threshold, this alarm is cleared. If **Trigger Count** is greater than **1**and the process usage is less than or equal to 90% of the threshold, this alarm is cleared.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12061	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

- Switch to user **omm** fails.
- New omm process cannot be created.

Possible Causes

- The alarm threshold is improperly configured.
- The maximum number of processes (including threads) that can be concurrently opened by user **omm** is inappropriate.
- An excessive number of threads are opened at the same time.

Procedure

Check whether the alarm threshold or alarm hit number is properly configured.

Step 1 On the FusionInsight Manager, change the alarm threshold and **Trigger Count** based on the actual CPU usage.

Specifically, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host**> **Process** > **omm Process Usage** to change Trigger Count, as shown in **Figure 10-11**.

NOTE

The alarm is generated when the process usage exceeds the threshold for the times specified by **Trigger Count**.

Figure 10-11 Setting Trigger Count

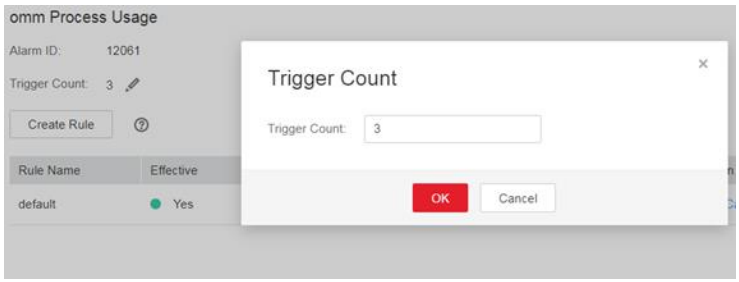

Set the alarm threshold based on the actual process usage. To check the process usage, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Host**> **Process** > **omm Process Usage**, as shown in **Figure 10-12**.

Figure 10-12 Setting an alarm threshold

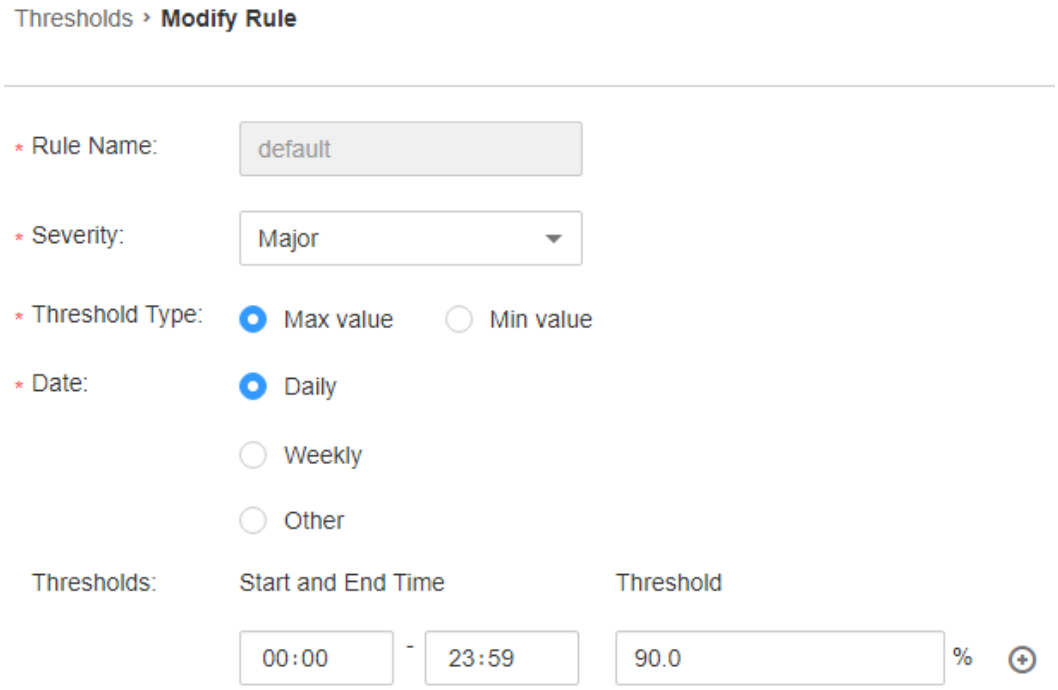

Step 2 2 minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 3**.

Check whether the maximum number of processes (including threads) opened by user omm is appropriate.

Step 3 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.

Step 4 Log in to the host where the alarm is generated as user **root**.

Step 5 Run the **su - omm** command to switch to user **omm**.

Step 6 Run the **ulimit -u** command to obtain the maximum number of threads that can be concurrently opened by user **omm** and check whether the number is greater than or equal to 60000.

- If it is, go to **Step 8**.
- If it is not, go to **Step 7**.

Step 7 Run the **ulimit -u 60000** command to change the maximum number to 60000. Two minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 12**.

Check whether an excessive number of processes are opened at the same time.

Step 8 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.

Step 9 Log in to the host where the alarm is generated as user **root**.

Step 10 Run the **ps -o nlwp, pid, lwp, args, -u omm|sort -n** command to check the numbers of threads used by the system. The result is sorted based on the thread number. Analyze the top 5 thread numbers and check whether the threads are incorrectly used. If they are, contact maintenance personnel to rectify the fault. If they are not, run the **ulimit -u** command to change the maximum number to be greater than 60000.

Step 11 Five minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 12**.

Collect fault information.

Step 12 On the FusionInsight Manager home page of the active clusters, choose **O&M** > **Log** > **Download**.

Step 13 Select **OmmServer** and **NodeAgent** from the **Service** and click **OK**.

Step 14 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 15 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.38 ALM-12062 OMS Parameter Configurations Mismatch with the Cluster Scale

Description

The system checks whether the OMS parameter configurations match with the cluster scale at each top hour. If the OMS parameter configurations do not meet the cluster scale requirements, the system generates this alarm. This alarm is automatically cleared when the OMS parameter configurations are modified.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12062	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The OMS configuration is not modified when the cluster is installed or the system capacity is expanded.

Possible Causes

The OMS parameter configurations mismatch with the cluster scale.

Procedure

Check whether the OMS parameter configurations match with the cluster scale.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.

Step 2 Log in to the host where the alarm is generated as user **root**.

Step 3 Run the **su - omm** command to switch to user **omm**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/controller/scriptlog/modify_manager_param.log** command to open the log file and search for the log file containing the following information: Current oms configurations cannot support *xx* nodes. In the information, *xx* indicates the number of nodes in the cluster.

Step 5 Optimize the current cluster configuration by following the instructions in **Optimizing Manager Configurations Based on the Number of Cluster Nodes**.

Step 6 One hour later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 Select **Controller** from the **Service** and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

Optimizing Manager Configurations Based on the Number of Cluster Nodes

Step 1 Log in to the active Manager node as user **omm**.

Step 2 Run the following command to switch the directory:

cd ${BIGDATA_HOME}/om-server/om/sbin

Step 3 Run the following command to view the current Manager configurations.

sh oms_config_info.sh -q

Step 4 Run the following command to specify the number of nodes in the current cluster.

Command format: **sh oms_config_info.sh -s** *number of nodes*

Example:

sh oms_config_info.sh -s 1000

Enter **y** as prompted.

```
The following configurations will be modified:
    Module      Parameter         Current          Target
    Controller  controller.Xmx    4096m        =>  16384m
    Controller  controller.Xms    1024m        =>  8192m
    Controller  controller.node.heartbeat.error.threshold     30000                =>   60000
    Pms         pms.mem           8192m        =>  10240m
Do you really want to do this operation? (y/n):
```

The configurations are updated successfully if the following information is displayed:

```
...
Operation has been completed. Now restarting OMS server.                [done]
Restarted oms server successfully.
```

NOTE

- OMS is automatically restarted during the configuration update process.
- Clusters with similar quantities of nodes have same Manager configurations. For example, when the number of nodes is changed from 100 to 101, no configuration item needs to be updated.

----End

10.39 ALM-12063 Unavailable Disk

Description

The system checks whether the data disk of the current host is available at the top of each hour. The system creates files, writes files, and deletes files in the mount directory of the disk. If the operations fail, the alarm is generated. If the operations succeed, the disk is available, and the alarm is cleared.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12063	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
DiskName	Specifies the disk for which the alarm is generated.

Impact on the System

Data read or write on the data disk fails, and services are abnormal.

Possible Causes

- The permission of the disk mount directory is abnormal.
- There are disk bad sectors.

Procedure

Check whether the permission of the disk mount directory is normal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host and **DiskName** for the disk for which the alarm is generated.

Step 2 Log in to the host where the alarm is generated as user **root**.

Step 3 Run the **df -h |grep DiskName** command to obtain the mount point and check whether the permission of the mount directory is unwritable or unreadable.

- If it is, go to **Step 4**.
- If it is not, go to **Step 8**.

NOTE

If the permission of the mount directory is 000 or the owner is **root**, the mount directory is unreadable and unwritable.

Step 4 Modify the directory permission.

Step 5 One hour later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Step 6 Contact hardware engineers to rectify the disk.

Step 7 One hour later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 9 Select **NodeAgent** from the **Service** and click **OK**.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.40 ALM-12064 Host Random Port Range Conflicts with Cluster Used Port

Alarm Description

The system checks whether the random port range of the host conflicts with the range of ports used by the Cluster system every hour. The alarm is generated if they conflict. The alarm is automatically cleared when the random port range of the host is changed to the normal range.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12064	Major	Yes

Parameters

Parameter	Description
Source	Specifiestheclusterorsystemforwhich-thealarmisgenerated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The default port of the Cluster system is occupied. As a result, some processes fail to be started.

Possible Causes

The random port range configuration is modified.

Procedure

Check the random port range of the system.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host for which the alarm is generated.

Step 2 Log in to the host where the alarm is generated as user **root**.

Step 3 Run the **cat /proc/sys/net/ipv4/ip_local_port_range** command to obtain the random port range of the host and check whether the minimum value is smaller than 32768.

- If it is, go to **Step 4**.
- If it is not, goto **Step 7**.

Step 4 Run the **vim /etc/sysctl.conf** command to change the value of **net.ipv4.ip_local_port_range** to **32768 61000**. If this parameter does not exist, add the following configuration: **net.ipv4.ip_local_port_range = 32768 61000**.

Step 5 Run the **sysctl -p /etc/sysctl.conf** command for the modification to take effect.

Step 6 One hour later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 Select **NodeAgent** for **Service** and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.41 ALM-12066 Trust Relationships Between Nodes Become Invalid

Alarm Description

The system checks whether the trust relationship between the active OMS node and other Agent nodes is normal every hour. The alarm is generated if the mutual trust fails. This alarm is automatically cleared after the fault is rectified.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12066	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Some operations on the management plane may be abnormal.

Possible Causes

- The **/etc/ssh/sshd_config** configuration file is damaged.
- The password of user **omm** has expired.

Handling Procedure

Check the status of the /etc/ssh/sshd_config configuration file.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm and click ⌄ to view the host list in the alarm details.

Step 2 Log in to the active OMS node as user **omm**.

Step 3 Run the **ssh** command, for example, **ssh *host2***, on each node in the alarm details to check whether the connection fails. (***host2*** is a node other than the OMS node in the alarm details.)

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 Open the **/etc/ssh/sshd_config** configuration file on host2 and check whether **AllowUsers** or **DenyUsers** is configured for other nodes.

- If yes, go to **Step 5**.
- If no, contact OS experts.

Step 5 Modify the whitelist or blacklist to ensure that user **omm** is in the whitelist or not in the blacklist. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check the status of the password of user omm.

Step 6 Check the interaction information of the **ssh** command.

- If the password of user **omm** is required, go to **Step 7**.
- If message "Enter passphrase for key '/home/omm/.ssh/id_rsa':" is displayed, go to **Step 9**.
- If the connection is set up successfully, run the following command to restart **ssh-agent** and then go to **Step 3**:

  ps -ef|grep ssh-agent |grep -v grep |awk '{print $2}' | xargs kill -9

Step 7 Check the trust list (**/home/omm/.ssh/authorized_keys**) of user **omm** on the OMS node and host2 node. Check whether the trust list contains the public key file (**/home/omm/.ssh/id_rsa.pub**) of user **omm** on the peer host.

- If yes, contact OS experts.
- If no, add the public key of user **omm** of the peer host to the trust list of the local host.

Step 8 Add the public key of user **omm** of the peer host to the trust list of the local host. Run the **ssh** command, for example, **ssh host2**, on each node in the alarm details to check whether the connection fails. (***host2*** is a node other than the OMS node in the alarm details.)

- If yes, go to **Step 9**.
- If no, check whether the alarm is cleared. If the alarm is cleared, no further action is required; otherwise, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Select **Controller** for **Service** and click **OK**.

Step 11 Click  in the upper right corner to set the log collection time range. Generally, the time range is 10 minutes before and after the alarm generation time. Click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

Perform the following steps to handle abnormal trust relationships between nodes:

NOTICE

- Perform this operation as user **omm**.
- If the network between nodes is disconnected, rectify the network fault first. Check whether the two nodes are connected to the same security group and whether **hosts.deny** and **hosts.allow** are set.

1. Run the **ssh-add -l** command on both nodes to check whether any identities exist.

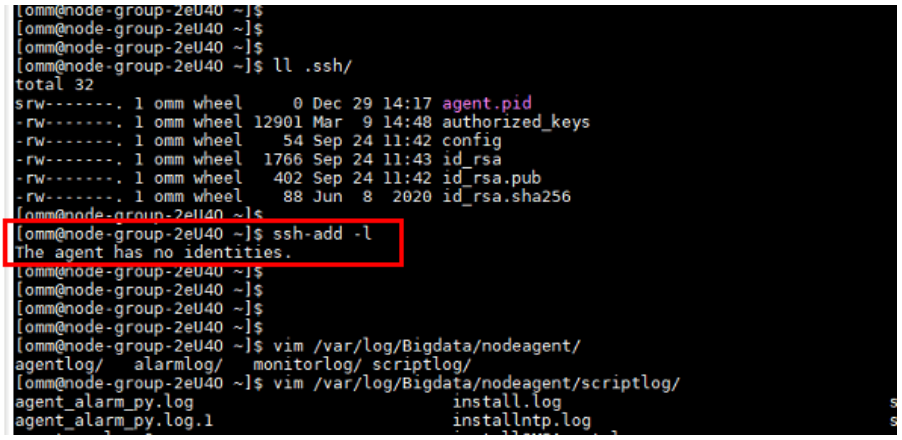

   - If yes, go to **4**.
   - If no, go to **2**.

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

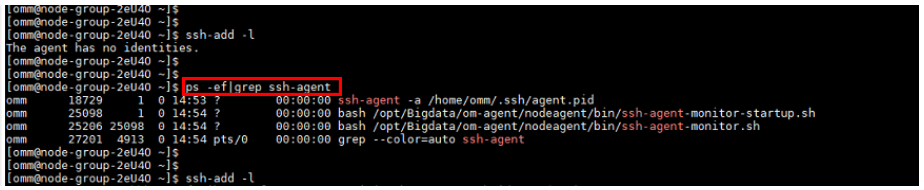

3. Run the **ssh-add -l** command to check whether the identities have been added. If yes, manually run the **ssh** command to check whether the trust relationship is normal.

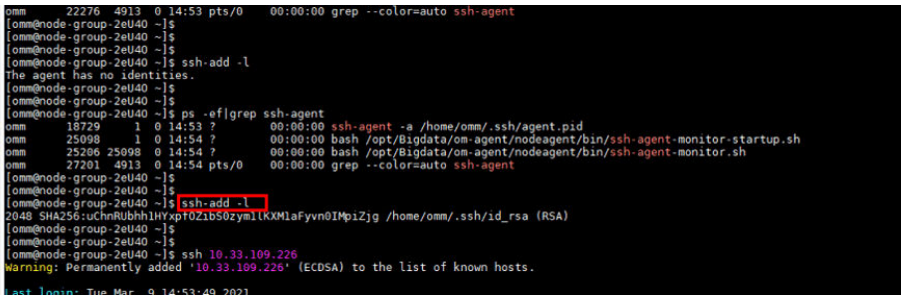

4. If identities exist, check whether the **/home/omm/.ssh/authorized_keys** file contains the information in the **/home/omm/.ssh/id_rsa.pub** file of the peer node. If it does not, manually add the information.
5. Check whether the permissions on the files in the **/home/omm/.ssh** directory are modified.
6. Check the **/var/log/Bigdata/nodeagent/scriptlog/ssh-agent-monitor.log** file.
7. If the **/home** directory of user **omm** is deleted, contact MRS support personnel for assistance.

10.42 ALM-12067 Tomcat Resource Is Abnormal

Alarm Description

HA checks the Tomcat resources of Manager every 85 seconds. This alarm is generated when HA detects that the Tomcat resources are abnormal for two consecutive times.

This alarm is cleared when HA detects that the Tomcat resources become normal.

Resource Type of Tomcat is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new Tomcat resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12067	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

- The active/standby Manager switchover occurs.
- The Tomcat process repeatedly restarts.

Possible Causes

- The Tomcat directory permission is abnormal, and the Tomcat process is abnormal.

Handling Procedure

Check whether the permission on the Tomcat directory is normal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the IP address of the host for which the alarm is generated.

Step 2 Log in to the alarm host as user **root**.

Step 3 Run the **su - omm** command to switch to user **omm**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/tomcat.log** command to check whether the Tomcat resource log contains keyword **Cannot find *XXX*** and rectify the file permission based on the keyword.

Step 5 After 5 minutes, check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 In the **Services** area, select **OmmServer** and **Tomcat**, and click **OK**.

Step 8 Click ✎ in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.43 ALM-12068 ACS Resource Exception

Alarm Description

HA checks the ACS resources of Manager every 80 seconds. This alarm is generated when HA detects that the ACS resources are abnormal for two consecutive times.

This alarm is cleared when HA detects that the ACS resources are normal.

Resource Type of ACS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new ACS resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12068	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

- The active/standby Manager switchover occurs.

- The ACS process repeatedly restarts, which may cause the FusionInsight Manager login failure.

Possible Causes

The ACS process is abnormal.

Handling Procedure

Check whether the ACS process is normal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the name of the host for which the alarm is generated.

Step 2 Log in to the alarm host as user **root**.

Step 3 Run the **su - omm** command and then **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** to check whether the status of the ACS resources managed by the HA is normal. In the single-node system, the ACS resource is in the normal state. In the dual-node system, the ACS resource is in the normal state on the active node and in the stopped state on the standby node.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/acs.log** command to check whether the ACS resource log of HA contains the keyword **ERROR**. If yes, analyze the logs to locate the resource exception cause and fix the exception.

Step 5 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.44 ALM-12069 AOS Resource Exception

Alarm Description

HA checks the AOS resources of Manager every 81 seconds. This alarm is generated when HA detects that the AOS resources are abnormal for two consecutive times.

This alarm is cleared when HA detects that the AOS resources become normal.

Resource Type of AOS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new AOS resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12069	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

- The active/standby Manager switchover occurs.
- The AOS process repeatedly restarts, which may cause the FusionInsight Manager login failure.

Possible Causes

The AOS process is abnormal.

Handling Procedure

Check whether the AOS process is normal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the name of the host for which the alarm is generated.

Step 2 Log in to the alarm host as user **root**.

Step 3 Run the **su - omm** command and then **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** to check whether the status of the AOS resources managed by the HA is normal. In the single-node system, the AOS resource is in the normal state. In the dual-node system, the AOS resource is in the normal state on the active node and in the stopped state on the standby node.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/aos.log** command to check whether the AOS resource log of HA contains the keyword **ERROR**. If yes, analyze the logs to locate the resource exception cause and fix the exception.

Step 5 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.45 ALM-12070 Controller Resource Is Abnormal

Alarm Description

HA checks the controller resources of Manager every 80 seconds. This alarm is generated when HA detects that the controller resources are abnormal for 2 consecutive times.

This alarm is cleared when the Controller resource is normal.

Resource Type of Controller is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new Controller resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12070	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The active/standby FusionInsight Manager switchover occurs.
- The Controller process repeatedly restarts, which may cause the FusionInsight Manager login failure.

Possible Causes

The Controller process is abnormal.

Procedure

Check whether the controller process is normal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su - omm** command to switch to user **omm**.Run the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the status of the Controller resources managed by the HA is normal. In the single-node system, the Controller resource is in the normal state. In the dual-node system, the Controller resource is in the normal state on the active node and in the stopped state on the standby node.

- If it is, go to **Step 6**.
- If it is not, go to **Step 4**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/controller.log** command to view the Controller resource logs, and run the **vi $BIGDATA_LOG_HOME/controller/controller.log** command to view the Controller running logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.

Step 5 Five minutes later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Select **Controller** and **OmmServe** for **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour before and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.46 ALM-12071 Httpd Resource Is Abnormal

Description

HA checks the httpd resources of Manager every 120 seconds. This alarm is generated when HA detects that the httpd resources are abnormal for 10 consecutive times.

This alarm is cleared when the httpd resource is normal.

Resource Type of httpd is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new httpd resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12071	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The active/standby FusionInsight Manager switchover occurs.
- The httpd process is repeatedly restarts, which may lead to the failure to visit the native service UI.

Possible Causes

The httpd process is abnormal.

Procedure

Check whether the httpd process is abnormal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su - omm** command to switch to user **omm**.

Step 4 Run the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the status of the httpd resources managed by the HA is normal. In the single-node system, the httpd resource is in the normal state. In the dual-node system, the httpd resource is in the normal state on the active node and in the stopped state on the standby node.

- If it is, go to **Step 7**.
- If it is not, go to **Step 5**.

Step 5 Run the **vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/httpd.log** command to view the httpd resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.

Step 6 Five minutes later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 Select **Controller** and **OmmServer** for **Service** and click **OK**.

Step 9 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.47 ALM-12072 FloatIP Resource Is Abnormal

Description

HA checks the floatip resources of Manager every 9 seconds. This alarm is generated when HA detects that the floatip resources are abnormal for 3 consecutive times.

This alarm is cleared when the FloatIP resource is normal.

Resource Type of FloatIP is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new FloatIP resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12072	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The active/standby FusionInsight Manager switchover occurs.
- The FloatIP process is repeatedly restarts, which may lead to the failure to visit the native service UI.

Possible Causes

- The floating IP address is abnormal.

Procedure

Check the floating IP address status of the active management node.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the address of the host for which the alarm is generated and the resource name.

Step 2 Log in to the active management node as user **root**.

Step 3 Run the following command, go to the **${BIGDATA_HOME}/om-server/om/sbin/** directory.

su - omm

cd ${BIGDATA_HOME}/om-server/om/sbin/

Step 4 Run the **sh status-oms.sh** command, and execute the **status-oms.sh** script to check whether the floating IP address of the active FusionInsight Manager is normal. View the command output, locate the row where **ResName** is **floatip**, and check whether the following information is displayed.

For example:

```
10-10-10-160 floatip Normal Normal Single_active
```

- If it is, go to **Step 8**.
- If it is not, go to **Step 5**.

Step 5 Run the **ifconfig** command to check whether the NIC with the floating IP address exists.

- If it does, go to **Step 8**.
- If it does not, go to **Step 6**.

Step 6 Run the **ifconfig** *NIC name Floating IPaddress* netmask *Subnet mask* command to reconfigure the NIC with the floating IP address. (For example, **ifconfig eth0 10.10.10.102 netmask 255.255.255.0**).

Step 7 Five minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 9 Select **Controller** and **OmmServer** for **Service** and click **OK**.

Step 10 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.48 ALM-12073 CEP Resource Is Abnormal

Description

HA checks the cep resources of Manager every 60 seconds. This alarm is generated when HA detects that the cep resources are abnormal for 2 consecutive times.

This alarm is cleared when the CEP resource is normal.

Resource Type of CEP is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new CEP resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12073	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The active/standby FusionInsight Manager switchover occurs.
- The CEP process repeatedly restarts, causing monitoring data to be abnormal.

Possible Causes

The CEP process is abnormal.

Procedure

Check whether the CEP process is abnormal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su -omm** command and then the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the status of the CEP resources managed by the HA is normal. In the single-node system, the CEP resource is in the normal state. In the dual-node system, the CEP resource is in the normal state on the active node and in the stopped state on the standby node.

- If it is, go to **Step 6**.
- If it is not, go to **Step 4**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/omm/oms/cep/cep.log** and **vi $BIGDATA_LOG_HOME/omm/oms/cep/scriptlog/cep_ha.log** commands to view the CEP resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.

Step 5 Five minutes later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Select **Controller** and **OmmServer** for **Service** and click **OK**.

Step 8 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.49 ALM-12074 FMS Resource Is Abnormal

Description

HA checks the fms resources of Manager every 60 seconds. This alarm is generated when HA detects that the fms resources are abnormal for 2 consecutive times.

This alarm is cleared when the FMS resource is normal.

Resource Type of FMS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new FMS resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12074	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The active/standby FusionInsight Manager switchover occurs.
- The FMS process repeatedly restarts. As a result, alarm information may fail to be reported.

Possible Causes

The FMS process is abnormal.

Procedure

Check whether the FMS process is abnormal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su -omm** command and then the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the status of the FMS resources managed by the HA is normal. In the single-node system, the FMS resource is in the normal state. In the dual-node system, the FMS resource is in the normal state on the active node and in the stopped state on the standby node.

- If it is, go to **Step 6**.
- If it is not, go to **Step 4**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/omm/oms/fms/fms.log** and **vi $BIGDATA_LOG_HOME/omm/oms/fms/scriptlog/fms_ha.log** commands to view the FMS resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.

Step 5 5 minutes later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**> **Log** > **Download**.

Step 7 Select **Controller** and **OmmServer** for **Service** and click **OK**.

Step 8 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.50 ALM-12075 PMS Resource Is Abnormal

Description

HA checks the pms resources of Manager every 55 seconds. This alarm is generated when HA detects that the pms resources are abnormal for three consecutive times.

This alarm is cleared when the PMS resource is normal.

Resource Type of PMS is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new PMS resources have been enabled on the new active FusionInsight Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby switchover.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12075	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The active/standby FusionInsight Manager switchover occurs.
- The PMS process repeatedly restarts, causing monitoring information to be abnormal.

Possible Causes

The PMS process is abnormal.

Procedure

Check whether the PMS process is abnormal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the name of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su -omm** command and then the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the status of the PMS resources managed by the HA is normal. In the single-node system, the PMS resource is in the normal state. In the dual-node system, the PMS resource is in the normal state on the active node and in the stopped state on the standby node.

- If it is, go to **Step 6**.
- If it is not, go to **Step 4**.

Step 4 Run the **vi $BIGDATA_LOG_HOME/omm/oms/pms/pms.log** and **vi $BIGDATA_LOG_HOME/omm/oms/pms/scriptlog/pms_ha.log** commands to view the PMS resource logs, check whether the keyword **ERROR** exists. Analyze the logs to locate and rectify the fault.

Step 5 Five minutes later, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**> **Log** > **Download**.

Step 7 Select **Controller** and **OmmServer** for **Service** and click **OK**.

Step 8 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 1 hour before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.51 ALM-12076 GaussDB Resource Is Abnormal

Description

HA checks the Manager database every 10 seconds. This alarm is generated when HA detects that the database is abnormal for 3 consecutive times.

This alarm is cleared when the database is normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12076	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If databases are abnormal, all core services and related service processes, such as alarms and monitoring functions, are affected.

Possible Causes

An exception occurs in the database.

Procedure

Check the database status of the active and standby management nodes.

Step 1 Log in to the active and standby management nodes respectively as user **root**. Run the **su - ommdba** command to switch to user **ommdba**, and then run the **gs_ctl query** command to check whether the following information is displayed in the command output.

Command output of the active management node:

```
Ha state:
    LOCAL_ROLE: Primary
    STATIC_CONNECTIONS          : 1
    DB_STATE                    : Normal
    DETAIL_INFORMATION          : user/password invalid
Senders info:
    No information
Receiver info:
    No information
```


Command output of the standby management node:

```
Ha state:
    LOCAL_ROLE: Standby
    STATIC_CONNECTIONS          : 1
    DB_STATE                    : Normal
    DETAIL_INFORMATION          : user/password invalid
Senders info:
    No information
Receiver info:
    No information
```

- If it is, go to **Step 3**.
- If it is not, go to **Step 2**.

Step 2 Contact the network administrator to check whether the network is faulty.

- If it is, go to **Step 3**.
- If it is not, go to **Step 5**.

Step 3 Five minutes later, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Step 4 Log in to the active and standby management nodes, run the **su -omm** command to switch to user **omm**, go to the **${BIGDATA_HOME} /om-server/om/sbin/** directory, and run the **status-oms.sh** script to check whether the floating IP addresses and GaussDB resources of the active and standby FusionInsight Managers are in the status shown in the following figure.

```
acs                 Normal              Normal          Single_active
aos                 Normal              Normal          Single_active
cep                 Normal              Normal          Single_active
controller          Normal              Normal          Single_active
feed_watchdog       Normal              Normal          Double_active
floatip             Normal              Normal          Single_active
fms                 Normal              Normal          Single_active
gaussDB             Active_normal       Normal          Active_standby
heartBeatCheck      Normal              Normal          Single_active
httpd               Normal              Normal          Single_active
iam                 Normal              Normal          Single_active
ntp                 Active_normal       Normal          Active_standby
okerberos           Normal              Normal          Double_active
oldap               Active_normal       Normal          Active_standby
pms                 Normal              Normal          Single_active
tomcat              Normal              Normal          Single_active
acs                 Stopped             Normal          Single_active
aos                 Stopped             Normal          Single_active
cep                 Stopped             Normal          Single_active
controller          Stopped             Normal          Single_active
feed_watchdog       Normal              Normal          Double_active
floatip             Stopped             Normal          Single_active
fms                 Stopped             Normal          Single_active
gaussDB             Standby_normal      Normal          Active_standby
heartBeatCheck      Stopped             Normal          Single_active
httpd               Stopped             Normal          Single_active
```

- If they are, find the alarm in the alarm list and manually clear the alarm.
- If they are not, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 6 Select **OmmServer** for **Service** and click **OK**.

Step 7 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 8 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.52 ALM-12077 User omm Expired

Description

The system starts at 00:00 every day to check whether user **omm** has expired every eight hours. This alarm is generated if the user account has expired.

This alarm is cleared when the expiration time of user **omm** is changed and the user account status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12077	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

User **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

Possible Causes

User **omm** has expired.

Procedure

Check whether user omm in the system has expired.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l omm**command to view the information about the password of user **omm**.

Step 2 View the value of **Account expires** to check whether the user configurations have expired.

NOTE

If the parameter value is **never**, the user configurations never expire.

- If they do, go to **Step 3**.
- If they do not, go to **Step 4**.

Step 3 Run the **chage -E 'yyyy-MM-dd' omm** command to set the expiration time of user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.53 ALM-12078 Password of User omm Expired

Description

The system starts at 00:00 every day to check whether the password of user **omm** has expired every 8 hours. This alarm is generated if the password has expired.

This alarm is cleared when the expiration time of user **omm** password is changed and the user password status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12078	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The password of user **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

Possible Causes

The password of user **omm** has expired.

Procedure

Check whether the password of user omm in the system has expired.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l omm**command to view the information about the password of user **omm**.

Step 2 View the value of **Password expires** to check whether the user configurations have expired.

NOTE

If the parameter value is **never**, the user configurations never expire.

- If they do, go to **Step 3**.
- If they do not, go to **Step 4**.

Step 3 Run the **chage -M '***days***' omm** command to set the validity period of the password for user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**> **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.54 ALM-12079 User omm Is About to Expire

Description

The system starts at 00:00 every day to check whether user **omm** is about to expire every 8 hours. This alarm is generated if the user account will expire no less than 15 days later.

This alarm is cleared when the expiration time of user **omm** is changed and the user account status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12079	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

User **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

Possible Causes

The account of user **omm** is about to expire.

Procedure

Check whether user omm is about to expire.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l omm** command to view the information about the password of user **omm**.

Step 2 View the value of **Account expires** to check whether the user configurations are about to expire.

NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to **Step 3**.
- If they are not, go to **Step 4**.

Step 3 Run the **chage -E** *'yyyy-MM-dd'* **omm** command to set the validity period of user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.55 ALM-12080 Password of User omm Is About to Expire

Description

The system starts at 00:00 every day to check whether the password of user **omm** is about to expire every 8 hours. This alarm is generated if the password will expire no less than 15 days later.

This alarm is cleared when the expiration time of user **omm** password is reset and the user password status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12080	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The password of user **omm** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

Possible Causes

The password of user **omm** is about to expire.

Procedure

Check whether the password of user omm in the system is about to expire.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l omm**command to view the information about the password of user **omm**.

Step 2 View the value of **Password expires** to check whether the user configurations are about to expire.

NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to **Step 3**.
- If they are not, go to **Step 4**.

Step 3 Run the **chage -M** '*days*' **omm** command to set the validity period of the password for user **omm**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**> **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.56 ALM-12081 User ommdba Expired

Description

The system starts at 00:00 every day to check whether user **ommdba** has expired every 8 hours. This alarm is generated if the user account has expired.

This alarm is cleared when the expiration time of user **ommdba** is reset and the user account status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12081	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The OMS database cannot be managed and data cannot be accessed.

Possible Causes

The account of user **ommdba** for the host has expired.

Procedure

Check whether user ommdba has expired.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about the password of user **ommdba**.

Step 2 View the value of **Account expires** to check whether the user configurations have expired.

NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password have expired.

- If they do, go to **Step 3**.
- If they do not, go to **Step 4**.

Step 3 Run the **chage -E** *'yyyy-MM-dd'* **omm** command to set the validity period of user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.57 ALM-12082 User ommdba Is About to Expire

Description

The system starts at 00:00 every day to check whether user **ommdba** is about to expire every 8 hours. This alarm is generated if the user account will expire no less than 15 days later.

This alarm is cleared when the expiration time of user **ommdba** is reset and the user account status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12082	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The OMS database cannot be managed and data cannot be accessed.

Possible Causes

The account of user **ommdba** for the host is about to expire.

Procedure

Check whether user ommdba is about to expire.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about user **ommdba**.

Step 2 View the value of **Account expires** to check whether the user configurations are about to expire.

NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to **Step 3**.
- If they are not, go to **Step 4**.

Step 3 Run the **chage -E** *'yyyy-MM-dd'* **ommdba** command to set the validity period of user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.58 ALM-12083 Password of User ommdba Is About to Expire

Description

The system starts at 00:00 every day to check whether the password of user **ommdba** is about to expire every 8 hours. This alarm is generated if the password is about to expire no less than 15 days later.

This alarm is cleared when the expiration time of user **ommdba** password is reset and the user password status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12083	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The OMS database cannot be managed and data cannot be accessed.

Possible Causes

The password of user **ommdba** is about to expire.

Procedure

Check whether the password of user ommdba in the system is about to expire.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about the password of user **ommdba**.

Step 2 View the value of **Password expires** to check whether the user configurations are about to expire.

NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password are about to expire within 15 days.

- If they are, go to **Step 3**.
- If they are not, go to **Step 4**.

Step 3 Run the **chage -M** *'days'* **ommdba** command to set the validity period of the password for user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.59 ALM-12084 Password of User ommdba Expired

Description

The system starts at 00:00 every day to check whether the password of user **ommdba** has expired every 8 hours. This alarm is generated if the password has expired.

This alarm is cleared when the expiration time of user **ommdba** password is reset and the user password status becomes normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12084	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The password of user **ommdba** has expired. The node trust relationship is unavailable, and FusionInsight Manager cannot manage the services.

Possible Causes

The password of user **ommdba** for the host has expired.

Procedure

Check whether the password of user ommdba in the system has expired.

Step 1 Log in to the faulty node as user **root**.

Run the **chage -l ommdba** command to view the information about the password of user **ommdba**.

Step 2 View the value of **Password expires** to check whether the user configurations have expired.

NOTE

If the parameter value is **never**, the user and password are valid permanently; if the value is a date, check whether the user and password have expired.

- If they do, go to **Step 3**.
- If they do not, go to **Step 4**.

Step 3 Run the **chage -M** *'days'* **ommdba** command to set the validity period of the password for user **ommdba**. Eight hours later, check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 5 Select **NodeAgent** for **Service** and click **OK**.

Step 6 Click  in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.60 ALM-12085 Service Audit Log Dump Failure

Description

The system dumps service audit logs at 03:00 every day and stores them on the OMS node. This alarm is generated when the dump fails. This alarm is cleared when the next dump succeeds.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12085	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The service audit logs may be lost.

Possible Causes

- The service audit logs are oversized.
- The OMS backup storage space is insufficient.
- The storage space of a host where the service is located is insufficient.

Procedure

Check whether the service audit logs are oversized.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and view the IP address of the host and additional information for which the alarm is generated.

Step 2 Log in to the host where the alarm is generated as user **root**.

Step 3 Run the **vi ${BIGDATA_LOG_HOME}/controller/scriptlog/getLogs.log** command to check whether the keyword "LOG SIZE is more than 5000MB" can be searched.

- If it can, go to **Step 4**.
- If it cannot, go to **Step 5**.

Step 4 Check whether the oversized service audit logs are caused by exceptions.

The OMS backup storage space is insufficient.

Step 5 Run the **vi ${BIGDATA_LOG_HOME}/controller/scriptlog/getLogs.log** command to check whether the keyword "Collect log failed, too many logs on" can be searched.

- If it can, obtain the host IP address following the keyword "Collect log failed, too many logs on", and go to **Step 6**.
- If it cannot, go to **Step 11**.

Step 6 Log in to the host with the IP address obtained in **Step 5** as user **root**.

Step 7 Run the **vi {BIGDATA_LOG_HOME}/nodeagent/scriptlog/collectLog.log** command to check whether the keyword "log size exceeds" can be searched.

- If it can, go to **Step 9**.
- If it cannot, go to **Step 8**.

Step 8 Check whether the alarm additional information contains the keyword "no enough space".

- If yes, go to **Step 9**.
- If no, go to**Step 11**.

Step 9 Perform the following operations to expand the disk capacity (only for MRS 3.1.2 and earlier versions) or reduce the maximum number of audit log backups:

- Expand the capacity of the OMS node.
- Run the following command to edit the file and decrease the value of **MAX_NUM_BK_AUDITLOG**.

  vi ${CONTROLLER_HOME}/etc/om/componentsauditlog.properties

Step 10 In the next execution period, 03:00, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 11**.

Check whether the space of the host where the service is located is insufficient.

Step 11 Run the **vi ${BIGDATA_LOG_HOME}/controller/scriptlog/getLogs.log** command to check whether the keyword "Collect log failed, no enough space on *hostIp*" can be searched.

- If it can, obtain the IP address of the abnormal host and go to **Step 12**.
- If it cannot, go to **Step 15**.

Step 12 Log in to the host with the IP address obtained as user **root**, and run the **df "$BIGDATA_HOME/tmp" -lP | tail -1 | awk '{print ($4/1024)}'** command to

obtain the remaining space of the host log directory. Check whether the value is less than 1000 MB.

- If it is, go to **Step 13**.
- If it is not, go to **Step 15**.

Step 13 Expand the capacity of the node

Step 14 In the next execution period, 03:00, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 15**.

Collect fault information.

Step 15 On FusionInsight Manager, choose **O&M**> **Log** > **Download**.

Step 16 Select **Controller** for **Service** and click **OK**.

Step 17 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 18 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.61 ALM-12087 System Is in the Upgrade Observation Period

Description

The system checks whether it is in the upgrade observation period at 00:00 every day and checks whether the duration that it has been in the upgrade observation state exceeds the preset upgrade observation period, 10 days by default. This alarm is generated when the system is in the upgrade observation period and the duration that the system has been in the upgrade observation state exceeds the preset period (10 days by default). This alarm is automatically cleared if the system exits the upgrade observation period after the user performs a rollback or submission.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12087	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Upgrade Observation Period (Days)	Specifies the days that the system is in the upgrade observation period.

Impact on the System

The next upgrade or patch installation will fail.

Possible Causes

The upgrade task is not submitted a specified period of time (10 days by default) after the system upgrade.

Procedure

Check whether the system is in the upgrade observation period.

Step 1 Log in to the active management node as user **root**.

Step 2 Run the following commands to switch to user **omm** and log in to the **omm** database:

su - omm

gsql -U omm -W *omm database password* **-p 20015**

Step 3 Run the **select * from OM_CLUSTERS** command to view cluster information.

Step 4 Check whether the value of **upgradObservationPeriod isON** is **true**, as shown in **Figure 10-13**.

- If it is, the system is in the upgrade observation period. Use the UpdateTool to submit the upgrade task. For details, see the upgrade guide of the corresponding version.
- If it is not, go to **Step 6**.

Figure 10-13 Cluster information

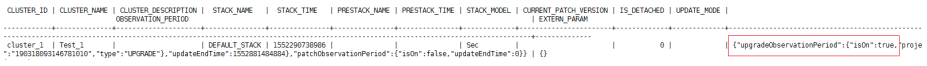

Step 5 In the early morning of the next day, check whether this alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 7 Select **Controller** from the **Service** and click **OK**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.62 ALM-12089 Inter-Node Network Is Abnormal

Description

The alarm module checks the network health status of nodes in the cluster every 10 seconds. This alarm is generated when the network between two nodes is unreachable or the network status is unstable.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12089	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Functions of some components, such as HDFS and ZooKeeper, are affected.

Possible Causes

- The node breaks down.
- The network is faulty.

Procedure

Check the network health status.

Step 1 In the alarm list on FusionInsight Manager, click the drop-down button of the alarm and view **Additional Information**. Record the source IP address and destination IP address of the node for which the alarm is reported.

Step 2 Log in to the node for which the alarm is reported. On the node, ping the target node to check whether the network between the two nodes is normal.

- If yes, go to **6**.
- If no, go to **3**.

Check the node status.

Step 3 On FusionInsight Manager, click **Host** and check whether the host list contains the faulty node to determine whether the faulty node has been removed from the cluster.

- If yes, go to **5**.
- If no, go to **4**.

Step 4 Check whether the faulty node is powered off.

- If yes, start the faulty node and go to **Step 2**.
- If no, contact related personnel to find root cause, if need to remove the faulty nodes from the cluster and go to **5**, otherwise go to **6**.

Step 5 Remove the file **$NODE_AGENT_HOME/etc/agent/hosts.ini** of all nodes in the cluster, and clean up the file **/var/log/Bigdata/unreachable/unreachable_ip_info.log**, and then manually clear the alarm.

Step 6 Wait for 30 seconds and checking if the alarm was been cleared.

- If yes, no further action is required.
- If no, go to **7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **OmmAgent** from the **Service** and click **OK**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.63 ALM-12091 Abnormal disaster Resources

Alarm Description

HA checks the disaster resources of Manager every 86 seconds. This alarm is generated when HA detects that the disaster resources have been abnormal for 10 consecutive times.

This alarm is cleared when HA detects that the disaster resources become normal.

Resource Type of disaster is **Single-active**. Active/Standby switchover will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new disaster resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12091	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- The active/standby Manager switchover occurs.
- The disaster process restarts repeatedly, which may cause active/standby DR to be unavailable.

Possible Causes

The disaster process is abnormal.

Handling Procedure

Check whether the disaster process is normal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the name of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su - omm** command to switch to user **omm**.

Step 4 Run the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the status of the disaster resources managed by the HA is normal. In the single-node system, the disaster resource is in the normal state. In the dual-node system, the disaster resource is in the normal state on the active node and in the stopped state on the standby node.

- If yes, go to **Step 7**.
- If no, go to **Step 5**.

Step 5 Run the **vi ${BIGDATA_LOG_HOME}/disaster/disaster.log** command to check whether the disaster resource log of HA contains the keyword **ERROR**. If yes, analyze the logs to locate the resource exception cause and fix the exception.

Step 6 Wait 5 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, select **Disaster** for the target cluster, and click **OK**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.64 ALM-12099 core dump Occurred

Description

GaussDB A manages the core file. It manages the lifecycle of core files generated when applications crash and manages alarm notification. This alarm is generated when a new core file is detected.

NOTE

This section applies to MRS 3.1.5 and later versions.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12099	Minor	No

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

This alarm indicates that some processes crashed. If a key process crashes, the cluster may be temporarily unavailable.

Possible Causes

Related processes crash.

Procedure

CAUTION

- Users' sensitive data may be involved in the following operations on parsing the core file stack information. Therefore, development or O&M personnel can perform these operations only after being authorized.
- The core file generated for the alarm is retained for 72 hours by default. If the period for storing the file exceeds 72 hours or the file size exceeds the preset value, the system automatically clears the file. Therefore, once this alarm is generated, contact O&M personnel in a timely manner.

Step 1 In the alarm list on the FusionInsight Manager page, click the row containing the alarm, and view the host IP address for which the alarm is generated in the alarm details. Then, view the path for storing the core file according to the **DumpedFilePath** attribute in the additional information.

Step 2 Log in to the host for which the alarm is generated as user **omm** and run the **gdb --version** command to check whether the gdb tool is installed on the host.

- If no, install the gdb tool and then go to **Step 3**.
- If yes, go to the **Step 3**.

Step 3 Use the gdb tool to view the stack details of the core file.

1. Go to the **DumpedFilePath** directory and find the core file.
2. Run the following commands to obtain the symbol table of the core file:

   source $BIGDATA_HOME/mppdb/.mppdbgs_profile

   cd ${BIGDATA_HOME}/FusionInsight_MPPDB_*XXX*/install/FusionInsight-MPPDB-*XXX*/package/MPPDB_ALL_PACKAGE

   tar -xzvf GaussDB-Kernel-V300R002C00-*Operating system*-64bit-symbol.tar.gz

cd symbols/bin/

Find the symbol table file whose name is the same as the process name in the alarm. For example, the symbol table file for the cm_agent process is **cm_agent.symbol**.

Copy the obtained symbol table to the **${GAUSSHOME}/bin** directory.

3. Run the **gdb --batch -n -ex thread -ex bt** *core file name* command to view the stack details of the core file.

Step 4 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system does not automatically clear this alarm, and you need to manually clear the alarm.

Related Information

None

10.65 ALM-12100 AD Service Connection Failed

Alarm Description

After a third-party active directory (AD) is interconnected, the third-party AD domain user can be synchronized using the synchronization period (60 minutes by default) or manually. During data synchronization, the AD service status will be checked. This alarm is generated when AD service unavailability is detected for three consecutive times. This alarm is cleared when AD service recovers.

NOTE

This section applies to MRS 3.1.5 and later versions.

Alarm Attributes

Alarm ID	Severity	Auto Clear
12100	Major	Yes

Alarm Parameters

Alarm Parameters	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the name of the service for which the alarm is generated.

Alarm Parameters	Description
RoleName	Specifies the name of the role for which the alarm is generated.
HostName	Specifies the name of the host for which the alarm is generated.

Impact on the System

When the alarm is generated, the AD service is unavailable, and AD domain user synchronization fails. An AD domain user cannot log in to FusionInsight Manager and execute services.

Possible Causes

- The configuration item for interconnecting with the third-party AD is incorrect.
- The network connection between FusionInsight and the third-party AD service is faulty.
- AD server fault
- AD service fault

Handling Procedure

Check the third-party AD configuration.

Step 1 On the **FusionInsight Manager** page, choose **System** > **Permission** > **Third-Party AD**. The third-party AD configuration page is displayed.

Step 2 Check whether the **AD IP Address**, **LDAP Port**, **Bind DN**, and **Bind DN Password** parameters are correctly set.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Modify the incorrect parameters, and then click **OK**.

Step 4 Choose **System** > **Permission** > **User** > **AD Domain User**, click **Manual Synchronization**, and check whether the message "Manual synchronization successfully." is displayed in the upper right corner of the page.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the third-party AD server and the network.

Step 5 Log in to the active management node as user **root**.

Step 6 On the host you have logged in to, ping the IP address of the third-party AD server to check whether the third-party AD server can be pinged.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 Run the following command to check whether the third-party AD service can be connected:

telnet *IP port*

IP indicates the IP address of the third-party AD server, and *port* indicates the port used by the third-party AD server.

- If yes, go to **Step 8**.
- If no, contact O&M personnel to check the network.

Step 8 Contact the third-party AD service administrator to check whether the AD service is normal.

- If yes, go to **Step 9**.
- If no, contact the third-party AD service administrator to rectify the AD server fault.

Collect the fault information.

Step 9 On the **FusionInsight Manager** page, choose **O&M** > **Log** > **Download**.

Step 10 In the **Service** area, select **Controller** under **OMS** and click **OK**.

Step 11 Click in the upper right corner to set **Start Date** and **End Date** to 10 minutes before and after the time when the alarm is generated, and click **Download**.

Step 12 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm. You do not need to manually clear it.

Reference

None

10.66 ALM-12101 AZ Unhealthy

Description

After the AZ DR function is enabled, the system checks the AZ health status every 5 minutes. This alarm is generated when the system detects that the AZ is subhealthy or unhealthy. This alarm is cleared when the AZ becomes healthy.

Attribute

Alarm ID	Alarm Severity	Auto Clear
12101	Critical	Yes

Parameters

Parameter	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
AZName	Specifies the AZ for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The health status of an AZ is determined by whether the health status of storage resources (HDFS), computing resources (Yarn), and key roles in the AZ exceeds the configured threshold.

An AZ is subhealthy when:

- The computing resources (Yarn) are unhealthy, but the storage resources (HDFS) are healthy. Tasks cannot be submitted to the local AZ, but data can still be read and written in the local AZ.
- The computing resources (Yarn) are healthy, but some storage resources (HDFS) are unhealthy. Tasks can be submitted to the local AZ, and some data can be read and written in the local AZ. This depends on the locality of data detected by Spark/Hive scheduling.

An AZ is unhealthy when:

- The computing resources (Yarn) are healthy, but the storage resources (HDFS) are unhealthy. Although tasks can be submitted to the local AZ, data cannot be read or written in the local AZ. As a result, the tasks submitted to the local AZ are invalid.
- The computing resources (Yarn) and storage resources (HDFS) are unhealthy. Tasks cannot be submitted to the local AZ, and data cannot be read or written in the local AZ.
- The health status of key roles except Yarn and HDFS is lower than the configured threshold.

Possible Causes

- The computing resources (Yarn) are unhealthy.
- The storage resources (HDFS) are unhealthy.
- Some storage resources (HDFS) are unhealthy.
- Key roles except Yarn and HDFS are unhealthy.

Procedure

Disable the DR drill.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Cross-AZ HA**. The Cross-AZ HA page is displayed.

Step 2 In the AZ DR list, check whether **Perform DR Drill** in the **Operation** column of the AZ whose health status is **Unhealthy** is gray.

- If yes, go to **Step 4**.
- If no, go to **Step 3**.

Step 3 Click **Restore** in the **Operation** column of the target AZ. Wait 2 minutes and refresh the page to view the health status of the AZ. Check whether the health status is normal.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect the fault information.

Step 4 Log in to the active management node as user **root**.

Step 5 View logs of unhealthy services.

- HDFS log files are stored in **/var/log/Bigdata/hdfs/nn/hdfs-az-state.log**.
- Yarn log files are stored in **/var/log/Bigdata/yarn/rm/yarn-az-state.log**.
- For other services, view the service health check logs in the corresponding service log directory.

Step 6 Contact O&M personnel and provide detailed log file information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.67 ALM-12102 AZ HA Component Is Not Deployed Based on DR Requirements

Alarm Description

The alarm module checks the deployment status of AZ HA components every 5 minutes. This alarm is generated when the components that support DR are not deployed based on DR requirements after AZ is enabled. This alarm is cleared when the components are deployed based on DR requirements.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12102	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Impact on the System

The cross-AZ HA capability of a single cluster is affected.

Possible Causes

The roles of the components that support DR are not deployed based on DR requirements.

Handling Procedure

Obtain alarm information.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, click ⌄ in the row that contains the alarm and view the roles that are not deployed based on DR requirements in **Additional Information**.

Redeploy the role instance.

Step 3 Choose **Cluster** > **Services** > *Name of the desired service* > **Instance**. On the instance page, redeploy or adjust the role instance.

Step 4 Check whether the alarm is cleared 10 minutes later.

- If yes, no further action is required.
- If no, contact O&M personnel.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.68 ALM-12103 Executor Resource Exception

Alarm Description

HA checks the Executor resources of Manager every 30 seconds. This alarm is generated when HA detects that the Executor resources are abnormal for two consecutive times.

This alarm is cleared when the Executor resources are normal.

Resource Type of Executor is **Single-active**. Active/standby will be triggered upon resource exceptions. When this alarm is generated, the active/standby switchover is complete and new Executor resources have been enabled on the new active Manager. In this case, this alarm is cleared. This alarm is used to notify users of the cause of the active/standby Manager switchover.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12103	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

- The active/standby Manager switchover occurs.
- The Executor process keeps restarting. As a result, the cluster page may fail to be accessed.

Possible Causes

The Executor process is abnormal.

Handling Procedure

Check whether the Executor process is abnormal.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the name of the host for which the alarm is generated.

Step 2 Log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su - omm** command to switch to user **omm**.

Step 4 Run the **sh ${BIGDATA_HOME}/om-server/OMS/workspace0/ha/module/hacom/script/status_ha.sh** command to check whether the status of the Executor resources managed by the HA is normal. In the single-node system, the Executor resource is in the normal state. In the dual-node system, the Executor resource is in the normal state on the active node and in the stopped state on the standby node.

- If yes, go to **Step 7**.
- If no, go to **Step 5**.

Step 5 Run the **vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/executor.log** command to check whether the Executor resource log of HA contains the keyword **ERROR**. If yes, analyze the log to locate the resource exception cause and fix the exception.

Step 6 After 5 minutes, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

10.69 ALM-12104 Abnormal Knox Resources

Alarm Description

HA checks the Knox resources of Manager every 70 seconds. This alarm is generated when HA detects that the Knox resources are abnormal for three consecutive times.

This alarm is cleared when HA detects that the Knox resources are normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12104	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Requests sent by upper-layer services by using Knox cannot be properly processed.

Possible Causes

The Knox process is abnormal.

Handling Procedure

Check whether the Knox process is normal.

Step 1 Log in to FusionInsight Manager. In the alarm list, locate the row that contains the alarm and view the name of the host for which the alarm is generated.

Step 2 Use PuTTY to log in to the host for which the alarm is generated as user **root**.

Step 3 Run the **su - omm** command to switch to user **omm**.

Step 4 Run the **sh ${BIGDATA_HOME}/om-server/om/sbin/status-oms.sh** command to check whether the status of the Knox resources managed by HA is normal. If the status is normal, the Knox resources are normal. Otherwise, the Knox resources are abnormal.

- If yes, go to **Step 7**.
- If no, go to **Step 5**.

Step 5 Run the **vi $BIGDATA_LOG_HOME/omm/oms/ha/scriptlog/knox.log** command to check whether the Knox resource log of HA contains the keyword **ERROR**. If yes, analyze the log to locate the resource exception cause and fix the exception.

Step 6 After 5 minutes, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 In the **Services** area, select **Controller** and **OmmServer**, and click **OK**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

10.70 ALM-12110 Failed to get ECS temporary AK/SK

Alarm Description

Meta calls the ECS API to obtain the AK/SK information every 5 minutes and caches the information. Before the AK/SK expires, Meta calls the API again to update it. This alarm is generated when Meta fails to call the API for three consecutive times.

This alarm is cleared when Meta successfully calls the ECS API.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12110	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

In storage-compute decoupling scenarios, the cluster cannot obtain the latest temporary AK/SK, which may lead to failure access to OBS.

Possible Causes

- The meta role of the MRS cluster is abnormal.
- The cluster has been bound to an agency and accessed OBS but has been unbound from the agency. As a result, the cluster has not been bound to any agency.

Handling Procedure

Check the status of the meta role.

Step 1 On FusionInsight Manager of the cluster, choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click ⌄ in the row containing the alarm, and determine the IP address of the host for which the alarm is generated.

Step 2 On FusionInsight Manager of the cluster, choose **Cluster** > **Services** > **meta**. On the page that is displayed, click the **Instances** tab, and check whether the meta role corresponding to the host for which the alarm is generated is normal.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Select the abnormal role, click **More**, and select **Restart Instance** to restart the abnormal meta role.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Log in to the host obtained in **Step 1** and check whether the **/var/log/Bigdata/meta/mrs-meta.log** file contains error information. If yes, rectify the fault based on the log information.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Rebind the cluster to an agency.

Step 7 Log in to the MRS management console.

Step 8 In the navigation pane on the left, choose **Active Clusters**. On the page that is displayed, click the cluster name to go to its overview page. Then, check whether the cluster is bound to an agency in the O&M management area.

- If yes, go to **Step 10**.
- If no, go to **Step 9**.

Step 9 Click **Manage Agency**. On the page that is displayed, rebind the cluster to an agency. Then check whether the alarm is cleared a few minutes later.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager of the active cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, select **meta** for the target cluster, and click **OK**.

Step 12 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

10.71 ALM-12172 Failed to Report Metrics to Cloud Eye

Alarm Description

After metric sharing is enabled for a cluster, the Controller periodically collects cluster metrics and reports them to Cloud Eye.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12172	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

MRS monitoring metrics are unavailable on Cloud Eye.

Possible Causes

- Failed to call Cloud Eye APIs due to insufficient permissions.
- Failed to report data to Cloud Eye due to network problems.
- Failed to report data to Cloud Eye due to internal errors.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click ⌄ in the row containing the alarm and view the additional information of the alarm.

Step 2 Rectify the fault based on the following scenarios:

- If "Call CES to send metrics fail. Permission exception" is displayed in the additional information, the token of the resource tenant is invalid. Restart the Controller and obtain the token again.
- If "Call CES to send metrics fail. Request CES error code xxx" is displayed in the additional information, an error occurs in the request to Cloud Eye. Check the network connectivity and authentication information.
- If "Call CES to send metrics fail. CES internal error code xxx" is displayed in the additional information, the Cloud Eye service encounters an internal error and is unavailable. Contact O&M personnel and send collected fault logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.72 ALM-12180 Suspended Disk I/O

Alarm Description

For MRS 3.3.0 and its later versions:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency reaches 6 seconds within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 1000 ms within 300 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:

  - By default, the system collects data every 3 seconds. The svctm latency reaches 3 seconds within 30 seconds in at least seven collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.
  - By default, the system collects data every 3 seconds. At least 50% of detected svctm take no less than 500 ms within 300 seconds.

The collection period is 3 seconds, and the detection period is 30 or 300 seconds. This alarm is automatically cleared when neither of the preceding conditions is met for three consecutive detection periods (30 or 300 seconds).

For versions earlier than MRS 3.3.0:

- For HDDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency exceeds 6 seconds within 30 seconds in at least 10 collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.
- For SSDs, the alarm is triggered when any of the following conditions is met:
  - By default, the system collects data every 3 seconds. The svctm latency exceeds 2 seconds within 30 seconds in at least 10 collection periods.
  - By default, the system collects data every 3 seconds. The disk queue depth (**avgqu-sz**) is greater than 0, the IOPS or bandwidth is 0, and **ioutil** is greater than 99% in at least 10 collection periods within 30 seconds.

This alarm is automatically cleared when the preceding conditions have not been met for 90s.

NOTE

For details about how to obtain and calculate related parameters, see **Related Information**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12180	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
DiskName	Specifies the disk for which the alarm was generated.

Impact on the System

A continuously high I/O usage may adversely affect service operations and result in service loss.

Possible Causes

The disk is aged.

Handling Procedure

Replace the disk.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**.

Step 2 View the detailed information about the alarm. Check the values of **HostName** and **DiskName** in the location information to obtain the information about the faulty disk for which the alarm is reported.

Step 3 Replace the hard disk.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Select **OMS** for **Service** and click **OK**.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

Obtain and calculate related parameters as follows:

- Run the following command in the OS to collect data:

  iostat -x -t 1 1

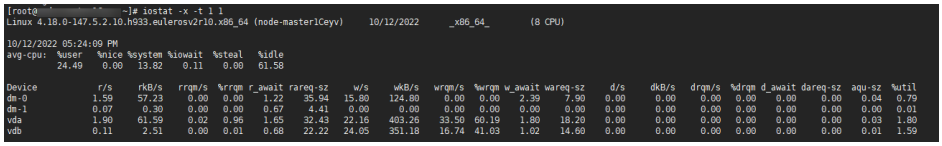

  Parameters are as follows:

  avgqu-sz indicates the disk queue depth.

  The sum of **r/s** and **w/s** is the IOPS.

  The sum of **rkB/s** and **wkB/s** is the bandwidth.

  %util is the **ioutil** value.

- MRS 3.1.0:

  Run the **iostat -x -t** command in the OS.

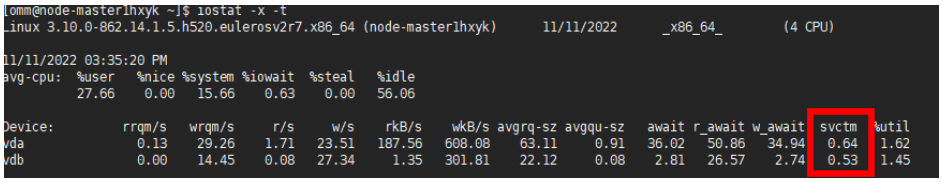

- Calculate **svctm** as follows in versions later than MRS 3.1.0:

  svctm = (tot_ticks_new - tot_ticks_old) / (rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old)

- Versions earlier than MRS 3.3.0: If **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, then **svctm = 0**.
- MRS 3.3.0 and its later versions:

  When the detection period is 30 seconds, if **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, then **svctm = 0**.

  When the detection period is 300 seconds and **rd_ios_new + wr_ios_new - rd_ios_old - wr_ios_old = 0**, if **tot_ticks_new - tot_ticks_old = 0**, then **svctm = 0**; otherwise, the value of **svctm** is infinite.

The parameters can be obtained as follows:

The system runs the **cat /proc/diskstats** command every 3 seconds to collect data. For example:

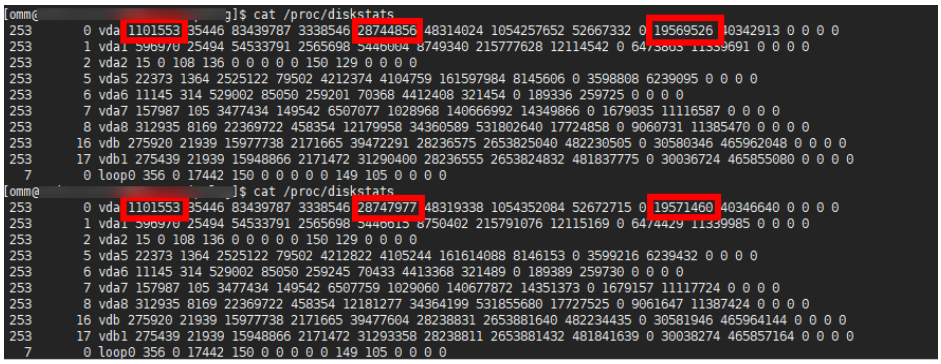

In the two commands:

In the data collected for the first time, the number in the fourth column is the **rd_ios_old** value, the number in the eighth column is the **wr_ios_old** value, and the number in the thirteenth column is the **tot_ticks_old** value.

In the data collected for the second time, the number in the fourth column is the **rd_ios_new** value, the number in the eighth column is the **wr_ios_new** value, and the number in the thirteenth column is the **tot_ticks_new** value.

In this case, the value of **svctm** is as follows:

(19571460 - 19569526) / (1101553 + 28747977 - 1101553 - 28744856) = 0.6197

10.73 ALM-12186 CGroup Task Usage Exceeds the Threshold

Alarm Description

The system checks the CGroup task usage of user **omm** every 5 minutes. This alarm is generated when the CGroup task usage exceeds 90%. This alarm is cleared when the CGroup task usage is less than or equal to 90%.

CGroup task usage = Number of used CGroup tasks/Maximum number of CGroup tasks

You can run the **systemctl status user-$(id -u).slice | grep limit | awk -F ' ' '{print $2}'** command as user **omm** to obtain the number of used CGroup tasks of this user and run the **echo $(systemctl status user-$(id -u).slice | grep limit | awk -F ' ' '{print $4}') | sed -e 's/)//g'** command to obtain the maximum number of CGroup tasks allowed for this user.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12186	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

- Failed to switch to user **omm**.
- Failed to create new **omm** processes.
- A faulty service or process cannot be restarted.

Possible Causes

The CGroup task usage exceeds 90%.

Handling Procedure

Check the maximum number of threads that can be concurrently opened by user omm is properly set.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click ⌄ in the row containing the alarm, and view the name of the host for which the alarm is generated in **Location**. Click the host name to view its IP address.

Step 2 Log in to the host for which the alarm is generated as user **omm**.

Step 3 Run the following command to obtain the maximum number of threads that can be concurrently opened by user **omm** and check whether this number is greater than or equal to **60000**:

systemctl status user-$(id -u).slice | grep limit

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Switch to user **root** and run the following command to change the value for user **omm** to **60000**:

systemctl set-property user-2000.slice TasksMax=60000

Step 5 Change the value of **UserTasksMax** in the **/etc/systemd/logind.conf** file to **60000**. (If the parameter is commented out, uncomment it.) Save the file, wait 5 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager of the cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, select **OmmServer** and **NodeAgent** for the target cluster, and click **OK**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.74 ALM-12187 Failed to Expand Disk Partition Capacity

Alarm Description

The system checks the disk space every 60 seconds. When detecting that the disk space is expanded, the system expands disk partition. This alarm is generated when the disk partition fails to be expanded.

This alarm is cleared when the system detects that the disk partition is successfully expanded.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12187	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
MountDirectoryName	Specifies the directory for which the alarm is generated.

Impact on the System

The expanded data disk space cannot be used to store data.

Possible Causes

- The growpart scale-out tool is not installed.
- The system fails to execute the command for expanding disk partition.

Handling Procedure

Check whether growpart is installed.

Step 1 Log in to FusionInsight Manager, click **O&M**, and choose **Alarm** > **Alarms** to view the alarm details. In the **Location** column, check the name of the host and mount directory for which the alarm is generated. Click the host name to view its IP address.

Step 2 Log in to the node for which the alarm is generated as user **root**.

Step 3 Run the following command to check whether growpart is installed:

which growpart

If information similar to the following is displayed, the growpart tool is installed. Otherwise, contact O&M personnel to install the growpart tool.

```
[root@xxx ~]#which growpart
/usr/bin/growpart
```

Step 4 Wait for 5 minutes, then choose **O&M**, and choose **Alarm** > **Alarms** on FusionInsight Manager. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Run the disk partition expansion command.

Step 5 Run the following command to view the disk and partition information:

lsblk

Search for the partition and the disk based on the mount directory name in the alarm location information, and check the disk and partition sizes.

In the following example, the mount directory is **/srv/BigData/data1**, the used disk is **/dev/vdb**, and the disk partition is **/dev/vdb1**.

```
[root@                    ~]# lsblk
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
loop0    7:0    0  1.5G  0 loop /media
vda    253:0    0  500G  0 disk
├─vda1 253:1    0  220G  0 part /
├─vda2 253:2    0    1K  0 part
├─vda5 253:5    0   10G  0 part /tmp
├─vda6 253:6    0   10G  0 part /var
├─vda7 253:7    0   60G  0 part /srv/BigData
└─vda8 253:8    0  180G  0 part /var/log
vdb    253:16   0  650G  0 disk
└─vdb1 253:17   0  600G  0 part /srv/BigData/data1
```

Step 6 Run the following command to expand the partition using growpart:

growpart *Data disk Partition number*

Run the following command:

growpart /dev/vdb 1

If information similar to the following is displayed, the execution is successful. If the execution fails, contact O&M personnel.

```
[root@host-                ~]# growpart /dev/vdb 1
CHANGED: partition=1 start=2048 old: size=1258287104 end=1258289152 new: size=1363146719 end=1363148767
```

Step 7 Run the following command to expand the file system size of the disk partition:

resize2fs *Disk partition*

Run the following command:

resize2fs /dev/vdb1

If information similar to the following is displayed, the execution is successful:

```
[root@host-                ~]# resize2fs /dev/vdb1
resize2fs 1.46.4 (18-Aug-2021)
Filesystem at /dev/vdb1 is mounted on /srv/BigData/data1; on-line resizing required
old_desc_blocks = 75, new_desc_blocks = 82
The filesystem on /dev/vdb1 is now 170393339 (4k) blocks long.
```

Step 8 Wait for 5 minutes, click **O&M**, and choose **Alarm** > **Alarms** on FusionInsight Manager. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, contact O&M personnel.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.75 ALM-12188 diskmgt Disk Monitoring Unavailable

Alarm Description

NodeAgent checks the status of the diskmgt disk monitoring service every 5 minutes. This alarm is generated when diskmgt disk monitoring is unavailable.

This alarm is cleared when the diskmgt disk monitoring service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12188	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

When diskmgt disk monitoring is unavailable, the read-only detection of the device partition file system, device partition loss detection, and disk partition scale-out detection cannot be performed.

Possible Causes

- The diskmgt disk monitoring service does not exist.
- The diskmgt disk monitoring service is not started.

Handling Procedure

Check whether the diskmgt disk monitoring service exists.

Step 1 Log in to FusionInsight Manager, click **O&M**, and choose **Alarm** > **Alarms** to view the alarm details. In the **Location** column, check the name of the host for which the alarm is generated. Click the host name to view its IP address.

Step 2 Log in to the node for which the alarm is generated as user **root**.

Step 3 Run the following command to check whether the core service file exists:

stat /usr/local/diskmgt/inner/diskmgtd

If the file does not exist, contact O&M personnel.

Start the diskmgt disk monitoring service.

Step 4 Run the following command to start the diskmgt disk monitoring service:

systemctl restart diskmgt

Step 5 Run the following command to check whether the diskmgt disk monitoring service is started:

systemctl status diskmgt

- If information similar to the following is displayed, the service is started successfully. Go to **Step 6**.

```
[root@host·                ~]# systemctl status diskmgt
● diskmgt.service - Disk monitor service
     Loaded: loaded (/usr/lib/systemd/system/diskmgt.service; enabled; vendor preset: disabled)
     Active: active (running) since Thu 2023-08-10 09:29:18 CST; 5 days ago
   Main PID: 33996 (diskmgtd)
      Tasks: 2 (limit: 407663)
     Memory: 5.6M
     CGroup: /system.slice/diskmgt.service
             ├─  33996 /bin/bash /usr/local/diskmgt/inner/diskmgtd
             └─1974506 sleep 8

Notice: journal has been rotated since unit was started, output may be incomplete.
```

- If no, contact O&M personnel.

Step 6 Wait for 5 minutes, click **O&M**, and choose **Alarm** > **Alarms** on FusionInsight Manager. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, contact O&M personnel.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.76 ALM-12190 Number of Knox Connections Exceeds the Threshold

Alarm Description

The system periodically checks the number of connections to all Knox topologies. This alarm is generated when the number of connections to a topology exceeds the threshold (90% by default). This alarm is automatically cleared when the number of connections to a topology falls below the threshold.

NOTE

This alarm applies to clusters of MRS 3.1.0 or later.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
12190	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Topology	Specifies the Knox topology for which the alarm was generated.

Impact on the System

The topology may reach the upper limit of connections and fail to forward requests, adversely affecting the MRS functions.

Possible Causes

Hue or Manager is too frequently used, but the default maximum number of Knox connections is small.

Handling Procedure

Step 1 Log in to active and standby OMS nodes as user **root**, respectively.

Step 2 Add the following configuration to the **gateway-site.xml** file on the active and standby OMS nodes to increase the number of thread pools:

vi /opt/knox/conf/gateway-site.xml

```
<property>
<name>gateway.httpclient.maxConnections</name>
<value>64</value>
</property>
```

Step 3 Log in to the active OMS node as user **omm** and run the following command to restart the Knox process:

sh /opt/knox/bin/restart-knox.sh

Step 4 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Contact O&M personnel to rectify the fault.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.77 ALM-13000 ZooKeeper Service Unavailable

Description

The system checks the ZooKeeper service status every 60 seconds. This alarm is generated when the ZooKeeper service is unavailable.

This alarm is cleared when the ZooKeeper service recovers.

Attribute

Alarm ID	Alarm Severity	Auto Clear
13000	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

ZooKeeper cannot provide coordination services for upper layer components and the components that depend on ZooKeeper may not run properly.

Possible Causes

- The DNS is installed on the ZooKeeper node.
- The network is faulty.
- The KrbServer service is abnormal.
- The ZooKeeper instance is abnormal.
- The disk capacity is insufficient.

Procedure

Check the DNS.

Step 1 Check whether the DNS is installed on the node where the ZooKeeper instance is located. On the Linux node where the ZooKeeper instance is located, run the **cat /etc/resolv.conf** command to check whether the file is empty.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 Run the **service named status** command to check whether the DNS is started.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Run the **service named stop** command to stop the DNS service. If "Shutting down name server BIND waiting for named to shut down (28s)" is displayed, the DNS service is stopped successfully. Comment out the content (if any) in **/etc/resolv.conf**.

Step 4 On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 5**.

Check the network status.

Step 5 On the Linux node where the ZooKeeper instance is located, run the **ping** command to check whether the host names of other nodes where the ZooKeeper instance is located can be pinged successfully.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

Step 6 Modify the IP addresses in **/etc/hosts** and add the host name and IP address mapping.

Step 7 Run the **ping** command again to check whether the host names of other nodes where the ZooKeeper instance is located can be pinged successfully.

- If yes, go to **Step 8**.
- If no, go to **Step 23**.

Step 8 On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check the KrbServer service status (Skip this step if the normal mode is used).

Step 9 On FusionInsight Manager, choose **Cluster >** *Name of the desired cluster* **> Services**.

Step 10 Check whether the KrbServer service is normal.

- If yes, go to **Step 13**.
- If no, go to **Step 11**.

Step 11 Perform operations based on "ALM-25500 KrbServer Service Unavailable" and check whether the KrbServer service is recovered.

- If yes, go to **Step 12**.
- If no, go to **Step 23**.

Step 12 On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Check the ZooKeeper service instance status.

Step 13 On FusionInsight Manager, choose **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper** > **quorumpeer**.

Step 14 Check whether the ZooKeeper instances are normal.

- If yes, go to **Step 18**.
- If no, go to **Step 15**.

Step 15 Select instances whose status is not good, and choose **More** > **Restart Instance**.

Step 16 Check whether the instance status is good after restart.

- If yes, go to **Step 17**.
- If no, go to **Step 18**.

Step 17 On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Check disk status.

Step 18 On FusionInsight Manager, choose **Cluster >** *Name of the desired cluster* **> Service** > **ZooKeeper** > **quorumpeer**, and check the node host information of the ZooKeeper instance.

Step 19 On FusionInsight Manager, click **Host**.

Step 20 In the **Disk** column, check whether the disk space of each node where ZooKeeper instances are located is insufficient (disk usage exceeds 80%).

- If yes, go to **Step 21**.
- If no, go to **Step 23**.

Step 21 Expand disk capacity. For details, see "ALM-12017 Insufficient Disk Capacity".

Step 22 On the **O&M > Alarm > Alarms** tab, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 23**.

Collect fault information.

Step 23 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 24 Select the following nodes in the required cluster from the **Service**: (KrbServer logs do not need to be downloaded in normal mode.)

- ZooKeeper
- KrbServer

Step 25 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 26 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.78 ALM-13001 Available ZooKeeper Connections Are Insufficient

Description

The system checks ZooKeeper connections every 60 seconds. This alarm is generated when the system detects that the number of used ZooKeeper instance connections exceeds the threshold (80% of the maximum connections).

When the **Trigger Count** is 1, this alarm is cleared when the number of used ZooKeeper instance connections is smaller than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the number of used ZooKeeper instance connections is smaller than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
13001	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host name for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Available ZooKeeper connections are insufficient. When the connection usage reaches 100%, external connections cannot be handled.

Possible Causes

The number of connections to the ZooKeeper node exceeds the threshold. Connection leakage occurs on some connection processes, or the maximum number of connections does not comply with the actual scenario.

Procedure

Check connection status.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms**. On the displayed interface, click the drop-down button of **Available ZooKeeper Connections Are Insufficient** and confirm the node IP address of the host for which the alarm is generated in the Location Information.

Step 2 Obtain the PID of the ZooKeeper process. Log in to the node involved in this alarm as user **root** and run the **pgrep -f proc_zookeeper** command.

Step 3 Check whether the PID can be correctly obtained.

- If yes, go to **Step 4**.
- If no, go to **Step 15**.

Step 4 Obtain all the IP addresses connected to the ZooKeeper instance and the number of connections and check 10 IP addresses with top connections. Run the following command based on the obtained PID: **lsof -i|grep** *$pid* **| awk '{print $9}' | cut -d : -f 2 | cut -d \>-f 2 | awk '{a[$1]++} END {for(i in a){print i,a[i] | "sort -r -g -k 2"}}' | head -10**. (The PID obtained in the preceding step is used.)

Step 5 Check whether node IP addresses and number of connections are successfully obtained.

- If yes, go to **Step 6**.
- If no, go to **Step 15**.

Step 6 Obtain the ID of the port connected to the process. Run the following command based on the obtained PID and IP address: **lsof -i|grep** *$pid* **| awk '{print $9}'|cut -d \> -f 2 |grep** *$IP*| **cut -d : -f 2**. (The PID and IP address obtained in the preceding step are used.)

Step 7 Check whether the port ID is successfully obtained.

- If yes, go to **Step 8**.
- If no, go to **Step 15**.

Step 8 Obtain the ID of the connected process. Log in to each IP address and run the following command based on the obtained port ID: **lsof -i|grep** *$port*. (The port ID obtained in the preceding step is used.)

Step 9 Check whether the process ID is successfully obtained.

- If yes, go to **Step 10**.
- If no, go to **Step 15**.

Step 10 Check whether connection leakage occurs on the process based on the obtained process ID.

- If yes, go to **Step 11**.

- If no, go to **Step 12**.

Step 11 Close the process where connection leakage occurs and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Step 12 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **quorumpeer** > **Performance** and increase the value of **maxCnxns** as required.

Figure 10-14 maxCnxns

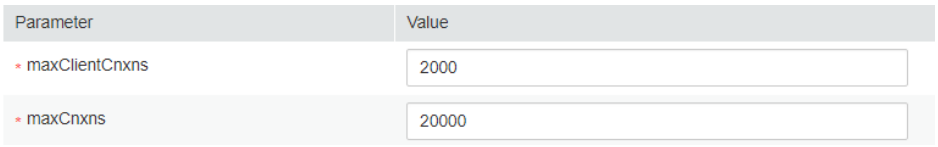

Step 13 Save the configuration and restart the ZooKeeper service.

Step 14 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Collect fault information.

Step 15 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 16 Select **ZooKeeper** in the required cluster from the **Service**:

Step 17 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact the O&M personnel and send the collected log information.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.79 ALM-13002 ZooKeeper Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the ZooKeeper service every 30 seconds. The alarm is generated when the direct memory usage of a ZooKeeper instance exceeds the threshold (80% of the maximum memory).

When the **Trigger Count** is 1, this alarm is cleared when the ZooKeeper Direct memory usage is less than the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the ZooKeeper Direct memory usage is less than 80% of the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
13002	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available direct memory of the ZooKeeper service is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the ZooKeeper instance is overused or the direct memory is inappropriately allocated.

Procedure

Check the direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms**. On the displayed interface, click the drop-down button of **ZooKeeper Direct Memory Usage Exceeds the Threshold**. Check the IP address of the instance that reports the alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper** > **Instance** > **quorumpeer(the IP address checked)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **CPU and Memory**, and select **ZooKeeper Heap And Direct Buffer Resource Percentage**, click **OK**.

Step 3 Check whether the used direct buffer memory of ZooKeeper reaches 80% of the maximum direct buffer memory specified for ZooKeeper.

- If yes, go to **Step 4**.
- If no, go to **Step 8**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **quorumpeer** > **System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC_OPTS** parameter.

- If yes, in the **GC_OPTS** parameter, delete "-XX:MaxDirectMemorySize" and go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Save the configuration and restart the ZooKeeper service.

Step 6 Check whether the **ALM-13004 ZooKeeper Heap Memory Usage Exceeds the Threshold** exists.

- If yes, handle the alarm by referring to **ALM-13004 ZooKeeper Heap Memory Usage Exceeds the Threshold**.
- If no, go to **Step 7**.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 9 Select **ZooKeeper** in the required cluster from the **Service**.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.80 ALM-13003 GC Duration of the ZooKeeper Process Exceeds the Threshold

Alarm Description

The system checks the garbage collection (GC) duration of the ZooKeeper process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
13003	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A long GC duration of the ZooKeeper process may interrupt the services.

Possible Causes

The heap memory of the ZooKeeper process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed page, click the drop-down list of **GC Duration of the ZooKeeper Process Exceeds the Threshold**. View the IP address of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Instance** > **quorumpeer**. Click the drop-down list in the upper right corner of **Chart**, choose **Customize** > **GC**, select **ZooKeeper GC Duration per Minute**, and click **OK** to check the GC duration statistics of the ZooKeeper process collected every minute.

Figure 10-15 ZooKeeper GC duration

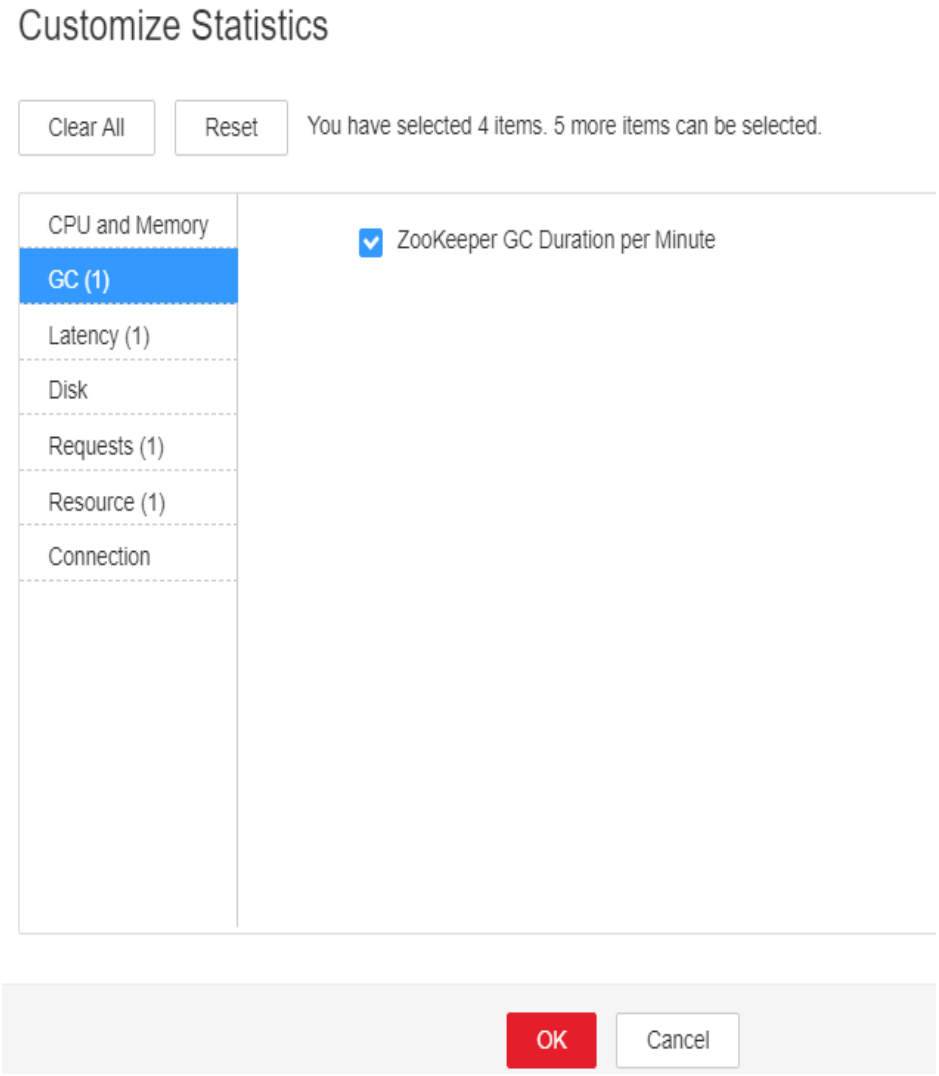

Step 3 Check whether the GC duration of the ZooKeeper process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 8**.

Step 4 Check whether memory leakage occurs in the application.

Step 5 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **quorumpeer** > **System**. Increase the value of the **GC_OPTS** parameter as required.

NOTE

Generally, **-Xmx** is twice of ZooKeeper data capacity. If the capacity of ZooKeeper reaches 2 GB, set **GC_OPTS** as follows:

-Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=512M -XX:MetaspaceSize=64M -XX:MaxMetaspaceSize=64M -XX:CMSFullGCsBeforeCompaction=1

Step 6 Save the configuration and restart the ZooKeeper service.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **ZooKeeper** for the target cluster.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.81 ALM-13004 ZooKeeper Heap Memory Usage Exceeds the Threshold

Description

The system checks the heap memory usage of the ZooKeeper service every 60 seconds. The alarm is generated when the heap memory usage of a ZooKeeper instance exceeds the threshold (95% of the maximum memory).

The alarm is cleared when the memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
13004	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available ZooKeeper heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the ZooKeeper instance is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, On the displayed interface, click the drop-down button of **ZooKeeper Heap Memory Usage Exceeds the Threshold** and confirm the node IP address of the host for which the alarm is generated in the Location Information.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper** > **Instance**, click **quorumpeer** in the **Role** column of the corresponding IP address. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **CPU and Memory**, and select **ZooKeeper Heap And Direct Buffer Resource Percentage**, click **OK**. Check the heap memory usage.

Step 3 Check whether the used heap memory of ZooKeeper reaches 95% of the maximum heap memory specified for ZooKeeper.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **quorumpeer** > **System**. Increase the value of **-Xmx** in **GC_OPTS** as required. The details are as follows:

1. On the **Instance** tab, click **quorumpeer** in the **Role** column of the corresponding IP address. Choose **Customize** > **CPU and Memory** in the

upper right corner, and select **ZooKeeper Heap And Direct Buffer Resource**, click **OK** to check the heap memory used by ZooKeeper.

2. Change the value of **-Xmx** in the **GC_OPTS** parameter based on the actual heap memory usage. Generally, the value is twice the size of the ZooKeeper data volume. For example, if 2 GB ZooKeeper heap memory is used, the following configurations are recommended: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=512M -XX:MetaspaceSize=64M -XX:MaxMetaspaceSize=64M -XX:CMSFullGCsBeforeCompaction=1

Step 5 Save the configuration and restart the ZooKeeper service.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **ZooKeeper** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.82 ALM-13005 Failed to Set the Quota of Top Directories of ZooKeeper Components

Description

The system sets quotas for each ZooKeeper top-level directory in the **customized.quota** configuration item and components every 5 hours. This alarm is generated when the system fails to set the quota for a directory.

This alarm is cleared when the setting succeeds after a failure.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13005	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
ServiceDirectory	Specifies the directory for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

Impact on the System

Components can write a large amount of data to the top-level directory of ZooKeeper. As a result, the ZooKeeper service is unavailable.

Possible Causes

The quota for the alarm directory is inappropriate.

Procedure

Check whether the quota for the alarm directory is appropriate.

Step 1 Log in to FusionInsight Manager, and choose **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper**. On the displayed page, choose **Configurations** > **All Configurations** > **Quota**. Check whether the directory for which the alarm is reported and its quota exist in the **customized.quota** configuration item.

- If yes, go to **Step 5**.
- If no, go to **Step 2**.

Step 2 Check whether the alarm directory for which the alarm is reported is in the following alarm list.

Table 10-1 Component alarm directory

Component	Alarm Directory
Hbase	/hbase

Component	Alarm Directory
Hive	/beelinesql
Yarn	/rmstore
Storm	/stormroot
Streaming	/storm
Kafka	/kafka

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Step 3 View the component of the alarm directory in the table, open the corresponding service page, and choose **Configurations** > **All Configurations**. On the displayed page, search for **zk.quota** in the upper right corner. The search result is the quota of the alarm directory.

Step 4 Check whether the quota of the alarm directory for which the alarm is reported is appropriate. The quota must be greater than or equal to the actual value, which can be obtained in **Trigger Condition**.

Step 5 Modify the **services.quota** value as prompted and save the configuration.

Step 6 After the time specified by **service.quotas.auto.check.cron.expression**, check whether the alarm is cleared.

The **service.quotas.auto.check.cron.expression** parameter indicates the scheduled expression used by ZooKeeper to set the directory quota. You can choose **Cluster > Services > ZooKeeper > Configurations > All Configurations** on FusionInsight Manager and set this parameter. The default value is */5 * * * *, indicating 5 minutes.

- If it is, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **ZooKeeper** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.83 ALM-13006 Znode Number or Capacity Exceeds the Threshold

Description

The system periodically detects the status of secondary Znode in the ZooKeeper service data directory every four hours. This alarm is generated when the number or capacity of secondary Znodes exceeds the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13006	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
ServiceDirectory	Specifies the directory for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

Impact on the System

A large amount of data is written to the ZooKeeper data directory. As a result, ZooKeeper cannot provide normal services.

Possible Causes

A large amount of data is written to the ZooKeeper data directory. The threshold is not appropriate.

Procedure

Check whether a large amount of data is written to the directory for which the alarm is generated.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed interface, click the drop-down button of **Znode Number or Capacity Exceeds the Threshold**. Confirm the Znode for which the alarm is generated in Location Information.

Step 2 Log in to FusionInsight Manager, open the ZooKeeper service interface, and select **Resource**. In the table **Used Resources (By Second-Level Znode)**, check whether a large amount of data is written to the top-level Znode for which the alarm is reported.

- If it is, go to **Step 3**.
- If it is not, go to **Step 4**.

Figure 10-16 Used Resources (By Second-Level Znode)

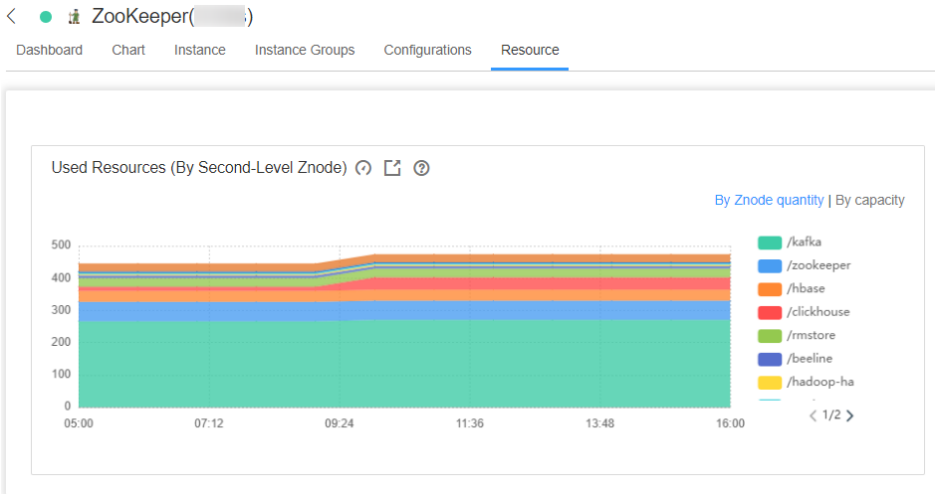

Step 3 Log in to the ZooKeeper client and delete the data in the top-level Znode.

Step 4 Log in to FusionInsight Manager and open the ZooKeeper service interface. On the **Resource** page, choose > **By Znode quantity** in **Used Resources (By Second-Level Znode)**. **Threshold Configuration of By Znode quantity** is displayed. Click **Modify** under **Operation**. Increase the threshold by referringto the value of **max.Znode.count** by choosing **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **Quota**.

Figure 10-17 Modify Rule

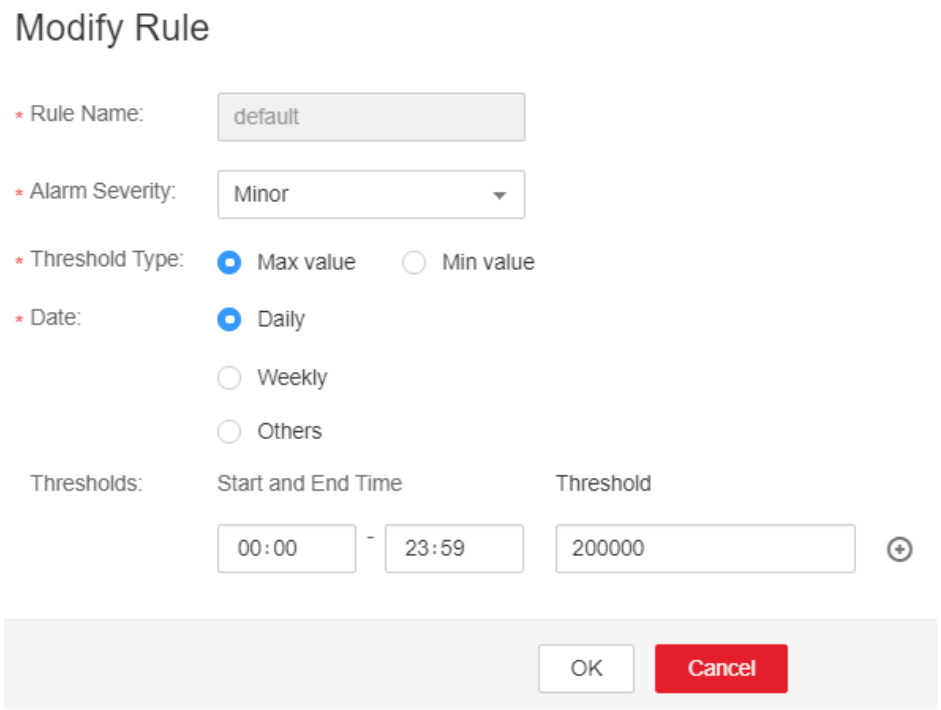

Step 5 In the **Used Resources (By Second-Level Znode)**, choose > **By capacity**. The **Threshold Settings** page of **By Capacity** is displayed. Click **Modify** under **Operation**. Increase the threshold by referring to the value of **max.data.size** by choosing **Cluster >** *Name of the desired cluster* **> Services** > **ZooKeeper** > **Configurations > All Configurations** > **Quota**.

Step 6 Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **ZooKeeper** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.84 ALM-13007 Available ZooKeeper Client Connections Are Insufficient

Description

The system periodically detects the number of active processes between the ZooKeeper client and the ZooKeeper server every 60 seconds. This alarm is generated when the number of connections exceeds the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13007	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the host name for which the alarm is generated.
ClientIP	Specifies the client IP address.
ServerIP	Specifies the server IP address.
Trigger Condition	Specifies the cause of the alarm.

Impact on the System

A large number of connections to ZooKeeper caused the ZooKeeper to be fully connected and unable to provide normal services.

Possible Causes

A large number of client processes are connected to ZooKeeper. The thresholds are not appropriate.

Procedure

Check whether there are a large number of client processes connected to ZooKeeper.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed interface, click the drop-down button of **Available ZooKeeper Client Connections Are Insufficient**. Confirm the node IP address of the host for which the alarm is generated in the Location Information.

Step 2 Open the ZooKeeper service interface, click **Resource** to enter the **Resource** page, and check whether the number of connections of the client with the IP address specified by **Number of Connections (By Client IP Address)** is large.

- If it is, go to **Step 3**.
- If it is not, go to **Step 4**.

Step 3 Check whether connection leakage occurs on the client process.

Step 4 Click in the **Number of Connections (by Client IP Address)** to enter the **Thresholds** page, and click **Modify** under **Operation**. Increase the threshold by referring to the value of **maxClientCnxns** by choosing **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Configurations** > **All Configurations** > **quorumpeer**.

Step 5 Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **ZooKeeper** in the required cluster from the **Service**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.85 ALM-13008 ZooKeeper Znode Usage Exceeds the Threshold

Description

The system checks the level-2 Znode status in the ZooKeeper data directory every hour. This alarm is generated when the system detects that the level-2 Znode usage exceeds the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13008	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
ServiceDirectory	Specifies the directory for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

Impact on the System

A large amount of data is written to the ZooKeeper data directory. As a result, ZooKeeper cannot provide services properly.

Possible Causes

- A large amount of data is written to the ZooKeeper data directory.
- The user-defined threshold is inappropriate.

Procedure

Check whether a large amount of data is written into the directory for which the alarm is generated.

Step 1 Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper**, and click **Resource**. Click **By Znode quantity** in **Used Resources (By Second-Level Znode)**, and check whether a large amount of data is written to the top Znode.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, select **Location** from the drop-down list box next to **ALM-13008 ZooKeeper Znode Quantity Usage Exceeds Threshold**, and obtain the Znode path in **ServiceDirectory**.

Step 3 Log in to the ZooKeeper client as a cluster user and delete unnecessary data from the Znode corresponding to the alarm.

Step 4 Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** > **Configurations** > **All Configurations**, and search for **max.znode.count**, which is the maximum number of ZooKeeper directories. The alarm threshold is 80% of this parameter. Increase the value of this parameter, click **Save**, and restart the service for the configuration to take effect.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 7 Select **ZooKeeper** in the required cluster from the **Service**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.86 ALM-13009 ZooKeeper Znode Capacity Usage Exceeds the Threshold

Alarm Description

The system checks the level-2 ZNode status in the ZooKeeper data directory every hour. This alarm is generated when the system detects that the capacity usage exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
13009	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
ServiceDirectory	Specifies the directory for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A large amount of data is written to the ZooKeeper data directory. As a result, ZooKeeper cannot provide services properly.

Possible Causes

- A large volume of data has been written to the ZooKeeper data directory.
- The threshold is improperly defined.

Handling Procedure

Check whether a large volume of data is written to the alarm directory.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. Click the drop-down list in the row containing **ALM-13009 ZooKeeper ZNode Capacity Usage Exceeds the Threshold**, and find the ZNode for which the alarm is generated in the **Location** area.

Step 2 Choose **Cluster** > **Services** > **ZooKeeper**. On the page that is displayed, click the **Resource** tab. In the **Used Resources (By Second-Level ZNode)** area, click **By capacity** and check whether a large amount of data is written to the top-level ZNode directory.

- If yes, record the directory to which a large amount of data is written and go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Check whether data in the directory can be deleted.

NOTICE

Deleting data from ZooKeeper is a high-risk operation. Exercise caution when performing this operation.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Log in to the ZooKeeper client and delete unnecessary data from the directory to which a large amount of data is written.

1. Log in to the ZooKeeper client installation directory, for example, **/opt/client**, and configure environment variables.

   cd /opt/client

   source bigdata_env

2. Run the following command to authenticate the user (skip this step for a cluster in normal mode):

   kinit *Component service user*

3. Run the following command to log in to the client tool:

   zkCli.sh -server <*Service IP address of the node where any ZooKeeper instance resides***>:<***Client port***>**

4. Run the following command to delete unnecessary data:

   delete *Path of the file to be deleted*

Step 5 Log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **ZooKeeper** > **Configurations** > **All Configurations**, and search for **max.data.size**. The value of **max.data.size** is the maximum capacity quota of the ZooKeeper directory. The unit is byte. Search for the **GC_OPTS** configuration item and check the value of **Xmx**.

Step 6 Compare the values of **max.data.size** and **Xmx*0.65**. The threshold is the smaller value multiplied by 80%. You can change the values of **max.data.size** and **Xmx*0.65** to increase the threshold.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **ZooKeeper** for the target cluster.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.87 ALM-13010 Znode Usage of a Directory with Quota Configured Exceeds the Threshold

Description

The system checks the Znode usage of all service directories with quota configured every hour. This alarm is generated when the system detects that the level-2 Znode usage exceeds the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
13010	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.

Name	Meaning
ServiceDirectory	Specifies the directory for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the cause of the alarm.

Impact on the System

A large amount of data is written to the ZooKeeper data directory. As a result, ZooKeeper cannot provide services properly.

Possible Causes

- A large amount of data is written to the ZooKeeper data directory.
- The user-defined threshold is inappropriate.

Procedure

Check whether a large amount of data is written into the directory for which the alarm is generated.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm > Alarms**. Confirm the Znode for which the alarm is generated in **Location** of this alarm.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** and click **Resource**. In **Used Resources (By Second-Level Znode)**, check whether a large amount of data is written into the top Znode.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Log in to FusionInsight Manager, choose **O&M > Alarm > Alarms**, select Location from the drop-down list box next to **ALM-13010 Znode Usage of a Directory with Quota Configured Exceeds the Threshold**, and obtain the Znode path in ServiceDirectory.

Step 4 Log in to the ZooKeeper client as a cluster user and delete unwanted data in the Znode for which the alarm is generated.

Step 5 Log in to FusionInsight Manager, and choose **Cluster** > *Name of the desired cluster* > **Services** > *Component of the top Znode for which the alarm isgenerated*. Choose **Configurations** > **All Configurations**, search for **zk.quota.number**, increase its value, click **Save**.

NOTICE

If the Component of the top Znode for which the alarm isgenerated is ClickHouse, change the value of **clickhouse.zookeeper.quota.node.count**.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **ZooKeeper** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.88 ALM-14000 HDFS Service Unavailable

Description

The system checks the NameService service status every 60 seconds. This alarm is generated when all the NameService services are abnormal and the system considers that the HDFS service is unavailable.

This alarm is cleared when at least one NameService service is normal and the system considers that the HDFS service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14000	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

HDFS fails to provide services for HDFS service-based upper-layer components, such as HBase and MapReduce. As a result, users cannot read or write files.

Possible Causes

- The ZooKeeper service is abnormal.
- All NameService services are abnormal.
- The number of service requests is too large, and the HDFS health check fails to read and write files.
- The health check fails due to HDFS FullGC.

Procedure

Check the ZooKeeper service status.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the Alarm page, check whether **ALM-13000 ZooKeeper Service Unavailable** is reported.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 See **ALM-13000 ZooKeeper Service Unavailable** to rectify the health status of ZooKeeper fault and check whether the **Running Status** of the ZooKeeper service restores to **Normal**.

- If yes, go to **Step 3**.
- If no, go to **Step 13**.

Step 3 On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Handle the NameService service exception alarm.

Step 4 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the Alarms page, check whether **ALM-14010 NameService Service Unavailable** is reported.

- If yes, go to **Step 5**.

- If no, go to **Step 7**.

Step 5 See **ALM-14010 NameService Service Unavailable** to handle the abnormal NameService services and check whether each NameService service exception alarm is cleared.

- If yes, go to **Step 6**.
- If no, go to **Step 13**.

Step 6 On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check whether the HDFS health check fails to read or write files due to a large number of service requests.

Step 7 On FusionInsight Manager, choose **O&M > Alarm > Alarms**, and check whether **ALM-14021 NameNode Average RPC Processing Time Exceeds the Threshold** or **ALM-14022 NameNode Average RPC Queuing Time Exceeds the Threshold** is generated.

- If yes, go to **Step 8**.
- If no, go to **Step 10**.

Step 8 Rectify the abnormal NameServices by following the handling methods of **ALM-14021 NameNode Average RPC Processing Time Exceeds the Threshold** and **ALM-14022 NameNode Average RPC Queuing Time Exceeds the Threshold**. Then, check whether the alarms are cleared.

- If yes, go to **Step 9**.
- If no, go to **Step 13**.

Step 9 On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check whether the health check fails due to HDFS FullGC.

Step 10 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. On the Alarms page, check whether **ALM-14014 NameNode GC Time Exceeds the Threshold** is reported.

- If yes, go to **Step 11**.
- If no, go to **Step 13**.

Step 11 See **ALM-14014 NameNode GC Time Exceeds the Threshold** to handle the abnormal NameService services and check whether each NameService service exception alarm is cleared.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 On the **O&M > Alarm > Alarms** page, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collect fault information.

Step 13 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 14 Select the following nodes in the required cluster from the **Service**:

- ZooKeeper
- HDFS

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.89 ALM-14001 HDFS Disk Usage Exceeds the Threshold

Description

The system checks the HDFS disk usage every 30 seconds and compares the actual HDFS disk usage with the threshold. The HDFS disk usage indicator has a default threshold, this alarm is generated when the value of the disk usage of a Hadoop distributed file system (HDFS) indicator exceeds the threshold.

To change the threshold, choose **O&M** > **Alarm > Thresholds** > *Name of the desired cluster* **> HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the value of the disk usage of HDFS cluster indicator is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the value of the disk usage of HDFS cluster indicator is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14001	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Writing Hadoop distributed file system (HDFS) data is affected.

Possible Causes

The disk space configured for the HDFS cluster is insufficient.

Procedure

Check the disk capacity and delete unnecessary files.

Step 1 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**.

Step 2 Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Disk**, and select **Percentage of HDFS Capacity** to check whether the HDFS disk usage exceeds the threshold (80% by default).

- If yes, go to **Step 3**.
- If no, go to **Step 11**.

Step 3 In the **Basic Information** area, click the **NameNode(Active)** of the failure NameService and the HDFS WebUI page is displayed.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 4 On the HDFS web user interface (WebUI), click **Datanodes** tab. In the **Block pool used** column, view the disk usage of all DataNodes to check whether the disk usage of any DataNode exceeds the threshold.

- If yes, go to **Step 6**.
- If no, go to **Step 11**.

Figure 10-18 Datanode Information

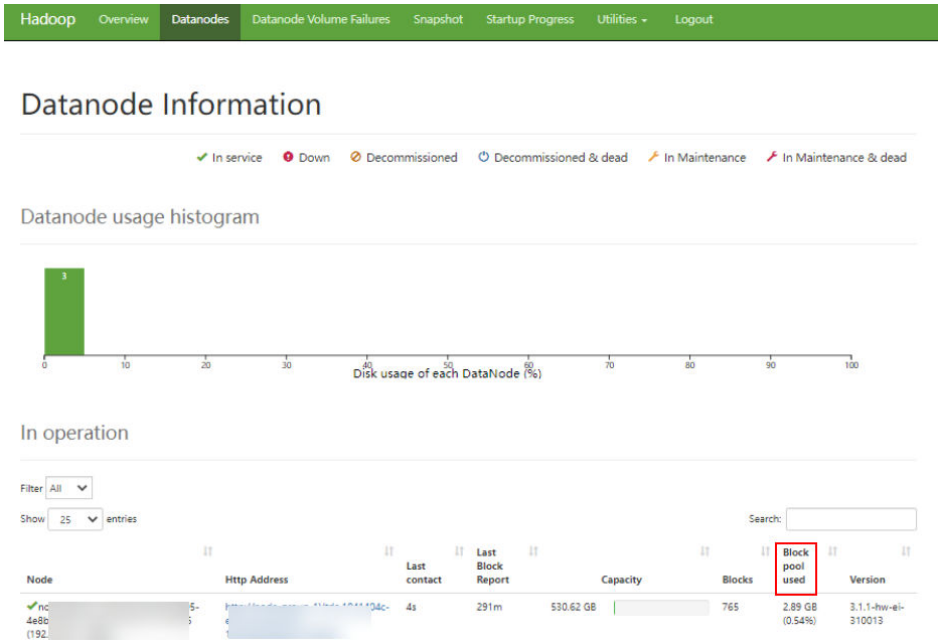

Step 5 Log in to the MRS client node as user **root**.

Step 6 Run **cd /opt/client** to switch to the client installation directory, and run **source bigdata_env**. If the cluster uses the security mode, perform security authentication. Run **kinit hdfs** and enter the password as prompted. Please obtain the password from the administrator.

Step 7 Run the **hdfs dfs -rm -r** *file or directory* command to delete unnecessary files.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Expand the system.

Step 9 Expand the disk capacity.

Step 10 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 12 Select the following nodes in the required cluster from the **Service**:

- ZooKeeper
- HDFS

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.90 ALM-14002 DataNode Disk Usage Exceeds the Threshold

Alarm Description

The system checks the DataNode disk usage every 30 seconds and compares the actual disk usage with the threshold. A default threshold range is provided for the DataNode disk usage. This alarm is generated when the DataNode disk usage exceeds the threshold.

To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **HDFS**.

If **Trigger Count** is **1**, this alarm is cleared when the DataNode disk usage is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the DataNode disk usage is less than or equal to 80% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14002	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Insufficient disk space will impact data write to HDFS.

Possible Causes

- The disk space configured for the HDFS cluster is insufficient.
- Data skew occurs among DataNodes.

Handling Procedure

Check whether the cluster disk capacity is insufficient.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and check whether the **ALM-14001 HDFS Disk Usage Exceeds the Threshold** alarm exists.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Handle the alarm by following the instructions in **ALM-14001 HDFS Disk Usage Exceeds the Threshold** and check whether the alarm is cleared.

- If yes, go to **Step 3**.
- If no, go to **Step 11**.

Step 3 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the balance status of DataNodes.

Step 4 On FusionInsight Manager, choose **Hosts**. Check whether the number of DataNodes on each rack is almost the same. If the difference is large, adjust the racks to which DataNodes belong to ensure that the number of DataNodes on each rack is almost the same. Restart the HDFS service for the settings to take effect.

Step 5 Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**.

Step 6 In the **Basic Information** area, click **NameNode(Active)**. The HDFS web UI is displayed.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 7 In the **Summary** area of the HDFS web UI, check whether the value of **Max** is 10% greater than that of **Median** in **DataNodes usages**.

- If yes, go to **Step 8**.
- If no, go to **Step 11**.

Step 8 Balance skewed data in the cluster. Log in to the MRS client as user **root**. If the cluster is in normal mode, run the **su - omm** command to switch to user **omm**. Run the **cd** command to go to the client installation directory and run the **source bigdata_env** command. If the cluster uses the security mode, perform security authentication. Run **kinit hdfs** and enter the password as prompted. Obtain the password from the MRS cluster administrator.

Step 9 Run the following command to balance data distribution:

hdfs balancer -threshold 10

Step 10 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect the fault information.

Step 11 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 12 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.91 ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold

Alarm Description

The system checks the lost blocks every 30 seconds and compares the actual lost blocks with the threshold. The lost blocks indicator has a default threshold. This alarm is generated when the number of lost HDFS blocks exceeds the threshold.

To change the threshold, choose **O&M** > **Alarm > Thresholds** > *Name of the desired cluster* **> HDFS**.

If **Trigger Count** is **1**, this alarm is cleared when the value of lost HDFS blocks is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the value of lost HDFS blocks is less than or equal to 90% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14003	Major **NOTE** The alarm severity in MRS 3.1.5 is **Critical**.	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Data stored in HDFS is lost. HDFS may enter the security mode and cannot provide write services. Lost block data cannot be restored.

Possible Causes

- The DataNode instance is abnormal.
- Data is deleted.

Handling Procedure

Check the DataNode instance.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance**.

Step 2 Check whether the **Running Status** of all DataNode instance is **Normal**.

- If yes, go to **Step 11**.
- If no, go to **Step 3**.

Step 3 Restart the DataNode instance and check whether the DataNode instance restarts successfully.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Delete the damaged file.

Step 5 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **NameNode(Active)**. On the WebUI page of the HDFS, view the information about lost blocks.

NOTE

- If a block is lost, a line in red is displayed on the WebUI.
- By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 6 The user checks whether the file containing the lost data block is useful.

NOTE

Files generated in directories **/mr-history**, **/tmp/hadoop-yarn**, and **/tmp/logs** during MapReduce task execution are unnecessary.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 The user checks whether the file containing the lost data block is backed up.

- If yes, go to **Step 8**.

- If no, go to **Step 11**.

Step 8 Log in to the HDFS client as user **root**. The user password is defined by the user before the installation. Contact the MRS cluster administrator to obtain the password. Run the following commands:

- Security mode:

  cd *Client installation directory*

  source bigdata_env

  kinit hdfs

- Normal mode:

  su - omm

  cd *Client installation directory*

  source bigdata_env

Step 9 On the node client, run **hdfs fsck / -delete** to delete the lost file. If the file where the lost block is located is a useful file, you need to write the file again to restore the data.

NOTE

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.

Step 10 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect the fault information.

Step 11 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 12 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 13 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.92 ALM-14006 Number of HDFS Files Exceeds the Threshold

Alarm Description

The system periodically checks the number of HDFS files every 30 seconds and compares the number of HDFS files with the threshold. This alarm is generated when the system detects that the number of HDFS files exceeds the threshold.

If **Trigger Count** is **1**, this alarm is cleared when the number of HDFS files is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the number of HDFS files is less than or equal to 90% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14006	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Disk storage space is insufficient, which may result in data import failure. The performance of the HDFS system is affected.

Possible Causes

The number of HDFS files exceeds the threshold.

Handling Procedure

Check the number of files in the system.

Step 1 On FusionInsight Manager, check the number of HDFS files. Specifically, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **File and Block**, and select **HDFS File** and **Total Blocks**.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations**, and search for the **GC_OPTS** parameter under **NameNode**.

Step 3 Configure the threshold of the number of configuration file objects. Specifically, change the value of **Xmx** (GB) in the **GC_OPTS** parameter. The threshold (specified by y) is calculated as follows: y = 0.2007 x Xmx - 0.6312, where x indicates the memory capacity Xmx (GB) and y indicates the number of files (unit: kW). Adjust the memory size as required.

Step 4 Confirm that the value of **GC_PROFILE** is **custom** so that the **GC_OPTS** configuration takes effect. Click **Save** and choose **More** > **Restart Instance** to restart the service.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether needless files exist in the system.

Step 6 Log in to the HDFS client as user **root**. Run **cd** to switch to the client installation directory, and run **source bigdata_env** to configure the environment variables.

If the cluster uses the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the MRS cluster administrator.

Step 7 Run **hdfs dfs -ls** *file or directory* to check whether the files in the directory can be deleted.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Run the **hdfs dfs -rm -r** *file or directory path* command. After deleting unnecessary files, wait until the files are retained in the recycle bin for a period longer than the value of **fs.trash.interval** on the NameNode. Then check whether the alarm is cleared.

NOTE

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.

- If yes, no further action is required.

- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 10 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

Configuration rules of the NameNode JVM parameter

Default value of the NameNode JVM parameter **GC_OPTS**:

-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M -XX:MetaspaceSize=128M -XX:MaxMetaspaceSize=128M -XX:+UseConcMarkSweepGC -XX:+CMSParallelRemarkEnabled -XX:CMSInitiatingOccupancyFraction=65 -XX:+PrintGCDetails -Dsun.rmi.dgc.client.gcInterval=0x7FFFFFFFFFFFFE -Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFFFFFFFE -XX:-OmitStackTraceInFastThrow -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=1M -Djdk.tls.ephemeralDHKeySize=3072 -Djdk.tls.rejectClientInitiatedRenegotiation=true -Djava.io.tmpdir=${Bigdata_tmp_dir}

The number of NameNode files is proportional to the used memory size of the NameNode. When file objects change, you need to change **-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M** in the default value. The following table lists the reference values.

Table 10-2 NameNode JVM configuration

Number of File Objects	Reference Value
10,000,000	-Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
20,000,000	-Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

Number of File Objects	Reference Value
50,000,000	-Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
100,000,000	-Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
200,000,000	-Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
300,000,000	-Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

10.93 ALM-14007 NameNode Heap Memory Usage Exceeds the Threshold

Description

The system checks the HDFS NameNode Heap Memory usage every 30 seconds and compares the actual Heap memory usage with the threshold. The HDFS NameNode Heap Memory usage has a default threshold. This alarm is generated when the HDFS NameNode Heap Memory usage exceeds the threshold.

You can change the threshold in **O&M** > **Alarm > Thresholds** > *Name of the desired cluster* **> HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the HDFS NameNode Heap memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the HDFS NameNode Heap memory usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14007	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The HDFS NameNode Heap Memory usage is too high, which affects the data read/write performance of the HDFS.

Possible Causes

The HDFS NameNode Heap Memory is insufficient.

Procedure

Delete unnecessary files.

Step 1 Log in to the HDFS client as user **root**. Run **cd** to switch to the client installation directory, and run **source bigdata_env**.

If the cluster uses the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

Step 2 Run the **hdfs dfs -rm -r** *file or directory* command to delete unnecessary files.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the NameNode JVM memory usage and configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS**.

Step 5 In the **Basic Information** area, click **NameNode(Active)** to go to the HDFS WebUI.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 6 On the HDFS WebUI, click the **Overview** tab. In **Summary**, check the numbers of files, directories, and blocks in the HDFS.

Step 7 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**. In **Search**, enter **GC_OPTS** to check the **GC_OPTS** memory parameter of **HDFS->NameNode**.

Adjust the configuration in the system.

Step 8 Check whether the memory is configured properly based on the number of files in **Step 6** and the NameNode Heap Memory parameters in **Step 7**.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

NOTE

The recommended mapping between the number of HDFS file objects (filesystem objects = files + blocks) and the JVM parameters configured for NameNode is as follows:

- If the number of file objects reaches 10,000,000, you are advised to set the JVM parameters as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the number of file objects reaches 20,000,000, you are advised to set the JVM parameters as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G
- If the number of file objects reaches 50,000,000, you are advised to set the JVM parameters as follows: -Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
- If the number of file objects reaches 100,000,000, you are advised to set the JVM parameters as follows: -Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
- If the number of file objects reaches 200,000,000, you are advised to set the JVM parameters as follows: -Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
- If the number of file objects reaches 300,000,000, you are advised to set the JVM parameters as follows: -Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

Step 9 Modify the heap memory parameters of the NameNode based on the mapping between the number of file objects and the memory. Click **Save** and choose **Dashboard** > **More** > **Restart Service**.

Step 10 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 12 Select the following nodes in the required cluster from the **Service**:

- ZooKeeper
- HDFS

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.94 ALM-14008 DataNode Heap Memory Usage Exceeds the Threshold

Description

The system checks the HDFS DataNode Heap Memory usage every 30 seconds and compares the actual Heap Memory usage with the threshold. The HDFS DataNode Heap Memory usage has a default threshold. This alarm is generated when the HDFS DataNode Heap Memory usage exceeds the threshold.

You can change the threshold in **O&M** > **Alarm > Thresholds** > *Name of the desired cluster* **> HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the HDFS DataNode Heap Memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the HDFS DataNode Heap Memory usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14008	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The HDFS DataNode Heap Memory usage is too high, which affects the data read/write performance of the HDFS.

Possible Causes

The HDFS DataNode Heap Memory is insufficient.

Procedure

Delete unnecessary files.

Step 1 Log in to the HDFS client as user **root**. Run **cd** to switch to the client installation directory, and run **source bigdata_env**.

If the cluster uses the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

Step 2 Run the **hdfs dfs -rm -r** *file or directory* command to delete unnecessary files.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the DataNode JVM memory usage and configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS**.

Step 5 In the **Basic Information** area, click **NameNode(Active)** to go to the HDFS WebUI.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 6 On the HDFS WebUI, click the **DataNodes** tab, and check the number of blocks of all DataNodes related to the alarm.

Step 7 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**. In **Search**, enter **GC_OPTS** to check the GC_OPTS memory parameter of **HDFS->DataNode**.

Adjust the configuration in the system.

Step 8 Check whether the memory is configured properly based on the number of block in **Step 6** and the DataNode Heap Memory parameters in **Step 7**.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

NOTE

The mapping between the average number of blocks of a DataNode instance and the DataNode memory is as follows:

- If the average number of blocks of a DataNode instance reaches 2,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the average number of blocks of a DataNode instance reaches 5,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

Step 9 Modify the heap memory parameters of the DataNode based on the mapping between the number of blocks and the memory. Click **Save** and choose **Dashboard** > **More** > **Restart Service**.

Step 10 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 12 Select **HDFS** in the required cluster from the **Service**.

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.95 ALM-14009 Number of Dead DataNodes Exceeds the Threshold

Description

The system periodically detects the number of dead DataNodes in the HDFS cluster every 30 seconds, and compares the number with the threshold. The

number of DataNodes in the Dead state has a default threshold. This alarm is generated when the number exceeds the threshold.

You can change the threshold in **O&M** > **Alarm > Thresholds** > *Name of the desired cluster* **> HDFS**.

When the **Trigger Count** is 1, this alarm is cleared when the number of Dead DataNodes is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the number of Dead DataNodes is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14009	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

DataNodes that are in the Dead state cannot provide HDFS services.

Possible Causes

- DataNodes are faulty or overloaded.
- The network between the NameNode and the DataNode is disconnected or busy.
- NameNodes are overloaded.

- The NameNodes are not restarted after the DataNode is deleted.

Procedure

Check whether DataNodes are faulty.

Step 1 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS**. The **HDFS Status** page is displayed.

Step 2 In the **Basic Information** area, click **NameNode(Active)** to go to the HDFS WebUI.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 3 On the HDFS WebUI, click the **Datanodes** tab. In the **In operation** area, click **Filter** to check whether **down** is in the drop-down list.

- If yes, select **down**, record the information about the filtered DataNodes, and go to **Step 4**.
- If no, go to **Step 8**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* > **Services** > **HDFS** > **Instance** to check whether recorded DataNodes exist in the instance list.

- If all recorded DataNodes exist, go to **Step 5**.
- If none of the recorded DataNodes exists, go to **Step 6**.
- If some of the recorded DataNodes exist, go to **Step 7**.

Step 5 Locate the DataNode instance, click **More** > **Restart Instance** to restart it and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Step 6 Select all NameNode instances, choose **More** > **Instance Rolling Restart** to restart them and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Step 7 Select all NameNode instances, choose **More** > **Instance Rolling Restart** to restart them. Locate the DataNode instance, click **More** > **Restart Instance** to restart it and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check the status of the network between the NameNode and the DataNode.

Step 8 Log in to the faulty DataNode on the management page as user **root**, and run the **ping** *IP address of the NameNode* command to check whether the network between the DataNode and the NameNode is abnormal.

On the FusionInsight Manager page, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance**. In the instance list, view the service plane IP address of the faulty DataNode.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 Rectify the network fault, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check whether the DataNode is overloaded.

Step 10 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** and check whether the alarm **ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold** exists.

- If yes, go to **Step 11**.
- If no, go to **Step 13**.

Step 11 See **ALM-14008 HDFS DataNode Memory Usage Exceeds the Threshold** to handle the alarm and check whether the alarm is cleared.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 Check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 13**.

Check whether the NameNode is overloaded.

Step 13 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** and check whether the alarm **ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold** exists.

- If yes, go to **Step 14**.
- If no, go to **Step 16**.

Step 14 See **ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold** to handle the alarm and check whether the alarm is cleared.

- If yes, go to **Step 15**.
- If no, go to **Step 16**.

Step 15 Check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 16**.

Collect fault information.

Step 16 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 17 Select **HDFS** in the required cluster from the **Service**.

Step 18 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.96 ALM-14010 NameService Service Is Abnormal

Alarm Description

The system checks the NameService service status every 180 seconds. This alarm is generated when the NameService service is unavailable.

This alarm is cleared when the NameService service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14010	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.

Impact on the System

HDFS fails to provide services for upper-layer components based on the NameService service, such as HBase and MapReduce. As a result, users cannot read or write files.

Possible Causes

- The KrbServer service is abnormal.
- The JournalNode is faulty.
- The DataNode is faulty.
- The disk capacity is insufficient.
- The NameNode enters safe mode.

Handling Procedure

Check the KrbServer service status.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 2 Check whether the KrbServer service exists.

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Step 3 Click **KrbServer**.

Step 4 Click **Instances**. On the KrbServer management page, select the faulty instance, and choose **More** > **Restart Instance**. Check whether the instance successfully restarts.

- If yes, go to **Step 5**.
- If no, go to **Step 24**.

Step 5 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check the JournalNode instance status.

Step 6 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 7 Choose **HDFS** > **Instances**.

Step 8 Check whether the **Running Status** of the JournalNode is **Normal**.

- If yes, go to **Step 11**.
- If no, go to **Step 9**.

Step 9 Select the faulty JournalNode, and choose **More** > **Restart Instance**. Check whether the JournalNode successfully restarts.

- If yes, go to **Step 10**.
- If no, go to **Step 24**.

Step 10 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Check the DataNode instance status.

Step 11 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**.

Step 12 Click **Instances** and check whether **Running Status** of all DataNodes is **Normal**.

- If yes, go to **Step 15**.
- If no, go to **Step 13**.

Step 13 Click **Instances**. On the DataNode management page, select the faulty instance, and choose **More** > **Restart Instance**. Check whether the DataNode successfully restarts.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Check disk status.

Step 15 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Host**.

Step 16 In the **Disk** column, check whether the disk space is insufficient.

- If yes, go to **Step 17**.
- If no, go to **Step 19**.

Step 17 Expand the disk capacity.

Step 18 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 19**.

Check whether NameNode is in the safe mode.

Step 19 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**. Click **NameNode(Active)** of the abnormal NameService. The NameNode web UI is displayed.

NOTE

By default, the admin user does not have the management rights of other components. If the page cannot be opened or the content is not completely displayed due to insufficient permission when you access the native page of a component, you can manually create a user with the management rights of the corresponding component to log in to the component.

Step 20 On the NameNode web UI, check whether "Safe mode is ON." is displayed.

Information behind **Safe mode is ON** is alarm information and is displayed based actual conditions.

- If yes, go to **Step 21**.
- If no, go to **Step 24**.

Figure 10-19 Overview

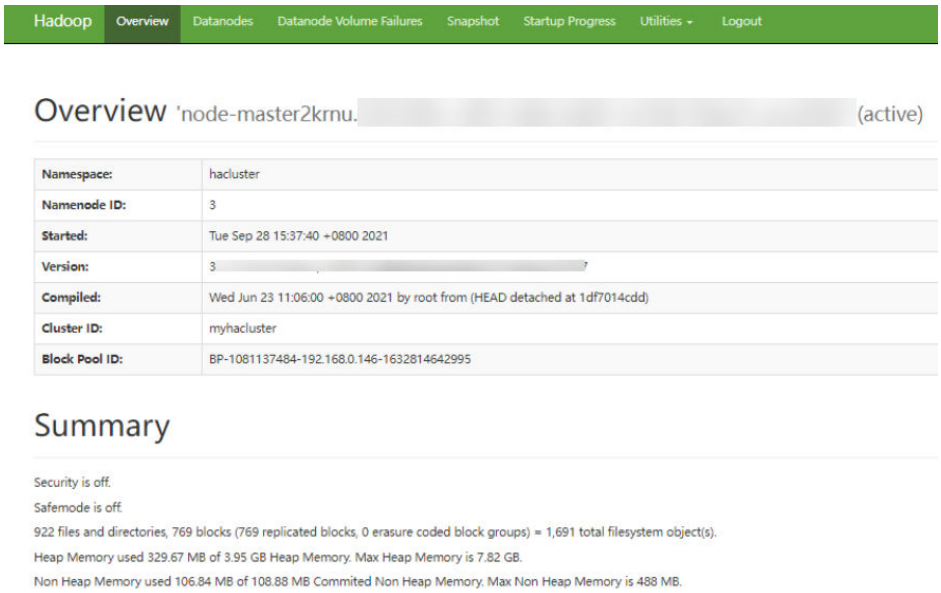

Step 21 Log in to the client as user **root**. Run the **cd** command to go to the client installation directory and run the **source bigdata_env** command. If the cluster uses the security mode, perform security authentication. Run the **kinit hdfs** command and enter the password as prompted. The password can be obtained from the MRS cluster administrator. If the cluster uses the non-security mode, log in as user **omm** and run the command. Ensure that user **omm** has the client execution permission.

Step 22 Run **hdfs dfsadmin -safemode leave**.

Step 23 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 24**.

Collect the fault information.

Step 24 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 25 In the **Service** area, select the following nodes of the desired cluster.

- ZooKeeper
- HDFS

Step 26 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 27 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.97 ALM-14011 DataNode Data Directory Is Not Configured Properly

Description

The DataNode parameter **dfs.datanode.data.dir** specifies DataNode data directories. This alarm is generated when a configured data directory cannot be created, a data directory uses the same disk as other critical directories in the system, or multiple directories use the same disk immediately.

This alarm is cleared when the DataNode data directory is configured properly and this DataNode for which the alarm is generated is restarted.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14011	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the DataNode data directory is mounted to the root directory or a critical directory, the disk space of the root directory or critical directory will be used up after long time running and the system will be faulty.

If the DataNode data directory is not configured properly, HDFS performance will deteriorate.

Possible Causes

- The DataNode data directory fails to be created.
- The DataNode data directory uses the same disk with critical directories, such as **/** or **/boot**.
- Multiple directories in the DataNode data directory use the same disk.

Procedure

Check the alarm cause and information about the DataNode for which the alarm is generated.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

Step 2 In **HostName** of **Location**, obtain the host name of the DataNode for which the alarm is generated.

Delete directories that do not comply with the disk plan from the DataNode data directory.

Step 3 Choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance**. In the instance list, click the DataNode instance on the node for which the alarm is generated.

Step 4 Click **Instance Configurations** and view the value of the DataNode parameter **dfs.datanode.data.dir**.

Step 5 Check whether all DataNode data directories are consistent with the disk plan.

- If yes, go to **Step 6**.
- If no, go to **Step 9**.

Step 6 Modify the DataNode parameter **dfs.datanode.data.dir** and delete the incorrect directories.

Step 7 Choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance** and restart the DataNode instance.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Log in to the DataNode for which the alarm is generated as user **root**.

- If the alarm cause is "The DataNode data directory fails to be created", go to **Step 10**.
- If the alarm cause is "The DataNode data directory uses the same disk with critical directories, such **/** or **/boot**", go to **Step 17**.
- If the alarm cause is "Multiple directories in the DataNode data directory uses the same disk", go to **Step 21**.

Check whether the DataNode data directory fails to be created.

Step 10 Run the **su - omm** command to switch to user **omm**.

Step 11 Run the **ls** command to check whether the directories exist in the DataNode data directory.

- If yes, go to **Step 26**.
- If no, go to **Step 12**.

Step 12 Run the **mkdir** *data directory* command to create the directory and check whether the directory can be successfully created.

- If yes, go to **Step 24**.
- If no, go to **Step 13**.

Step 13 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** to check whether alarm **ALM-12017 Insufficient Disk Capacity** exists.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Adjust the disk capacity and check whether alarm **ALM-12017 Insufficient Disk Capacity** is cleared. For details, see **ALM-12017 Insufficient Disk Capacity**.

- If yes, go to **Step 12**.
- If no, go to **Step 15**.

Step 15 Check whether user **omm** has the **rwx** or **x** permission of all the upper-layer directories of the directory. (For example, for **/tmp/abc/**, user **omm** has the **x** permission for directory **tmp** and the **rwx** permission for directory **abc**.)

- If yes, go to **Step 24**.
- If no, go to **Step 16**.

Step 16 Run the **chmod u+rwx** *path* or **chmod u+x** *path* command as user **root** to assign the **rwx** or **x** permission of these directories to user **omm**. Then go to **Step 12**.

Check whether the DataNode data directory use the same disk as other critical directories in the system.

Step 17 Run the **df** command to obtain the disk mounting information of each directory in the DataNode data directory.

Step 18 Check whether the directories mounted to the disk are critical directories, such as **/** or **/boot**.

- If yes, go to **Step 19**.
- If no, go to **Step 24**.

Step 19 Change the value of the DataNode parameter **dfs.datanode.data.dir** and delete the directories that use the same disk as critical directories.

Step 20 Go to **Step 24**.

Check whether multiple directories in the DataNode data directory use the same disk.

Step 21 Run the **df** command to obtain the disk mounting information of each directory in the DataNode data directory. Record the mounted directory in the command output.

Step 22 Modify the DataNode node parameters **dfs.datanode.data.dir** to reserve only one directory among the directories that mounted to the same disk directory.

Step 23 Go to **Step 24**.

Restart the DataNode and check whether the alarm is cleared.

Step 24 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance** and restart the DataNode instance

Step 25 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 26**.

Collect fault information.

Step 26 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 27 Select **HDFS** in the required cluster from the **Service**.

Step 28 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 29 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.98 ALM-14012 JournalNode Is Out of Synchronization

Description

On the active NameNode, the system checks the data consistency of all JournalNodes in the cluster every 5 minutes. This alarm is generated when the data on a JournalNode is inconsistent with the data on the other JournalNodes.

This alarm is cleared in 5 minutes after the data on JournalNodes is consistent.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14012	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.

Impact on the System

When a JournalNode is working incorrectly, the data on the node becomes inconsistent with that on the other JournalNodes. If data on more than half of JournalNodes is inconsistent, the NameNode cannot work correctly, making the HDFS service unavailable.

Possible Causes

- The JournalNode instance does not exist (deleted or migrated).
- The JournalNode instance has not been started or has been stopped.
- The JournalNode instance is working incorrectly.
- The network of the JournalNode is unreachable.

Procedure

Check whether the JournalNode instance has been started up.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

Step 2 Check **Location** and obtain the IP address of the JournalNode for which the alarm is generated.

Step 3 Choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance**. In the instance list, check whether the JournalNode instance exists on the node for which the alarm is generated.

- If yes, go to **Step 5**.
- If no, go to **Step 4**.

Step 4 Choose **O&M** > **Alarm** > **Alarms**. In the alarm list, click **Clear** in the **Operation** column of the alarm. In the dialog box that is displayed, click **OK**. No further action is needed.

Step 5 Click the JournalNode instance and check whether its **Configuration Status** is **Synchronized**.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 Select the JournalNode instance and choose **Start Instance** to start the instance.

Step 7 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Check whether the JournalNode instance is working correctly.

Step 8 Check whether **Running Status** of the JournalNode instance is **Normal**.

- If yes, go to **Step 11**.
- If no, go to **Step 9**.

Step 9 Select the JournalNode instance and choose **More** > **Restart Instance** to start the instance.

Step 10 After 5 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Check whether the network of the JournalNode is reachable.

Step 11 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance** to check the service IP address of the active NameNode.

Step 12 Log in to the active NameNode as user **root**.

Step 13 Run the **ping** command to check whether a timeout occurs or the network is unreachable between the active NameNode and the JournalNode.

ping *service IP address of the JournalNode*

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Contact the network administrator to rectify the network fault and check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 15**.

Collect fault information.

Step 15 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 16 Select **HDFS** in the required cluster from the **Service**.

Step 17 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.99 ALM-14013 Failed to Update the NameNode FsImage File

Description

HDFS metadata is stored in the FsImage file of the NameNode data directory, which is specified by the **dfs.namenode.name.dir** configuration item. The standby NameNode periodically combines existing FsImage files and Editlog files stored in the JournalNode to generate a new FsImage file, and then pushes the new FsImage file to the data directory of the active NameNode. This period is specified by the **dfs.namenode.checkpoint.period** configuration item of HDFS. The default value is 3600s, namely, one hour. If the FsImage file in the data directory of the active NameNode is not updated, the HDFS metadata combination function is abnormal and requires rectification.

On the active NameNode, the system checks the FsImage file information every five minutes. This alarm is generated when no FsImage file is generated within three combination periods.

This alarm is cleared when a new FsImage file is generated and pushed to the active NameNode, which indicates that the HDFS metadata combination function can be properly used.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14013	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService for which the alarm is generated.

Impact on the System

If the FsImage file in the data directory of the active NameNode is not updated, the HDFS metadata combination function is abnormal and requires rectification. If it is not rectified, the Editlog files increase continuously after HDFS runs for a period. In this case, HDFS restart is time-consuming because a large number of Editlog files need to be loaded. In addition, this alarm also indicates that the standby NameNode is abnormal and the NameNode high availability (HA) mechanism becomes invalid. When the active NameNode is faulty, the HDFS service becomes unavailable.

Possible Causes

- The standby NameNode is stopped.
- The standby NameNode instance is working incorrectly.
- The standby NameNode fails to generate a new FsImage file.
- Space of the data directory on the standby NameNode is insufficient.
- The standby NameNode fails to push the FsImage file to the active NameNode.
- Space of the data directory on the active NameNode is insufficient.

Procedure

Check whether the standby NameNode is stopped.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

Step 2 View **Location** and obtain the host name of the active NameNode for which the alarm is generated and name of the NameService where the active NameNode resides.

Step 3 Choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance**, find the standby NameNode instance of the NameService in the instance list, and check whether its **Configuration Status** is **Synchronized**.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Select the standby NameNode instance, choose **Start Instance**, and wait until the startup is complete.

Step 5 Wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the NameNode instance is working correctly.

Step 6 Check whether **Running Status** of the standby NameNode instance is **Normal**.

- If yes, go to **Step 9**.
- If no, go to **Step 7**.

Step 7 Select the standby NameNode instance, choose **More** > **Restart Instance**, and wait until the startup is complete.

Step 8 Wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 30**.

Check whether the standby NameNode fails to generate a new FsImage file.

Step 9 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**, and search and obtain the value of **dfs.namenode.checkpoint.period**. This value is the period of NameNode metadata combination.

Step 10 Choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance** and obtain the service IP addresses of the active and standby NameNodes of the NameService for which the alarm is generated.

Step 11 Click the **NameNode(***xx***,Standy)** and **Instance Configurations** to obtain the value of **dfs.namenode.name.dir**. This value is the FsImage storage directory of the standby NameNode.

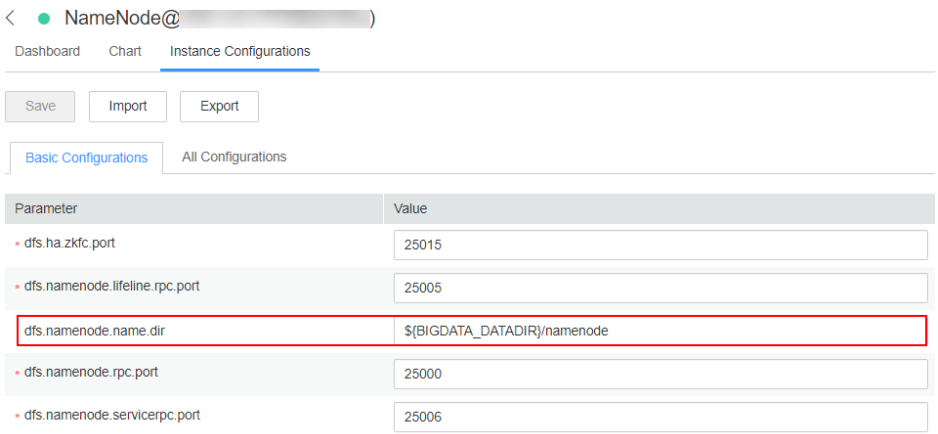

Step 12 Log in to the standby NameNode as user **root** or **omm**.

Step 13 Go to the FsImage storage directory and check the generation time of the newest FsImage file.

cd *Storage directory of the standby NameNode***/current**

stat -c %y $(ls -t | grep "fsimage_[0-9]*$" | head -1)

Step 14 Run the **date** command to obtain the current system time.

Step 15 Calculate the time difference between the generation time of the newest FsImage file and the current system time and check whether the time difference is greater than three times of the metadata combination period.

- If yes, go to **Step 16**.
- If no, go to **Step 20**.

Step 16 The metadata combination function of the standby NameNode is faulty. Run the following command to check whether the fault is caused by insufficient storage space.

Go to the FsImage storage directory and check the size of the newest FsImage file (in MB).

cd *Storage directory of the standby NameNode*/**current**

du -m $(ls -t | grep "fsimage_[0-9]*$" | head -1) | awk '{print $1}'

Step 17 Run the following command to check the available disk space of the standby NameNode (in MB).

df -m ./ | awk 'END{print $4}'

Step 18 Compare the FsImage file size and the available disk space and determine whether another FsImage file can be stored on the disk.

- If yes, go to **Step 7**.
- If no, go to **Step 19**.

Step 19 Clear the redundant files on the disk where the directory resides to reserve sufficient space for metadata. After the clearance, wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 20**.

Check whether the standby NameNode fails to push the FsImage file to the active NameNode.

Step 20 Log in to the standby NameNode as user **root**.

Step 21 Run the **su - omm** command to switch to user **omm**.

Step 22 Run the following command to check whether the standby NameNode can push the file to the active NameNode.

tmpFile=/tmp/tmp_test_$(date +%s)

echo "test" > $tmpFile

scp $tmpFile *Service IP address of the active NameNode*:**/tmp**

- If yes, go to **Step 24**.
- If no, go to **Step 23**.

Step 23 When the standby NameNode fails to push data to the active NameNode as user **omm**, contact the system administrator to handle the fault. Wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 24**.

Check whether space on the data directory of the active NameNode is insufficient.

Step 24 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance**, click the active NameNode of the NameService for which the alarm is generated, and then click **Instance Configurations** to obtain the value of **dfs.namenode.name.dir**. This value is the FsImage storage directory of the active NameNode.

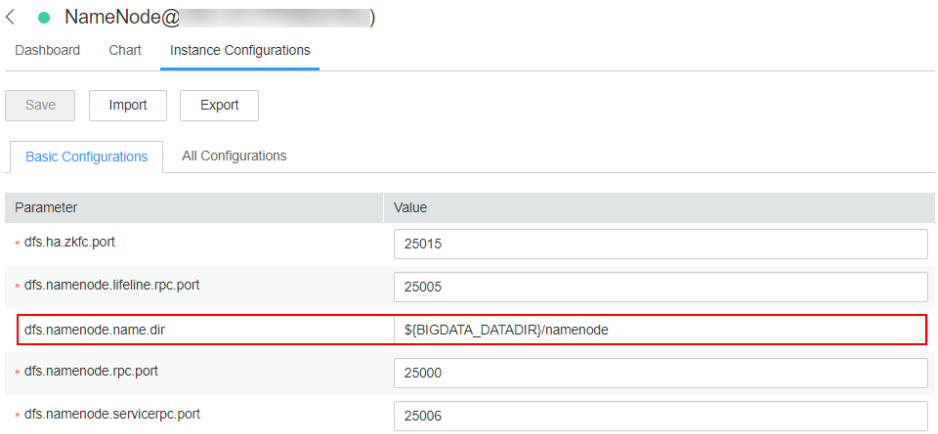

Step 25 Log in to the active NameNode as user **root** or **omm**.

Step 26 Go to the FsImage storage directory and check the size of the newest FsImage file (in MB).

cd *Storage directory of the active NameNode*/**current**

du -m $(ls -t | grep "fsimage_[0-9]*$" | head -1) | awk '{print $1}'

Step 27 Run the following command to check the available disk space of the active NameNode (in MB).

df -m ./ | awk 'END{print $4}'

Step 28 Compare the FsImage file size and the available disk space and determine whether another FsImage file can be stored on the disk.

- If yes, go to **Step 30**.
- If no, go to **Step 29**.

Step 29 Clear the redundant files on the disk where the directory resides to reserve sufficient space for metadata. After the clearance, wait for a NameNode metadata combination period and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 30**.

Collect fault information.

Step 30 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 31 Select **NameNode** in the required cluster from the **Service**.

Step 32 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 33 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.100 ALM-14014 NameNode GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the NameNode process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14014	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

A long GC duration of the NameNode process may interrupt the services.

Possible Causes

The heap memory of the NameNode instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms**. On the displayed interface, click the drop-down button of **ALM-14014 NameNode GC Time Exceeds the Threshold.** Then check the role name in **Location** and confirm the IP adress of the instance.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Instance** > **NameNode (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Garbage Collection**, and select **NameNode Garbage Collection (GC)** to check the GC duration statistics of the NameNode process collected every minute.

Figure 10-20 NameNode Garbage Collection (GC)

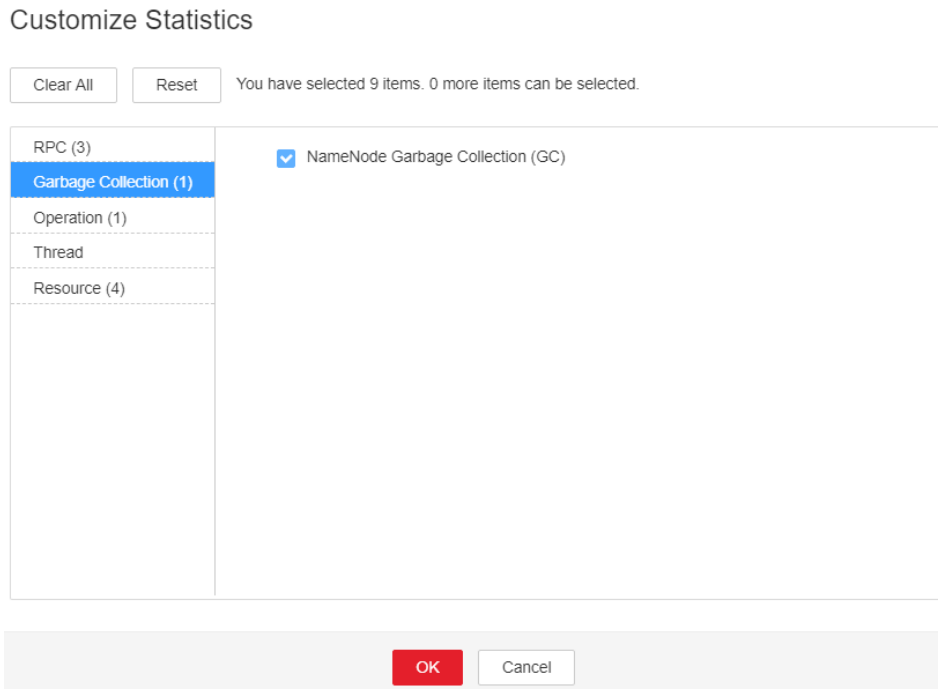

Step 3 Check whether the GC duration of the NameNode process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS** > **Configurations** > **All Configurations** > **NameNode** > **System** to increase the value of **GC_OPTS** parameter as required.

NOTE

The recommended mapping between the number of HDFS file objects (filesystem objects = files + blocks) and the JVM parameters configured for NameNode is as follows:

- If the number of file objects reaches 10,000,000, you are advised to set the JVM parameters as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the number of file objects reaches 20,000,000, you are advised to set the JVM parameters as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G
- If the number of file objects reaches 50,000,000, you are advised to set the JVM parameters as follows: -Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
- If the number of file objects reaches 100,000,000, you are advised to set the JVM parameters as follows: -Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
- If the number of file objects reaches 200,000,000, you are advised to set the JVM parameters as follows: -Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
- If the number of file objects reaches 300,000,000, you are advised to set the JVM parameters as follows: -Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

Step 5 Save the configuration and restart the NameNode instance.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **NameNode** in the required cluster from the **Service**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.101 ALM-14015 DataNode GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the DataNode process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14015	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

A long GC duration of the DataNode process may interrupt the services.

Possible Causes

The heap memory of the DataNode instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms**. On the displayed interface, click the drop-down button of **ALM-14015 DataNode GC Time Exceeds the Threshold**. Then check the role name in **Location** and confirm the IP adress of the instance.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance** > **DataNode (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Garbage Collection**, and select **DataNode Garbage Collection (GC)** to check the GC duration statistics of the DataNode process collected every minute.

Figure 10-21 DataNode Garbage Collection (GC)

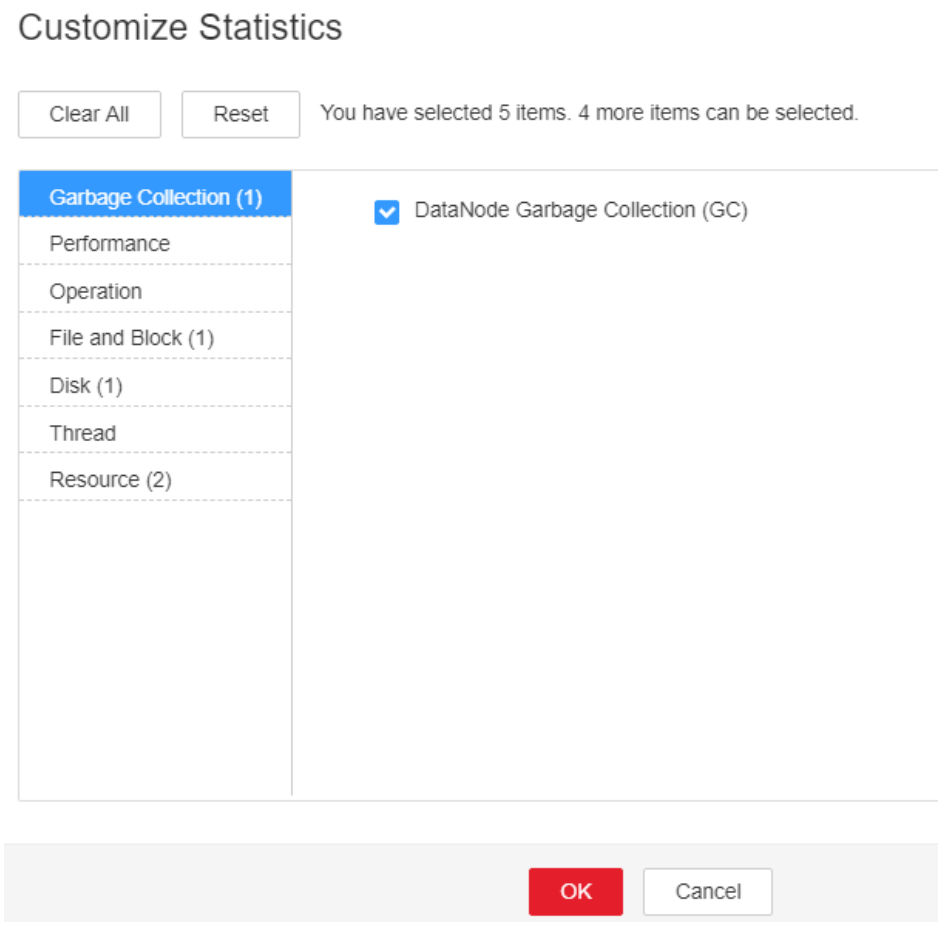

Step 3 Check whether the GC duration of the DataNode process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations** > **DataNode** > **System** to increase the value of **GC_OPTS** parameter as required.

NOTE

The mapping between the average number of blocks of a DataNode instance and the DataNode memory is as follows:

- If the average number of blocks of a DataNode instance reaches 2,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the average number of blocks of a DataNode instance reaches 5,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

Step 5 Save the configuration and restart the DataNode instance.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **DataNode** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.102 ALM-14016 DataNode Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of HDFS every 30 seconds. This alarm is generated when the direct memory usage of DataNode instances exceeds the threshold (90% of the maximum memory).

This alarm is automatically cleared when the direct memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14016	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the available direct memory of DataNode instances is insufficient, a memory overflow may occur and the service breaks down.

Possible Causes

The direct memory of DataNode instances is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

Step 1 On the **Home** page of FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, click the drop-down list in the row containing **ALM-14016 DataNode Direct Memory Usage Exceeds the Threshold**, and view the role name and IP address of the instance for which the alarm is generated in the **Location** area.

Step 2 On the **Home** page of FusionInsight Manager, choose **Cluster** > **Services** > **HDFS**. On the page that is displayed, click the **Instance** tab. In the instance list, select **DataNode** (IP address of the instance for which this alarm is generated). Click the drop-down list in the upper right corner of the chart, choose **Customize** > **Resource**, and select **DataNode Memory** to check the direct memory usage.

Step 3 Check whether the used direct memory of a DataNode instance reaches 90% (default threshold) of the maximum direct memory allocated to it.

- If yes, go to **Step 4**.
- If no, go to **Step 8**.

Step 4 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **HDFS** > **Configurations** > **All Configurations** > **DataNode** > **System**. Check whether **-XX:MaxDirectMemorySize** exists in the **GC_OPTS** parameter.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Adjust the value of **-XX:MaxDirectMemorySize**.

1. In **GC_OPTS**, check the value of **-Xmx** and check whether the node memory is sufficient.

   NOTE

   You can determine whether the node memory is sufficient based on the actual environment. For example, you can use the following method:

   Use the IP address to log in to the instance for which the alarm is generated as user **root** and run the **free -g** command to check the value of **Mem** in the **free** column. The value indicates the available memory of the node. In the following example, the available memory of the node is 4 GB.

   ```
                total        used        free      shared  buff/cache   available
   Mem:            112          48           4          10          58          46
   ......
   ```

   If the value of **Mem** is at least that of **-Xmx**, the node memory is sufficient. If the value of **Mem** is less than that of **-Xmx**, the node memory is insufficient.

   - If yes, change the value of **-XX:MaxDirectMemorySize** to that of **-Xmx**.
   - If no, increase **-XX:MaxDirectMemorySize** to a value no larger than that of **Mem**.

2. Save the configuration and restart the DataNode instances.

Step 6 Check whether **ALM-14008 DataNode Heap Memory Usage Exceeds the Threshold** exists.

- If yes, rectify the fault by referring to **ALM-14008 DataNode Heap Memory Usage Exceeds the Threshold**.
- If no, go to **Step 7**.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **DataNode** for the target cluster.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.103 ALM-14017 NameNode Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the HDFS service every 30 seconds. This alarm is generated when the direct memory usage of a NameNode instance exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14017	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available direct memory of the HDFS service is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the NameNode instance is overused or the direct memory is inappropriately allocated.

Procedure

Check the direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms.** On the displayed interface, click the drop-down button of **ALM-14017 NameNode Direct Memory Usage Exceeds the Threshold**. Then check the role name in **Location** and confirm the IP adress of the instance.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Instance** > **NameNode (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Resource**, and select **NameNode Memory** to check the direct memory usage.

Step 3 Check whether the used direct memory of NameNode reaches 90% of the maximum direct memory specified for NameNode by default.

- If yes, go to **Step 4**.
- If no, go to **Step 8**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations** > **NameNode** > **System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC_OPTS** parameter.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 In the **GC_OPTS** parameter, delete "-XX:MaxDirectMemorySize". Save the configuration and restart the NameNode instance.

Step 6 Check whether the **ALM-14007 NameNode Heap Memory Usage Exceeds the Threshold** exists.

- If yes, handle the alarm by referring to **ALM-14007 NameNode Heap Memory Usage Exceeds the Threshold**.
- If no, go to **Step 7**.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 9 Select **NameNode** in the required cluster from the **Service**.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.104 ALM-14018 NameNode Non-heap Memory Usage Exceeds the Threshold

Description

The system checks the non-heap memory usage of the HDFS NameNode every 30 seconds and compares the actual usage with the threshold. The non-heap memory usage of the HDFS NameNode has a default threshold. This alarm is generated when the non-heap memory usage of the HDFS NameNode exceeds the threshold.

Users can choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> HDFS** to change the threshold.

This alarm is cleared when the no-heap memory usage of the HDFS NameNode is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14018	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the memory usage of the HDFS NameNode is too high, data read/write performance of HDFS will be affected.

Possible Causes

Non-heap memory of the HDFS NameNode is insufficient.

Procedure

Delete unnecessary files.

Step 1 Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory, and run the **source bigdata_env** command.

If the cluster adopts the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

Step 2 Run the **hdfs dfs -rm -r** *file or directory path* command to delete unnecessary files.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the NameNode JVM non-heap memory usage and configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HDFS**. The HDFS status page is displayed.

Step 5 In the **Basic Information** area, click **NameNode(Active)**. The HDFS WebUI is displayed.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 6 On the HDFS WebUI, click the **Overview** tab. In **Summary**, check the numbers of files, directories, and blocks in HDFS.

Step 7 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**. In **Search**, enter **GC_OPTS** to check the **GC_OPTS** non-heap memory parameter of **HDFS->NameNode**.

Adjust system configurations.

Step 8 Check whether the non-heap memory is properly configured based on the number of file objects in **Step 6** and the non-heap parameters configured for NameNode in **Step 7**.

- If yes, go to **Step 9**.
- If no, go to **Step 12**.

NOTE

The recommended mapping between the number of HDFS file objects (filesystem objects = files + blocks) and the JVM parameters configured for NameNode is as follows:

- If the number of file objects reaches 10,000,000, you are advised to set the JVM parameters as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the number of file objects reaches 20,000,000, you are advised to set the JVM parameters as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G
- If the number of file objects reaches 50,000,000, you are advised to set the JVM parameters as follows: -Xms32G -Xmx32G -XX:NewSize=3G -XX:MaxNewSize=3G
- If the number of file objects reaches 100,000,000, you are advised to set the JVM parameters as follows: -Xms64G -Xmx64G -XX:NewSize=6G -XX:MaxNewSize=6G
- If the number of file objects reaches 200,000,000, you are advised to set the JVM parameters as follows: -Xms96G -Xmx96G -XX:NewSize=9G -XX:MaxNewSize=9G
- If the number of file objects reaches 300,000,000, you are advised to set the JVM parameters as follows: -Xms164G -Xmx164G -XX:NewSize=12G -XX:MaxNewSize=12G

Step 9 Modify the **GC_OPTS** parameter of the NameNode based on the mapping between the number of file objects and non-heap memory.

Step 10 Save the configuration and click **Dashboard** > **More** > **Restart Service**.

Step 11 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Collect fault information.

Step 12 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 13 Select the following services in the required cluster from the **Service**.

- ZooKeeper
- HDFS

Step 14 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 15 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.105 ALM-14019 DataNode Non-heap Memory Usage Exceeds the Threshold

Description

The system checks the non-heap memory usage of the HDFS DataNode every 30 seconds and compares the actual usage with the threshold. The non-heap memory usage of the HDFS DataNode has a default threshold. This alarm is generated when the non-heap memory usage of the HDFS DataNode exceeds the threshold.

Users can choose **O&M > Alarm > Thresholds>** *Name of the desired cluster* **> HDFS** to change the threshold.

This alarm is cleared when the no-heap memory usage of the HDFS DataNode is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14019	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the memory usage of the HDFS DataNode is too high, data read/write performance of HDFS will be affected.

Possible Causes

Non-heap memory of the HDFS DataNode is insufficient.

Procedure

Delete unnecessary files.

Step 1 Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory, and run the **source bigdata_env** command.

If the cluster adopts the security mode, perform security authentication.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

Step 2 Run the **hdfs dfs -rm -r** *file or directory path* command to delete unnecessary files.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the DataNode JVM non-heap memory usage and configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS**.

Step 5 In the **Basic Information** area, click **NameNode(Active)**. The HDFS WebUI is displayed.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 6 On the HDFS WebUI, click the **Datanodes** tab to view the number of blocks of all DataNodes that report alarms.

Step 7 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**. In **Search**,

enter **GC_OPTS** to check the **GC_OPTS** non-heap memory parameter of **HDFS->DataNode**.

Adjust system configurations.

Step 8 Check whether the memory is properly configured based on the number of blocks in **Step 6** and the memory parameters configured for DataNode in **Step 7**.

- If yes, go to **Step 9**.
- If no, go to **Step 12**.

NOTE

The mapping between the average number of blocks of a DataNode instance and the DataNode memory is as follows:

- If the average number of blocks of a DataNode instance reaches 2,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
- If the average number of blocks of a DataNode instance reaches 5,000,000, the reference values of the JVM parameters of the DataNode are as follows: -Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

Step 9 Modify the **GC_OPTS** parameter of the DataNode based on the mapping between the number of blocks and memory.

Step 10 Save the configuration and click **Dashboard** > **More** > **Restart Service**.

Step 11 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Collect fault information.

Step 12 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 13 Select the following services in the required cluster from the **Service**.

- ZooKeeper
- HDFS

Step 14 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 15 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.106 ALM-14020 Number of Entries in the HDFS Directory Exceeds the Threshold

Description

The system obtains the number of subfiles and subdirectories in a specified directory every hour and checks whether it reaches the percentage of the threshold (the maximum number of subfiles and subdirectories in an HDFS directory, the threshold for triggering an alarm is **90%** by default). If it exceeds the percentage of the threshold, an alarm is triggered.

When the number of subfiles and subdirectories in the directory the alarm is lower than the percentage of the threshold, the alarm is automatically cleared. When the monitoring switch is disabled, alarms corresponding to all directories are cleared. If a directory is removed from the monitoring list, alarms corresponding to the directory are cleared.

NOTE

- The **dfs.namenode.fs-limits.max-directory-items** parameter specifies the maximum number of subfiles and subdirectories in the HDFS directory. Its default value is **1048576**. If the number of subfiles and subdirectories in a directory exceeds the parameter value, subfiles and subdirectories cannot be created in the directory.
- The **dfs.namenode.directory-items.monitor** parameter specifies the list of directories to be monitored. Its default value is **/tmp,/SparkJobHistory,/mr-history**.
- The **dfs.namenode.directory-items.monitor.enabled** parameter is used to enable or disable the monitoring switch. Its default value is **true**, which means the monitoring switch is enabled by default.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14020	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.

Name	Meaning
Directory	Specifies the directory for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the number of entries in the monitored directory exceeds 90% of the threshold, an alarm is triggered, but entries can be added to the directory. Once the maximum threshold is exceeded, entries will fail to be added to the directory.

Possible Causes

The number of entries in the monitored directory exceeds 90% of the threshold.

Procedure

Check whether unnecessary files exist in the system.

Step 1 Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory, and run the **source bigdata_env** command to set the environment variables.

If the cluster is in security mode, security authentication is required.

Run the **kinit hdfs** command and enter the password as prompted. Obtain the password from the administrator.

Step 2 Run the following command to check whether files and directories in the directory with the alarm can be deleted:

hdfs dfs -ls *Directory with the alarm*

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Run the following command to delete unnecessary files.

hdfs dfs -rm -r -f *File or directory path*

NOTE

Deleting a file or folder is a high-risk operation. Ensure that the file or folder is no longer required before performing this operation.

Step 4 Wait 1 hour and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the threshold is correctly configured.

Step 5 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**. Search for the **dfs.namenode.fs-limits.max-directory-items** parameter and check whether the parameter value is appropriate.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

Step 6 Increase the parameter value.

Step 7 Save the configuration and click **Dashboard** > **More** > **Restart Service**.

Step 8 Wait 1 hour and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 10 Select **HDFS** in the required cluster from the **Service**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.107 ALM-14021 NameNode Average RPC Processing Time Exceeds the Threshold

Description

The system checks the average RPC processing time of NameNode every 30 seconds, and compares the actual average RPC processing time with the threshold (default value: 100 ms). This alarm is generated when the system detects that the average RPC processing time exceeds the threshold for several consecutive times (10 times by default).

You can choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> HDFS** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the average RPC processing time of NameNode is less than or equal to the threshold. When the

Trigger Count is greater than 1, this alarm is cleared when the average RPC processing time of NameNode is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14021	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

NameNode cannot process the RPC requests from HDFS clients, upper-layer services that depend on HDFS, and DataNode in a timely manner. Specifically, the services that access HDFS run slowly or the HDFS service is unavailable.

Possible Causes

- The CPU performance of NameNode nodes is insufficient and therefore NameNode nodes cannot process messages in a timely manner.
- The configured NameNode memory is too small and frame freezing occurs on the JVM due to frequent full garbage collection.
- NameNode parameters are not configured properly, so NameNode cannot make full use of system performance.

Procedure

Obtain alarm information.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

Step 2 Check the alarm. Obtain the host name of the NameNode node involved in this alarm from the **HostName** information of **Location**. Then obtain the name of the NameService node involved in this alarm from the **NameServiceName** information of **Location**.

Check whether the threshold is too small.

Step 3 Check the status of the services that depend on HDFS. Check whether the services run slowly or task execution times out.

- If yes, go to **Step 8**.
- If no, go to **Step 4**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **RPC**, and select **Average Time of Active NameNode RPC Processing** and click **OK**.

Figure 10-22 Average Time of Active NameNode RPC Processing

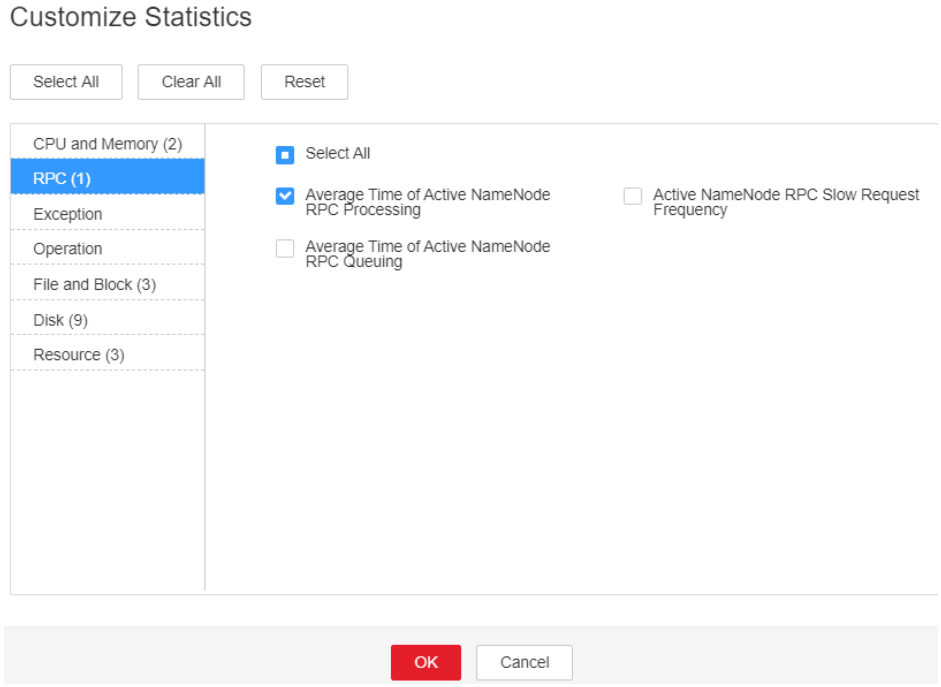

Step 5 On the **Average Time of Active NameNode RPC Processing** monitoring page, obtain the value of the NameService node involved in this alarm.

Step 6 On the FusionInsight Manager portal, choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> HDFS**. Locate **Average Time of Active NameNode RPC Processing** and click the **Modify** in the **Operation** column of the default rule. The **Modify Rule** page is displayed. Change **Threshold** to 150% of the peak value within one day before and after the alarm is generated. Click **OK** to save the new threshold.

Figure 10-23 Modify Rule

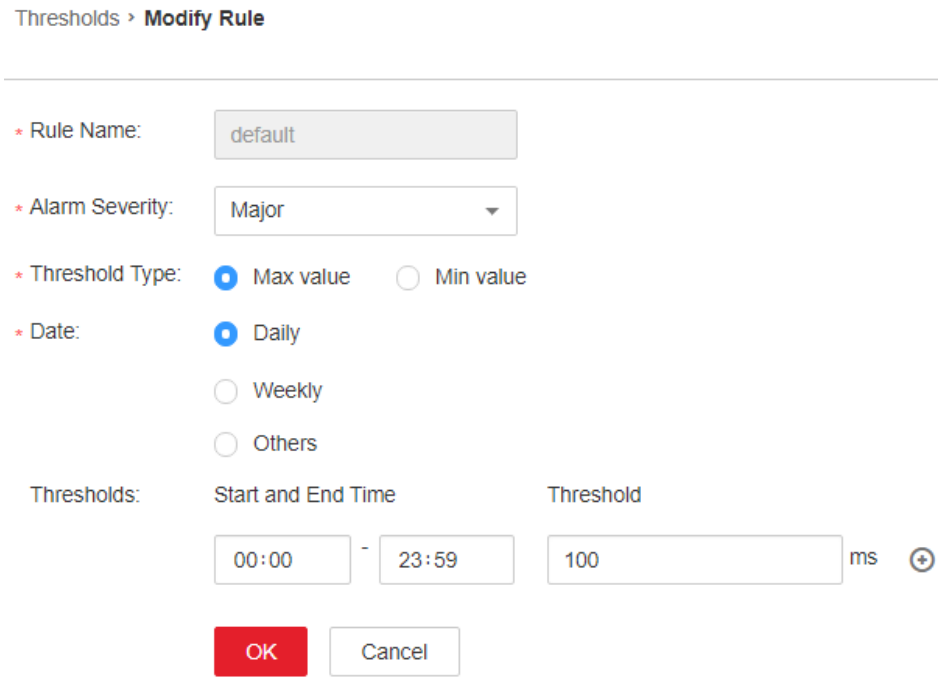

Step 7 Wait for 5 minutes and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check whether the CPU performance of the NameNode node is sufficient.

Step 8 On the FusionInsight Manager portal, click **O&M > Alarm >Alarms** and check whether **ALM-12016 CPU Usage Exceeds the Threshold** is generated for the NameNode node.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

Step 9 Handle **ALM-12016 CPU Usage Exceeds the Threshold** by taking recommended actions.

Step 10 Wait for 10 minutes and check whether alarm 14021 is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Check whether the memory of the NameNode node is too small.

Step 11 On the FusionInsight Manager portal, click **O&M > Alarm >Alarms** and check whether **ALM-14007 HDFS NameNode Heap Memory Usage Exceeds the Threshold** is generated for the NameNode node.

- If yes, go to **Step 12**.
- If no, go to **Step 14**.

Step 12 Handle **ALM-14007 HDFS NameNode Heap Memory Usage Exceeds the Threshold** by taking recommended actions.

Step 13 Wait for 10 minutes and check whether alarm 14021 is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Check whether NameNode parameters are configured properly.

Step 14 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**. Search for parameter **dfs.namenode.handler.count** and view its value. If the value is less than or equal to 128, change it to **128**. If the value is greater than 128 but less than 192, change it to **192**.

Step 15 Search for parameter **ipc.server.read.threadpool.size** and view its value. If the value is less than 5, change it to **5**.

Step 16 Click **Save** and click **OK**.

Step 17 On the **Instance** page of HDFS, select the standby NameNode of NameService involved in this alarm and choose **More** > **Restart Instance**. Enter the password and click **OK**. Wait until the standby NameNode is started up.

Step 18 On the **Instance** page of HDFS, select the active NameNode of NameService involved in this alarm and choose **More** > **Restart Instance**. Enter the password and click **OK**. Wait until the active NameNode is started up.

Step 19 Wait for 1 hour and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 20**.

Collect fault information.

Step 20 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 21 Select the following node in the required cluster from the **Service**.

- HDFS

Step 22 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 23 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.108 ALM-14022 NameNode Average RPC Queuing Time Exceeds the Threshold

Description

The system checks the average RPC queuing time of NameNode every 30 seconds, and compares the actual average RPC queuing time with the threshold (default value: 200 ms). This alarm is generated when the system detects that the average RPC queuing time exceeds the threshold for several consecutive times (10 times by default).

You can choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* > **HDFS** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the average RPC queuing time of NameNode is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the average RPC queuing time of NameNode is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14022	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

NameNode cannot process the RPC requests from HDFS clients, upper-layer services that depend on HDFS, and DataNode in a timely manner. Specifically, the services that access HDFS run slowly or the HDFS service is unavailable.

Possible Causes

- The CPU performance of NameNode nodes is insufficient and therefore NameNode nodes cannot process messages in a timely manner.
- The configured NameNode memory is too small and frame freezing occurs on the JVM due to frequent full garbage collection.
- NameNode parameters are not configured properly, so NameNode cannot make full use of system performance.
- The volume of services that access HDFS is too large and therefore NameNode is overloaded.

Procedure

Obtain alarm information.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. In the alarm list, click the alarm.

Step 2 Check the alarm. Obtain the alarm generation time from **Generated**. Obtain the host name of the NameNode node involved in this alarm from the **HostName** information of **Location**. Then obtain the name of the NameService node involved in this alarm from the **NameServiceName** information of **Location**.

Check whether the threshold is too small.

Step 3 Check the status of the services that depend on HDFS. Check whether the services run slowly or task execution times out.

- If yes, go to **Step 8**.
- If no, go to **Step 4**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **RPC**, and select **Average Time of Active NameNode RPC Queuing** and click **OK**.

Figure 10-24 Average Time of Active NameNode RPC Queuing

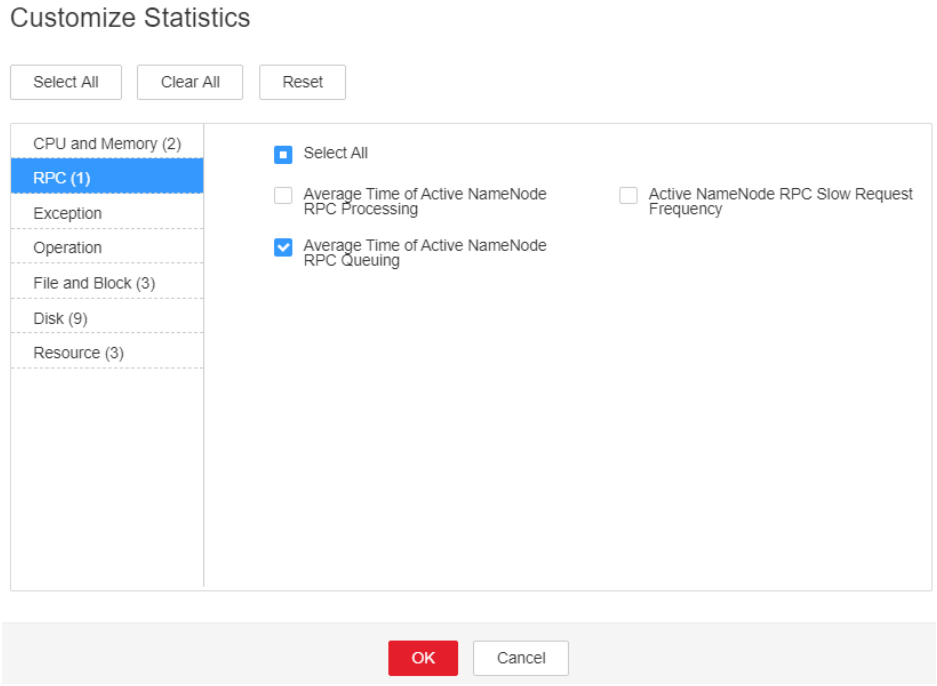

Step 5 On the **Average Time of Active NameNode RPC Queuing** monitoring page, obtain the value of the NameService node involved in this alarm.

Step 6 On the FusionInsight Manager portal, choose **O&M > Alarm > Thresholds > *Name of the desired cluster* > HDFS**. Locate **Average Time of Active NameNode RPC Queuing** and click the **Modify** in the **Operation** column of the default rule. The **Modify Rule** page is displayed. Change **Threshold** to 150% of the monitored value. Click **OK** to save the new threshold.

Step 7 Wait for 1 minute and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check whether the CPU performance of the NameNode node is sufficient.

Step 8 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and check whether **ALM-12016 HDFS NameNode Memory Usage Exceeds the Threshold** is generated.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

Step 9 Handle **ALM-12016 CPU Usage Exceeds the Threshold** by taking recommended actions.

Step 10 Wait for 10 minutes and check whether alarm 14022 is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Check whether the memory of the NameNode node is too small.

Step 11 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and check whether **ALM-14007 HDFS NameNode Memory Usage Exceeds the Threshold** is generated.

- If yes, go to **Step 12**.
- If no, go to **Step 14**.

Step 12 Handle **ALM-14007 CPU Usage Exceeds the Threshold** by taking recommended actions.

Step 13 Wait for 10 minutes and check whether alarm 14022 is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Check whether NameNode parameters are configured properly.

Step 14 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS** > **Configurations** > **All Configurations**. Search for parameter **dfs.namenode.handler.count** and view its value. If the value is less than or equal to 128, change it to **128**. If the value is greater than 128 but less than 192, change it to **192**.

Step 15 Search for parameter **ipc.server.read.threadpool.size** and view its value. If the value is less than 5, change it to **5**.

Step 16 Click **Save**, and click **OK**.

Step 17 On the **Instance** page of HDFS, select the standby NameNode of NameService involved in this alarm and choose **More** > **Restart Instance**. Enter the password and click **OK**. Wait until the standby NameNode is started up.

Step 18 On the **Instance** page of HDFS, select the active NameNode of NameService involved in this alarm and choose **More** > **Restart Instance**. Enter the password and click **OK**. Wait until the active NameNode is started up.

Step 19 Wait for 1 hour and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 20**.

Check whether the HDFS workload changes and reduce the workload properly.

Step 20 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **HDFS**. Click the drop-down menu in the upper right corner of **Chart**, click **Customize**, select **Average Time of Active NameNode RPC Queuing** and click **OK**.

Figure 10-25 Average Time of Active NameNode RPC Queuing

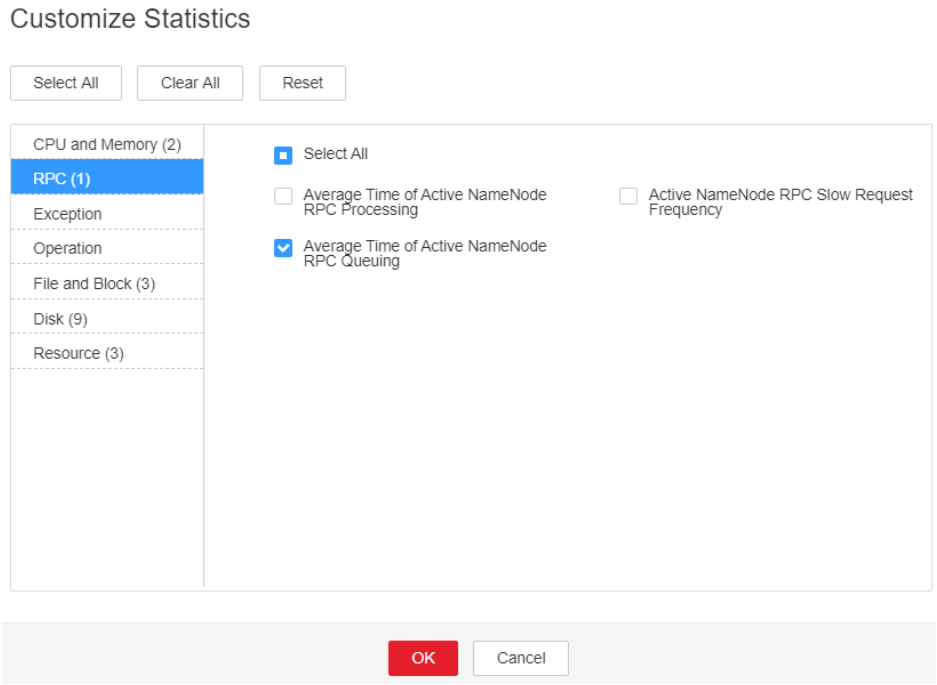

Step 21 Click . The **Details** page is displayed.

Step 22 Set the monitoring data display period, from 5 days before the alarm generation time to the alarm generation time. Click **OK**.

Step 23 On the **Average RPC Queuing Time** monitoring page, check whether the point in time when the queuing time increases abruptly exists.

- If yes, go to **Step 24**.
- If no, go to **Step 27**.

Step 24 Confirm and check the point in time. Check whether a new task frequently accesses HDFS and whether the access frequency can be reduced.

Step 25 If a Balancer task starts at the point in time, stop the task or specify a node for the task to reduce the HDFS workload.

Step 26 Wait for 1 hour and then check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 27**.

Collect fault information.

Step 27 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 28 Select **HDFS** in the required cluster from the **Service**.

Step 29 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 30 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.109 ALM-14023 Percentage of Total Reserved Disk Space for Replicas Exceeds the Threshold

Description

The system checks the percentage of total reserved disk space for replicas (Total reserved disk space for replicas/(Total reserved disk space for replicas + Total remaining disk space)) every 30 seconds and compares the actual percentage with the threshold (**90%** by default). This alarm is generated when the percentage of total reserved disk space for replicas exceeds the threshold for multiple consecutive times (**Trigger Count**).

The alarm is cleared in the following two scenarios: The value of **Trigger Count** is **1** and the percentage of total reserved disk space for replicas is less than or equal to the threshold; the value of **Trigger Count** is greater than **1** and the percentage of total reserved disk space for replicas is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14023	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
NameServiceName	Specifies the NameService service for which the alarm is generated.

Name	Meaning
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The performance of writing data to HDFS is affected. If all remaining DataNode space is reserved for replicas, writing HDFS data fails.

Possible Causes

- The alarm threshold is improperly configured.
- The disk space configured for the HDFS cluster is insufficient.
- The volume of services that access HDFS is too large and therefore DataNode is overloaded.

Procedure

Check whether the alarm threshold is appropriate.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> HDFS > Disk > Percentage of Reserved Space for Replicas of Unused Space** to check whether the alarm threshold is appropriate. (The default threshold is **90%**. Users can change it as required.)

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> HDFS > Disk > Percentage of Reserved Space for Replicas of Unused Space** and Click **Modify,** change the threshold based on the actual usage.

Figure 10-26 Modify Thresholds

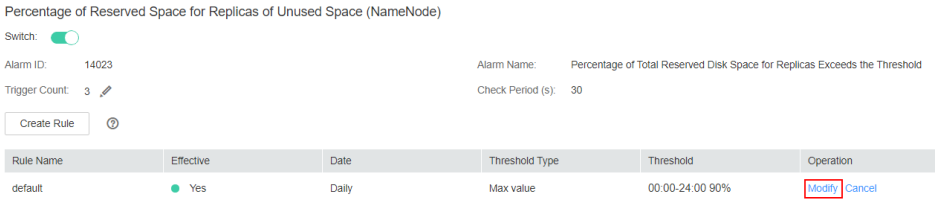

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether an alarm indicating insufficient disk space is generated.

Step 4 On the FusionInsight Manager portal, check whether **ALM-14001 HDFS Disk Usage Exceeds the Threshold** or **ALM-14002 DataNode Disk Usage Exceeds the Threshold** exists on the **O&M > Alarm > Alarms** page.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Handle the alarm by referring to instructions in **ALM-14001 HDFS Disk Usage Exceeds the Threshold** or **ALM-14002 DataNode Disk Usage Exceeds the Threshold** and check whether the alarm is cleared.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Expand the DataNode capacity.

Step 7 Expand the DataNode capacity.

Step 8 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 10 Select **HDFS** in the required cluster from the **Service**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.110 ALM-14024 Tenant Space Usage Exceeds the Threshold

Description

The system checks the space usage (used space of each directory/space allocated to each directory) of each directory associated with a tenant every hour and compares the space usage of each directory with the threshold set for the directory. This alarm is generated when the space usage exceeds the threshold.

This alarm is cleared when the space usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14024	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TenantName	Specifies the tenant for which the alarm is generated.
DirectoryName	Specifies the directory for which the alarm is generated.
Trigger condition	Specifies the threshold for triggering the alarm.

Impact on the System

This alarm is generated if the space usage of the tenant directory exceeds the custom threshold. File writing to the directory is not affected. If the used space exceeds the maximum storage space allocated to the directory, the HDFS fails to write data to the directory.

Possible Causes

- The alarm threshold is improperly configured.
- The space allocated to the tenant is improper.

Procedure

Check whether the alarm threshold is appropriate.

Step 1 View the alarm location information to obtain the tenant name and tenant directory for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether the storage space threshold configured for the tenant directory for which the alarm is generated is proper. (The default value 90% is a proper value. You can set it based on the site requirements.)

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 On the **Resources** page, click **Modify** to modify or delete the storage space threshold.

Step 4 About one minute later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the space allocated to the tenant is appropriate.

Step 5 On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether the storage space quota of the tenant directory for which the alarm is generated is proper based on the actual service status of the tenant directory.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 On the **Resources** page, click **Modify** to modify the storage space quota.

Step 7 About one minute later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 9 Select **HDFS** in the required cluster and **NodeAgent** under **Manager** from the **Service**.

Step 10 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.111 ALM-14025 Tenant File Object Usage Exceeds the Threshold

Description

The system checks the file object usage (used file objects of each directory/number of file objects allocated to each directory) of each directory associated with a tenant every hour and compares the file object usage of each directory with the threshold set for the directory. This alarm is generated when the file object usage exceeds the threshold.

This alarm is cleared when the file object usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
14025	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
TenantName	Specifies the tenant for which the alarm is generated.
DirectoryName	Specifies the directory for which the alarm is generated.
Trigger condition	Specifies the threshold for triggering the alarm.

Impact on the System

This alarm is generated if the usage of file objects in a tenant directory exceeds the custom threshold. File writing to the directory is not affected. If the number of

used file objects exceeds the maximum number of file objects allocated to the directory, the HDFS fails to write data to the directory.

Possible Causes

- The alarm threshold is improperly configured.
- The maximum number of file objects allocated to the tenant directory is inappropriate.

Procedure

Check whether the alarm threshold is appropriate.

Step 1 View the alarm location information to obtain the tenant name and tenant directory for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether the file object threshold configured for the tenant directory for which the alarm is generated is proper. (The default value 90% is a proper value. You can set it based on the site requirements.)

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 On the **Resources** page, click **Modify** to modify or delete the file object threshold of the tenant directory for which the alarm is generated.

Step 4 About one minute later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the maximum number of file objects allocated to the tenant is appropriate.

Step 5 On the FusionInsight Manager portal, choose the **Tenant Resources** page, select the tenant for which the alarm is generated, and click **Resources**. Check whether the maximum number of file objects configured for the tenant directory for which the alarm is generated is proper based on the actual service status of the tenant directory.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 On the **Resources** page, click **Modify** to modify or delete the maximum number of file objects configured for the tenant directory.

Step 7 About one minute later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 9 Select **HDFS** in the required cluster and **NodeAgent** under **Manager** from the **Service**.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.112 ALM-14026 Blocks on DataNode Exceed the Threshold

Alarm Description

The system checks the number of blocks on each DataNode every 30 seconds. This alarm is generated when the number of blocks on the DataNode exceeds the threshold.

If **Trigger Count** is **1** and the number of blocks on the DataNode is less than or equal to the threshold, this alarm is cleared. If **Trigger Count** is greater than **1** and the number of blocks on the DataNode is less than or equal to 90% of the threshold, this alarm is cleared.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14026	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If this alarm is reported, there are too many blocks on the DataNode. In this case, data writing into the HDFS may fail due to insufficient disk space.

Possible Causes

- The alarm threshold is improperly configured.
- Data skew occurs among DataNodes.
- The disk space configured for the HDFS cluster is insufficient.

Handling Procedure

Change the threshold.

Step 1 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **HDFS**. Then choose **Configurations** > **All Configurations**. On the displayed page, find the **GC_OPTS** parameter under **HDFS->DataNode**.

Step 2 Set the threshold of the DataNode blocks. Specifically, change the value of **Xmx** of the **GC_OPTS** parameter. **Xmx** specifies the memory, and each GB memory supports a maximum of 500,000 DataNode blocks. Set the memory as required. Confirm that **GC_PROFILE** is set to **custom** and save the configuration.

Step 3 Choose **Cluster**, click the name of the desired cluster, and choose **HDFS** > **Instance**. Select the DataNode instance whose status is **Expired**, click **More**, and select **Restart Instance** to make the **GC_OPTS** configuration take effect.

Step 4 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether associated alarms are reported.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and check whether the **ALM-14002 DataNode Disk Usage Exceeds the Threshold** alarm exists.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 Handle the alarm by following the instructions in **ALM-14002 DataNode Disk Usage Exceeds the Threshold** and check whether the alarm is cleared.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 8**.

Expand the DataNode capacity.

Step 8 Expand the DataNode capacity.

Step 9 On FusionInsight Manager, wait for 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

Configuration rules of the DataNode JVM parameter.

Default value of the DataNode JVM parameter **GC_OPTS**:

-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M -XX:MetaspaceSize=128M -XX:MaxMetaspaceSize=128M -XX:+UseConcMarkSweepGC -XX:+CMSParallelRemarkEnabled -XX:CMSInitiatingOccupancyFraction=65 -XX:+PrintGCDetails -Dsun.rmi.dgc.client.gcInterval=0x7FFFFFFFFFFFFE -Dsun.rmi.dgc.server.gcInterval=0x7FFFFFFFFFFFFE -XX:-OmitStackTraceInFastThrow -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10 -XX:GCLogFileSize=1M -Djdk.tls.ephemeralDHKeySize=2048

Average number of saved blocks = Number of HDFS blocks x 3/Number of DataNodes

NOTE

To obtain the number of HDFS blocks, log in to FusionInsight Manager, choose **Cluster** > **Services** > **HDFS**, click **NameNode(*xxx*,Active)** to the right of **NameNode Web UI** to go to the native HDFS web UI, and then view the information in the **Summary** area.

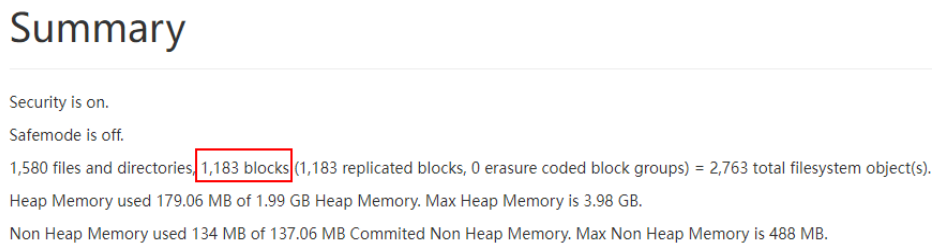

If the average number of blocks on a single DataNode has changed, modify **-Xms2G -Xmx4G -XX:NewSize=128M -XX:MaxNewSize=256M** in the default value. The following table lists the reference values.

Table 10-3 DataNode JVM configuration

Average Number of Blocks in a DataNode Instance	Reference Value
2,000,000	-Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=512M
5,000,000	-Xms12G -Xmx12G -XX:NewSize=1G -XX:MaxNewSize=1G

Xmx specifies memory which corresponds to the threshold of the number of DataNode blocks, and each GB memory supports a maximum of 500,000 DataNode blocks. Set the memory as required.

10.113 ALM-14027 DataNode Disk Fault

Alarm Description

The system checks the disk status on DataNodes every 60 seconds. This alarm is generated when a disk is faulty.

After all faulty disks on the DataNode are recovered, you need to manually clear the alarm and restart the DataNode.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14027	Major	No

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Failed Volumes	Specifies the list of faulty disks.

Impact on the System

If this alarm is reported, there are abnormal disk partitions on the DataNode. This may cause the loss of written files.

Possible Causes

- The hard disk is faulty.
- The disk permissions are configured improperly.

Handling Procedure

Check whether a disk alarm is generated.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** and check whether **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault** exists.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Rectify the fault by referring to the handling procedure of **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault**. Then, check whether the alarm is cleared.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Modify disk permissions.

Step 4 Choose **O&M** > **Alarm** > **Alarms** and view **Location** and **Additional Information** of the alarm to obtain the location of the faulty disk.

Step 5 Log in to the node for which the alarm is generated as user **root**. Go to the directory where the faulty disk is located, and run the **ll** command to check whether the permission of the faulty disk is **711** and whether the user is **omm**.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 Modify the permission of the faulty disk. For example, if the faulty disk is **data1**, run the following commands:

chown omm:wheel data1

chmod 711 data1

Step 7 In the alarm list on Manager, click **Clear** in the **Operation** column of the alarm to manually clear the alarm. Choose **Cluster** > **Services** > **HDFS** > **Instance**, select the DataNode, choose **More** > **Restart Instance**, wait for 5 minutes, and check whether a new alarm is reported.

- If no, no further action is required.
- If yes, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **HDFS** and **OMS** for the target cluster.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault is rectified, the system does not automatically clear this alarm and you need to manually clear the alarm.

Related Information

None

10.114 ALM-14028 Number of Blocks to Be Supplemented Exceeds the Threshold

Alarm Description

The system checks the number of blocks to be supplemented every 30 seconds and compares the number with the threshold. The number of blocks to be

supplemented has a default threshold. This alarm is generated when the number of blocks to be supplemented exceeds the threshold.

You can change the threshold specified by **Blocks Under Replicated (NameNode)** by choosing **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **HDFS** > **File and Block**.

If **Trigger Count** is set to **1** and the number of blocks to be supplemented is less than or equal to the threshold, this alarm is cleared. If **Trigger Count** is greater than **1** and the number of blocks to be supplemented is less than or equal to 90% of the threshold, this alarm is cleared.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14028	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Data stored in HDFS is lost. HDFS may enter the security mode and cannot provide write services. Lost block data cannot be restored.

Possible Causes

- The DataNode instance is abnormal.
- Data is deleted.
- The number of replicas written into the file is greater than the number of DataNodes.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, check whether alarm **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold** is generated.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 Rectify the fault according to the handling procedure of **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold**. Five minutes later, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Log in to the HDFS client as user **root**. The user password is defined by the user before the installation. Contact the MRS cluster administrator to obtain the password. Run the following commands:

- Security mode:

  cd *Client installation directory*

  source bigdata_env

  kinit hdfs

- Normal mode:

  su - omm

  cd *Client installation directory*

  source bigdata_env

Step 4 Run the **hdfs fsck / >> fsck.log** command to obtain the status of the current cluster.

Step 5 Run the following command to count the number (*M*) of blocks to be replicated:

cat fsck.log | grep "Under-replicated"

Step 6 Run the following command to count the number (*N*) of blocks to be replicated in the **/tmp/hadoop-yarn/staging/** directory:

cat fsck.log | grep "Under replicated" | grep "/tmp/hadoop-yarn/staging/" | wc -l

NOTE

/tmp/hadoop-yarn/staging/ is the default directory. If the directory is modified, obtain it from the configuration item **yarn.app.mapreduce.am.staging-dir** in the **mapred-site.xml** file.

Step 7 Check whether the percentage of *N* is greater than 50% (N/M > 50%).

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Run the following command to reconfigure the number of file replicas in the directory (set the number of file replicas to the number of DataNodes or the default number of file replicas):

hdfs dfs -setrep -w *Number of file replicas*/tmp/hadoop-yarn/staging/

NOTE

To obtain the default number of file replicas:

Log in to FusionInsight Manager, choose **Cluster** > **Services** > **HDFS** > **Configurations** > **All Configurations**, and search for the **dfs.replication** parameter. The value of this parameter is the default number of file replicas.

Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.115 ALM-14029 Number of Blocks in a Replica Exceeds the Threshold

Alarm Description

The system checks the number of blocks in a single replica every four hours and compares the number with the threshold. There is a threshold for the number of blocks in a single replica. This alarm is generated when the actual number of blocks in a single replica exceeds the threshold.

This alarm is cleared when the number of blocks to be supplemented is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14029	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
NameServiceName	Specifies the NameService for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Replica data is prone to be lost when a node is faulty. Too many files of a single replica affect the security of the HDFS file system.

Possible Causes

- The DataNode is faulty.
- The disk is faulty.
- Files are written to a single replica.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, check whether alarm **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold** is generated.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 Rectify the fault according to the handling procedure of **ALM-14003 Number of Lost HDFS Blocks Exceeds the Threshold**. In the next detection period, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Check whether files of a single replica have been written into the service.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 Log in to the HDFS client as user **root**. The user password is defined by the user before the installation. Contact the MRS cluster administrator to obtain the password. Run the following commands:

- Security mode:

  cd *Client installation directory*

  source bigdata_env

  kinit hdfs

- Normal mode:

  su - omm

  cd *Client installation directory*

  source bigdata_env

Step 5 Run the following command on the client node to increase the number of replicas for a single replica file:

hdfs dfs -setrep -w *file replica number file name or file path*

Step 6 In the next detection period, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.116 ALM-14030 HDFS Allows Write of Single-Replica Data

Alarm Description

This alarm is generated when **dfs.single.replication.enable** is set to **true**, indicating that HDFS is configured to allow write of single-replica data.

This alarm is cleared when this function is disabled on HDFS.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14030	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Impact on the System

If this configuration is enabled on the server and the number of HDFS replicas configured on the client is 1, single-replica data can be written to HDFS. Data of a single replica may be lost. Therefore, the system does not allow write of single-replica data by default. If a service requires single-replica data write to a directory, modify the HDFS configuration item **dfs.single.replication.exclude.pattern**.

Possible Causes

The HDFS configuration item **dfs.single.replication.enable** is set to **true**.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **HDFS**. On the page that is displayed, click the **Configurations** tab then the **All Configurations** sub-tab.

Step 2 Search for **dfs.single.replication.enable** in the search box, change the value of the configuration item to **false**, and click **Save**.

Step 3 Wait for about 10 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 6 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.117 ALM-14031 DataNode Process Is Abnormal

Alarm Description

The DataNode process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14031	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

Handling Procedure

Check whether the process is in the D, Z, or T state.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.

- If the alarm is not in the list, no further action is required.
- If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in **Step 2**.

Step 2 Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.

Step 3 Run the following command to check the process state:

ps ww -eo stat,cmd| grep -w org.apache.hadoop.hdfs.server.datanode.DataNode | grep -v grep | awk '{print $1}'

Step 4 Check whether the command output contains any abnormal state (D, Z, or T).

- If the output contains any abnormal state, go to **Step 5**.
- If the output does not contain abnormal states, go to **Step 7**.

Step 5 Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.118 ALM-14032 JournalNode Process Is Abnormal

Alarm Description

The JournalNode process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14032	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

Handling Procedure

Check whether the process is in the D, Z, or T state.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.

- If the alarm is not in the list, no further action is required.
- If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in **Step 2**.

Step 2 Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.

Step 3 Run the following command to check the process state:

ps ww -eo stat,cmd| grep -w org.apache.hadoop.hdfs.qjournal.server.JournalNode | grep -v grep | awk '{print$1}'

Step 4 Check whether the command output contains any abnormal state (D, Z, or T).

- If the output contains any abnormal state, go to **Step 5**.
- If the output does not contain abnormal states, go to **Step 7**.

Step 5 Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.119 ALM-14033 ZKFC Process Is Abnormal

Alarm Description

The ZKFC process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14033	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

Handling Procedure

Check whether the process is in the D, Z, or T state.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.

- If the alarm is not in the list, no further action is required.
- If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in **Step 2**.

Step 2 Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.

Step 3 Run the following command to check whether the process state is abnormal:

ps ww -eo stat,cmd| grep -w org.apache.hadoop.hdfs.tools.DFSZKFailoverController | grep -v grep | awk '{print$1}'

Step 4 Check whether the command output contains any abnormal state (D, Z, or T).

- If the output contains any abnormal state, go to **Step 5**.
- If the output does not contain abnormal states, go to **Step 7**.

Step 5 Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.120 ALM-14034 Router Process Is Abnormal

Alarm Description

The Router process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14034	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

Handling Procedure

Check whether the process is in the D, Z, or T state.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.

- If the alarm is not in the list, no further action is required.
- If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in **Step 2**.

Step 2 Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.

Step 3 Run the following command to check whether the process state is abnormal:

ps ww -eo stat,cmd| grep -w org.apache.hadoop.hdfs.server.federation.router.DFSRouter | grep -v grep | awk '{print$1}'

Step 4 Check whether the command output contains any abnormal state (D, Z, or T).

- If the output contains any abnormal state, go to **Step 5**.
- If the output does not contain abnormal states, go to **Step 7**.

Step 5 Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.121 ALM-14035 HttpFS Process Is Abnormal

Alarm Description

The HttpFS process checks the process status every 20 seconds. This alarm is generated when the process status is abnormal and does not recover for a long time.

This alarm is cleared when the process status recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
14035	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the process status is abnormal, the process cannot provide services properly. As a result, the entire service may become abnormal.

Possible Causes

The host responds slowly to I/O (disk I/O and network I/O) requests and some processes are in the D state and Z state. The process may also be suspended and enter the T state.

Handling Procedure

Check whether the process is in the D, Z, or T state.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Wait for about 10 minutes and check whether the alarm is automatically cleared.

- If the alarm is not in the list, no further action is required.
- If the alarm is in the list, view the alarm details and record the IP address of the host where the alarm is generated. Run the command in **Step 2**.

Step 2 Log in to the host where the alarm is generated as the **root** user and run the **su - omm** command to switch to the **omm** user.

Step 3 Run the following command to check whether the process state is abnormal:

ps ww -eo stat,cmd| grep -w org.apache.hadoop.fs.http.server.HttpFSServerWebServer | grep -v grep | awk '{print$1}'

Step 4 Check whether the command output contains any abnormal state (D, Z, or T).

- If the output contains any abnormal state, go to **Step 5**.
- If the output does not contain abnormal states, go to **Step 7**.

Step 5 Switch to user **root** and run the **reboot** command to restart the host for which the alarm is generated. (Restarting a host is risky. Ensure that the service process is normal after the restart.)

Step 6 Wait 5 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HDFS** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.122 ALM-16000 Percentage of Sessions Connected to the HiveServer to Maximum Number Allowed Exceeds the Threshold

Description

The system detects the percentage of sessions connected to the HiveServer to the maximum number of allowed sessions every 30 seconds. This indicator can be viewed on the **Cluster** > *Name of the desired cluster* > **Services** > **Hive > Instance** > *HiveServer instance***.** This alarm is generated when the percentage exceeds the default value **90%**.

To change the threshold, choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* > **Hive > Percentage of Sessions Connected to the HiveServer to Maximum Number of Sessions Allowed by the HiveServer**.

When the **Trigger Count** is 1, this alarm is cleared when the percentage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the percentage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16000	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If a connection alarm is generated, too many sessions are connected to Hive and new connections are unavailable.

Possible Causes

Too many clients are connected to HiveServer.

Procedure

Increase the maximum number of connections to Hive.

Step 1 On the FusionInsight Manager portal, Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations >All Configurations**.

Step 2 Search for **hive.server.session.control.maxconnections** and increase the value of this parameter. If the value of this parameter is **A**, the threshold is **B**, and the number of sessions connected to the HiveServer is **C**, adjust the value of this parameter according to **A x B > C**. To view the number of sessions connected to the HiveServer, check the value of **Statistics for Sessions of the HiveServer** on the Hive monitoring page.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 5 Select **Hive** in the required cluster from the **Service**.

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.123 ALM-16001 Hive Warehouse Space Usage Exceeds the Threshold

Description

This alarm is generated when the Hive warehouse space usage exceeds the specified threshold (85% by default). The system checks the Hive data warehouse space usage every 30s. The indicator **Percentage of HDFS Space Used by Hive to the Available Space** can be viewed on the Hive service monitoring page.

To change the threshold, choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> Hive > Percentage of HDFS Space Used by Hive to the Available Space**.

When the **Trigger Count** is 1, this alarm is cleared when the Hive warehouse space usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the Hive warehouse space usage is less than or equal to 90% of the threshold.

NOTE

The administrator can reduce the warehouse space usage by expanding the warehouse capacity or releasing the used space.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16001	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The system fails to write data, which causes data loss.

Possible Causes

- The upper limit of the HDFS capacity available for Hive is too small.
- The HDFS space is insufficient.
- Some data nodes break down.

Procedure

Expand the system configuration.

Step 1 Analyze the cluster HDFS capacity usage and increase the upper limit of the HDFS capacity available for Hive.

Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations > All Configurations**, find **hive.metastore.warehouse.size.percent**, and increase its value so that larger HDFS capacity will be available for Hive. Assume that the value of the configuration item is A, the total HDFS storage space is B, the threshold is C, and the HDFS space used by Hive is D. The adjustment policy is A x B x C > D. The total HDFS storage space can be viewed on the HDFS NameNode page. The HDFS space used by Hive can be viewed on the Hive monitoring page.

Step 2 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Expand the system.

Step 3 Expand the system.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the data node is normal.

Step 5 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms**.

Step 6 Check whether "ALM-12006 Node Fault", "ALM-12007 Process Fault", or "ALM-14002 DataNode Disk Usage Exceeds the Threshold" exist.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Clear the alarm by following the steps provided in "ALM-12006 Node Fault", "ALM-12007 Process Fault", and "ALM-14002 DataNode Disk Usage Exceeds the Threshold".

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 10 Select **Hive** in the required cluster from the **Service**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.124 ALM-16002 Hive SQL Execution Success Rate Is Lower Than the Threshold

Description

The system checks the percentage of the HQL statements that are executed successfully in every 30 seconds. The formula is: Percentage of HQL statements that are executed successfully = Number of HQL statements that are executed successfully by Hive in a specified period/Total number of HQL statements that are executed by Hive. This indicator can be viewed on the **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** > *HiveServer instance* **.** The default threshold of the percentage of HQL statements that are executed successfully is **90%**. An alarm is reported when the percentage is lower than the **90%**. Users can view the name of the host where an alarm is generated in the location information about the alarm. The IP address of the host is the IP address of the HiveServer node.

Users can modify the threshold by choosing **O&M > Alarm > Thresholds >** *Name of the desired cluster* > **Hive** > **Percentage of HQL Statements That Are Executed Successfully by Hive**.

This alarm is cleared when the execution success rate is higher than 110% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16002	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The system configuration and performance cannot meet service processing requirements.

Possible Causes

- A syntax error occurs in HQL statements.
- The HBase service is abnormal when a Hive on HBase task is performed.
- The Spark service is abnormal when a Hive on Spark task is performed.
- The dependent basic services, such as HDFS, Yarn, and ZooKeeper, are abnormal.

Procedure

Check whether the HQL statements comply with syntax.

Step 1 On the FusionInsight Manager page, choose **O&M** > **Alarm** to view the alarm details and obtain the node where the alarm is generated.

Step 2 Use the Hive client to log in to the HiveServer node where an alarm is reported. Query the HQL syntax provided by Apache, and check whether the HQL commands are correct.

- If yes, go to **Step 4**.
- If no, go to **Step 3**.

NOTE

To view the user who runs an incorrect statement, you can download the hiveserver audit log file of the HiveServer node where this alarm is generated. **Start Data** and **End Data** are 10 minutes before and after the alarm generation time respectively. Open the log file and search for the **Result=FAIL** keyword to filter the log information about the incorrect statement, and then view the user who runs the incorrect statement according to **UserName** in the log information.

Step 3 Enter the correct HQL statements, and check whether the command can be properly executed.

- If yes, go to **Step 12**.
- If no, go to **Step 4**.

Check whether the HBase service is abnormal.

Step 4 Check whether an Hive on HBase task is performed with the user who runs the HQL command.

- If yes, go to **Step 5**.
- If no, go to **Step 8**.

Step 5 On the FusionInsight Manager page, click **Cluster** > *Name of the desired cluster* > **Services**, check whether the HBase service is normal in the service list.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 Choose **O&M** > **Alarm**, check the related alarms displayed on the alarm page and clear them according to related alarm help.

Step 7 Enter the correct HQL statements, and check whether the command can be properly executed.

- If yes, go to **Step 12**.
- If no, go to **Step 8**.

Check whether the HDFS, Yarn, and ZooKeeper are normal.

Step 8 On the FusionInsight Manager portal, click **Cluster** > *Name of the desired cluster* > **Services**.

Step 9 In the service list, check whether the services, such as HDFS, Yarn, and ZooKeeper are normal.

- If yes, go to **Step 12**.
- If no, go to **Step 10**.

Step 10 Check the related alarms displayed on the alarm page and clear them according to related alarm help.

Step 11 Enter the correct HQL statements, and check whether the command can be properly executed.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 After 1 minute, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collect fault information.

Step 13 On the FusionInsight Manager home page, choose **O&M** > **Log** > **Download**.

Step 14 Select the following nodes in the required cluster from the **Service**:

- MapReduce
- Hive

Step 15 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.125 ALM-16003 Background Thread Usage Exceeds the Threshold

Description

The system checks the background thread usage in every 30 seconds. This alarm is generated when the usage of the background thread pool of Hive exceeds the threshold, 90% by default.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16003	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

There are too many background threads, so the newly submitted task cannot run in time.

Possible Causes

The usage of the background thread pool of Hive is excessively high when:

- There are many tasks executed in the background thread pool of HiveServer.
- The capacity of the background thread pool of HiveServer is too small.

Procedure

Check the number of tasks executed in the background thread pool of HiveServer.

Step 1 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive**. On the displayed page, click **HiveServer Instance** and check values of **Background Thread Count** and **Background Thread Usage**.

Figure 10-27 Background

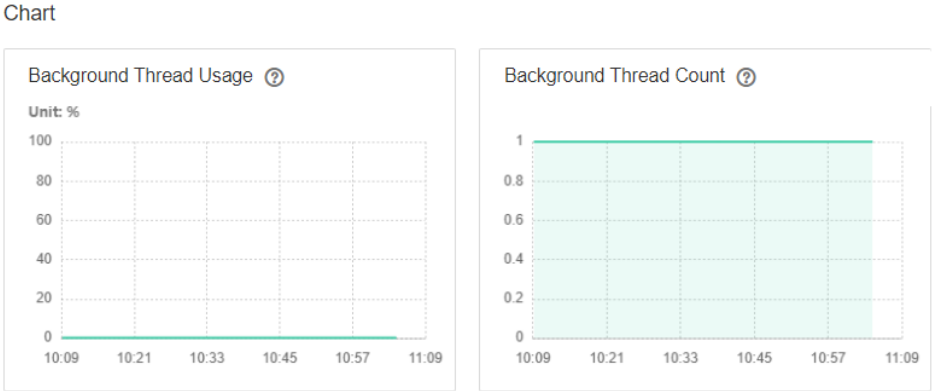

Step 2 Check whether the number of background threads in the latest half an hour is excessively high. (By default, the queue number is 100, and the thread number is considered as high if it is 90 or larger.)

- If it is, go to **Step 3**.
- If it is not, go to **Step 5**.

Step 3 Adjust the number of tasks submitted to the background thread pool. (For example, cancel some time-consuming tasks with low performance.)

Step 4 Check whether the values of Background Thread Count and Background Thread Usage decrease.

- If it is, go to **Step 7**.
- If it is not, go to **Step 5**.

Check the capacity of the HiveServer background thread pool.

Step 5 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive**. On the displayed page, click **HiveServer Instance** and check values of Background Thread Count and Background Thread Usage.

Figure 10-28 Background

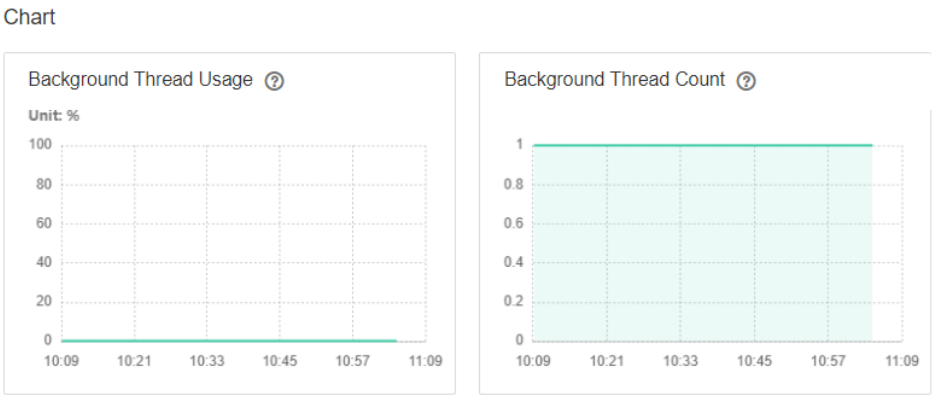

Step 6 Increase the value of **hive.server2.async.exec.threads** in the **${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/1_23_HiveServer/etc/hive-site.xml** file. For example, increase the value by 20%.

Step 7 Save the modification.

Step 8 Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 10 Select **Hive** in the required cluster from the **Service**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.126 ALM-16004 Hive Service Unavailable

Description

This alarm is generated when the HiveServer service is unavailable. The system checks the HiveServer service status every 60 seconds.

This alarm is cleared when the HiveServer service is normal.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16004	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system cannot provide data loading, query, and extraction services.

Possible Causes

- Hive service unavailability may be related to the faults of the Hive process as well as basic services, such as ZooKeeper, Hadoop distributed file system (HDFS), Yarn, and DBService.
  - The ZooKeeper service is abnormal.
  - The HDFS service is abnormal.
  - The Yarn service is abnormal.
  - The DBService service is abnormal.
  - The Hive service process is abnormal. If the alarm is caused by Hive process fault, the alarm report has a delay of about 5 minutes.
- The network communication between the Hive and basic services is interrupted.
- The permission on the HDFS temporary directory of Hive is abnormal.
- The local disk space of the Hive node is insufficient.

Procedure

Check the HiveServer/MetaStore process status.

Step 1 On the FusionInsight Manager portal, click **Cluster >** *Name of the desired cluster* **> Services** > **Hive** > **Instance**. In the Hive instance list, check whether the HiveServer or MetaStore instances are in the Unknown state.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 In the Hive instance list, choose **More** > **Restart Instance** to restart the HiveServer/MetaStore process.

Step 3 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the ZooKeeper service status.

Step 4 On the FusionInsight Manager, check whether the alarm list contains **Process Fault**.

- If yes, go to **Step 5**.
- If no, go to **Step 8**.

Step 5 In the **Process Fault**, check whether **ServiceName** is **ZooKeeper**.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 Rectify the fault by following the steps provided in "ALM-12007 Process Fault".

Step 7 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check the HDFS service status.

Step 8 On the FusionInsight Manager, check whether the alarm list contains **HDFS Service Unavailable**.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

Step 9 Rectify the fault by following the steps provided in "ALM-14000 HDFS Service Unavailable".

Step 10 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Check the Yarn service status.

Step 11 In FusionInsight Manager alarm list, check whether **Yarn Service Unavailable** is generated.

- If yes, go to **Step 12**.
- If no, go to **Step 14**.

Step 12 Rectify the fault. For details, see "ALM-18000 Yarn Service Unavailable".

Step 13 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Check the DBService service status.

Step 14 In FusionInsight Manager alarm list, check whether **DBService Service Unavailable** is generated.

- If yes, go to **Step 15**.
- If no, go to **Step 17**.

Step 15 Rectify the fault. For details, see "ALM-27001 DBService Service Unavailable".

Step 16 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 17**.

Check the network connection between the Hive and ZooKeeper, HDFS, Yarn, and DBService.

Step 17 On the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive**.

Step 18 Click **Instance**.

The HiveServer instance list is displayed.

Step 19 Click **Host Name** in the row of **HiveServer**.

The active HiveServer host status page is displayed.

Step 20 Record the IP address under **Basic Information**.

Step 21 Use the IP address obtained in **Step 20** to log in to the host where the active HiveServer runs as user **omm**.

Step 22 Run the **ping** command to check whether communication between the host that runs the active HiveServer and the hosts that run the ZooKeeper, HDFS, Yarn, and DBService services is normal. (Obtain the IP addresses of the hosts that run the ZooKeeper, HDFS, Yarn, and DBService services in the same way as that for obtaining the IP address of the active HiveServer.)

- If yes, go to **Step 31**.
- If no, go to **Step 23**.

Step 23 Contact the administrator to restore the network.

Step 24 In the alarm list, check whether **Hive Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 31**.

Check the permission on the HDFS temporary directory.

Step 25 Log in to the node where the HDFS client is located and run the following command to go to the HDFS client installation directory:

cd *Client installation directory*

source bigdata_env

kinit *user with the supergroup permission* (Skip this step for common clusters.)

Step 26 Run the following command to check whether the permission on the data warehouse directory is 770:

hdfs dfs -ls /tmp | grep hive-scratch

- If yes, go to **Step 29**.
- If no, go to **Step 27**.

Step 27 Run the following command to restore the default data warehouse permission:

hdfs dfs -chmod 770 /tmp/hive-scratch

Step 28 Wait for several minutes and check whether the Hive Service Unavailable alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 29**.

Check whether the local disk space is normal.

Step 29 Run the **df -h** command to check the root directory and check whether the disk usage of the **/srv**, **/var**, and **/opt** directories exceeds 95%.

- If yes, go to **Step 30**.
- If no, go to **Step 31**.

Step 30 Clear unnecessary information in the corresponding directory to ensure that the available disk space is greater than 80%. Wait for several minutes and check whether the Hive Service Unavailable alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 31**.

Collect fault information.

Step 31 On the FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 32 Select the following nodes in the required cluster from the **Service**:

- ZooKeeper
- HDFS
- Yarn
- DBService
- Hive

Step 33 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 34 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.127 ALM-16005 The Heap Memory Usage of the Hive Process Exceeds the Threshold

Description

The system checks the Hive service status every 30 seconds. The alarm is generated when the heap memory usage of an Hive service exceeds the threshold (95% of the maximum memory).

Users can choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> Hive** to change the threshold.

The alarm is cleared when the heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16005	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

Impact on the System

When the heap memory usage of Hive is overhigh, the performance of Hive task operation is affected. In addition, a memory overflow may occur so that the Hive service is unavailable.

Possible Causes

The heap memory of the Hive instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16005**. Then check the role name in **Location** and confirm the IP adress of the instance.

- If the role for which the alarm is generated is HiveServer, go to **Step 2**.
- If the role for which the alarm is generated is MetaStore, go to **Step 3**.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* **> Services** > **Hive** > **Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **HiveServer Memory Usage Statistics** and click **OK**, check whether the used heap memory of the HiveServer service reaches the threshold(default value: 95%) of the maximum heap memory specified for HiveServer.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-29 HiveServer Memory Usage Statistics

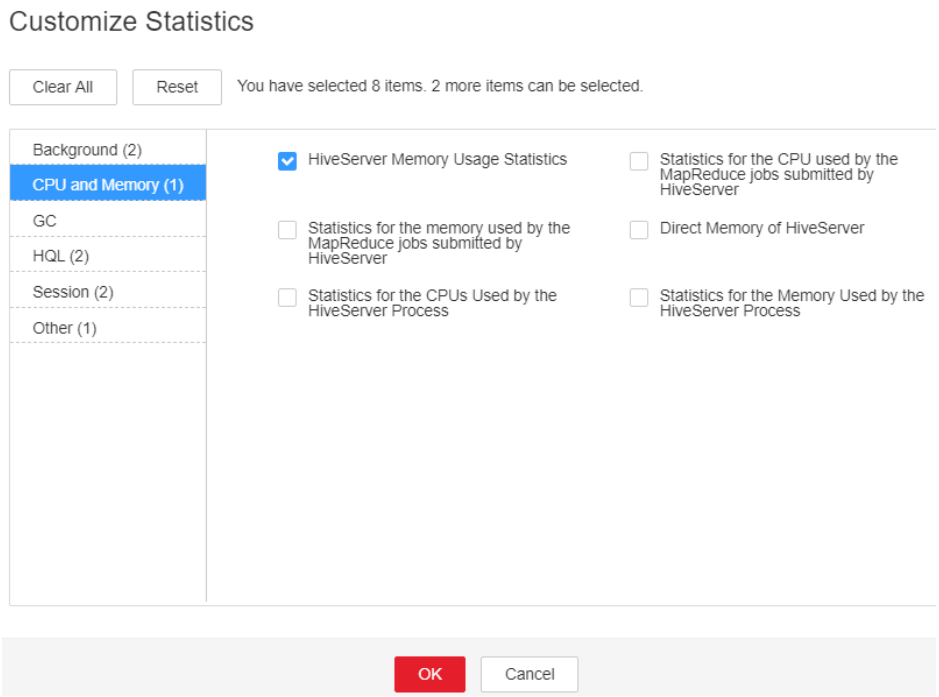

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **MetaStore Memory Usage Statistics** and click **OK**, check whether the used heap memory of the MetaStore service reaches the threshold(default value: 95%) of the maximum heap memory specified for MetaStore.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-30 MetaStore Memory Usage Statistics

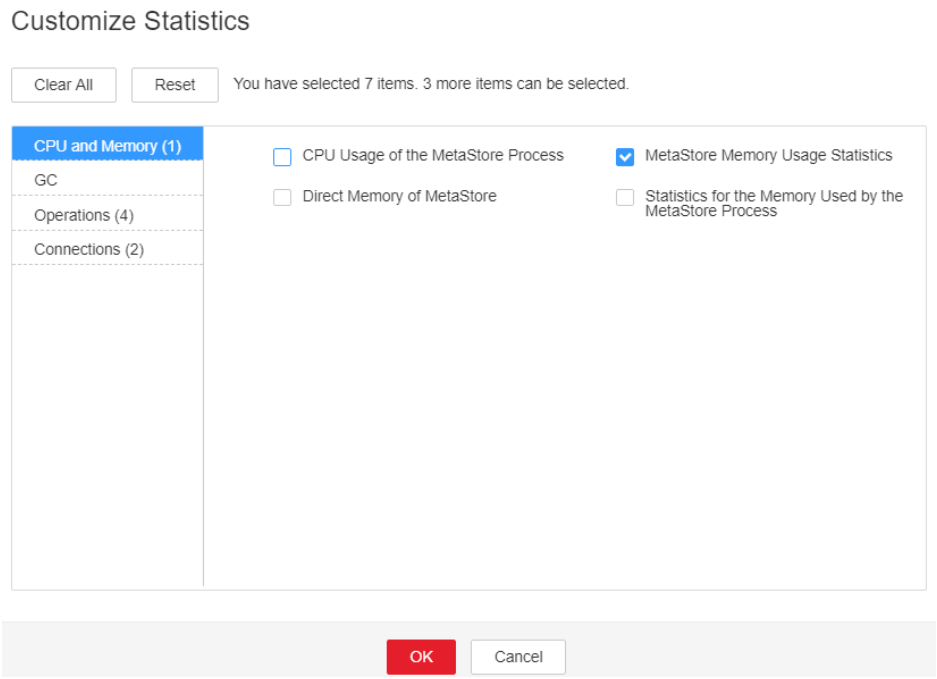

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations > All Configurations**. Choose **HiveServer/MetaStore** > **JVM**. Adjust the value of **-Xmx** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- When the heap memory used by the HiveServer process reaches the threshold (default value: 95%) of the maximum heap memory set by the HiveServer process, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2GB by default, change the value of **-Xmx** to 4GB. You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically. On the FusionInsight Manager home page, choose **O&M**> **Alarm**> **Thresholds** > *Name of the desired cluster* **> Hive** > **CPU and Memory** > **HiveServer Heap Memory Usage Statistics (HiveServer)** to view **Threshold**.

Suggestions for GC parameter settings for the MetaServer:

- When the heap memory used by the MetaStore process reaches the threshold (default value: 95%) of the maximum heap memory set by the MetaStore process, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2GB by default, change the value of **-Xmx** to 4GB. On the FusionInsight Manager home page, choose **O&M**> **Alarm**> **Thresholds** > *Name of the desired cluster* **> Hive** > **CPU and Memory** > **MetaStore Heap Memory Usage Statistics (MetaStore)** to view **Threshold**.
- You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically.

Step 5 Click **More > Restart Service** to restart the service.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **Hive** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.128 ALM-16006 The Direct Memory Usage of the Hive Process Exceeds the Threshold

Description

The system checks the Hive service status every 30 seconds. The alarm is generated when the direct memory usage of an Hive service exceeds the threshold (95% of the maximum memory).

Users can choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* **> Hive** to change the threshold.

The alarm is cleared when the direct memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16006	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the direct memory usage of Hive is overhigh, the performance of Hive task operation is affected. In addition, a memory overflow may occur so that the Hive service is unavailable.

Possible Causes

The direct memory of the Hive instance on the node is overused or the direct memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check direct memory usage.

Step 1 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16006**. Then check the role name in **Location** and confirm the IP adress of the instance.

- If the role for which the alarm is generated is HiveServer, go to **Step 2**.
- If the role for which the alarm is generated is MetaStore, go to **Step 3**.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services > Hive > Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **HiveServer Memory Usage Statistics** and click **OK**, check whether the used direct memory of the HiveServer service reaches the threshold(default value: 95%) of the maximum direct memory specified for HiveServer.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-31 HiveServer Memory Usage Statistics

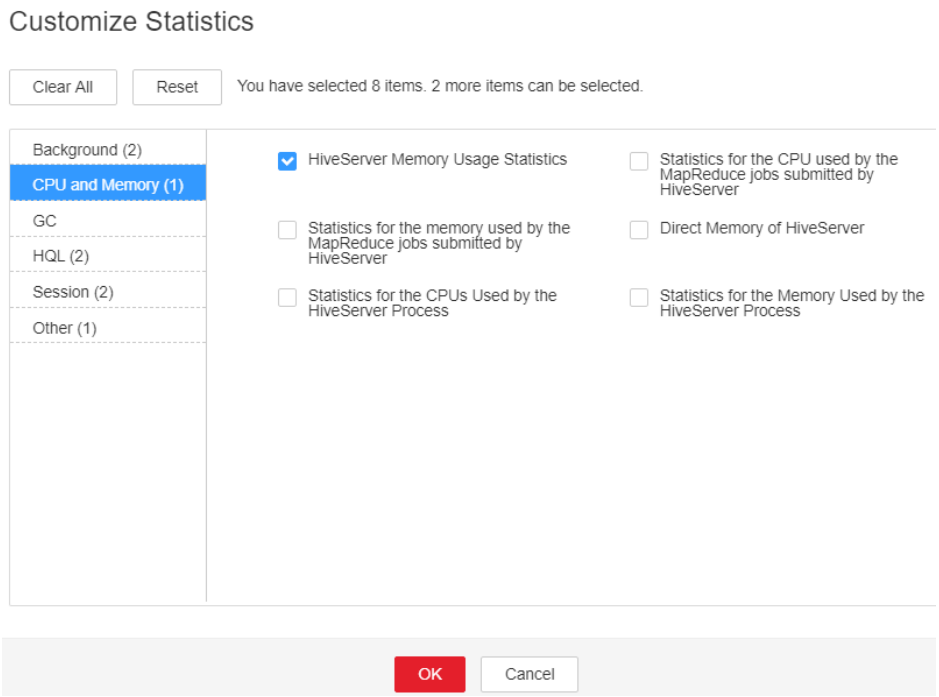

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **MetaStore Memory Usage Statistics** and click **OK**, check whether the used direct memory of the MetaStore service reaches the threshold(default value: 95%) of the maximum direct memory specified for MetaStore.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-32 MetaStore Memory Usage Statistics

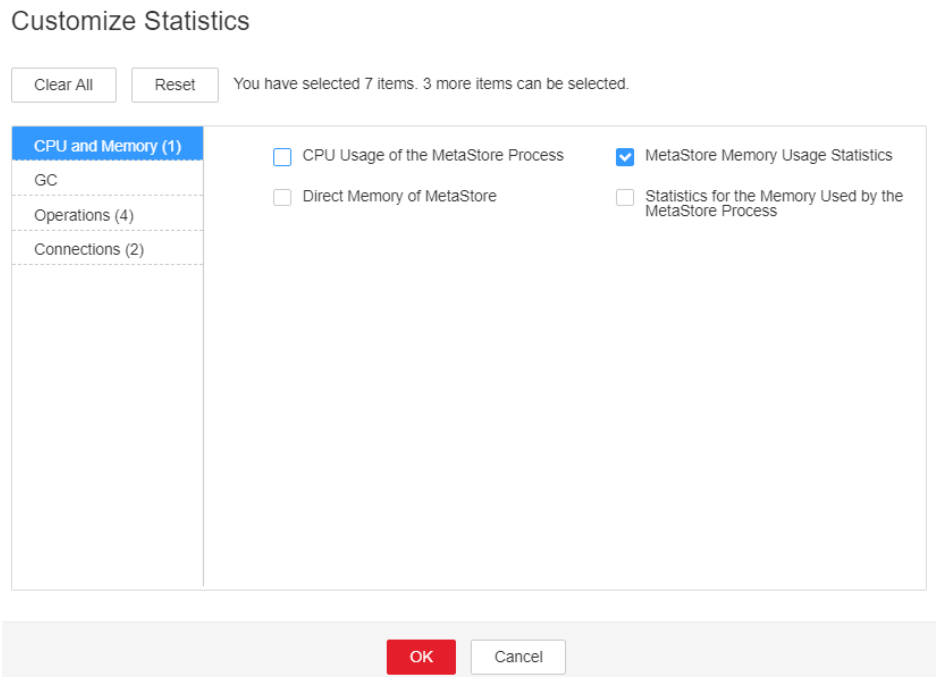

Step 4 On the FusionInsight Manager portal, choose **Cluster** >*Name of the desired cluster*> **Services** > **Hive** > **Configurations > All Configurations**. Choose **HiveServer/MetaStore** > **JVM**. Adjust the value of **-XX:MaxDirectMemorySize** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- It is recommended that you set the value of **-XX:MaxDirectMemorySize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 8 GB, **-XX:MaxDirectMemorySize** is set to 1024 MB. If **-Xmx** is set to 4 GB, **-XX:MaxDirectMemorySize** is set to 512 MB. It is recommended that the value of **-XX:MaxDirectMemorySize** be greater than or equal to 512 MB.

Suggestions for GC parameter settings for the MetaServer:

- It is recommended that you set the value of **-XX:MaxDirectMemorySize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 8 GB, **-XX:MaxDirectMemorySize** is set to 1024 MB. If **-Xmx** is set to 4 GB, **-XX:MaxDirectMemorySize** is set to 512 MB. It is recommended that the value of **-XX:MaxDirectMemorySize** be greater than or equal to 512 MB.

Step 5 Click **More > Restart Service** to restart the service.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **Hive** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.129 ALM-16007 Hive GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) time of the Hive service every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (exceeds 12 seconds for three consecutive checks.) To change the threshold, choose **O&M > Alarm > Thresholds >** *Name of the desired cluster* > **Hive**. This alarm is cleared when the Hive GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16007	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the GC time exceeds the threshold, Hive data read and write are affected.

Possible Causes

The memory of Hive instances is overused, the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC time.

Step 1 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16007**. Then check the role name in **Location** and confirm the IP adress of the instance.

- If the role for which the alarm is generated is HiveServer, go to **Step 2**.
- If the role for which the alarm is generated is MetaStore, go to **Step 3**.

Step 2 On the FusionInsight Manager portal, choose **Cluster** >*Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **GC**, and select **Garbage Collection (GC) Time of HiveServer** and click **OK** to check whether the GC time is longer than 12 seconds.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-33 Garbage Collection (GC) Time of HiveServer

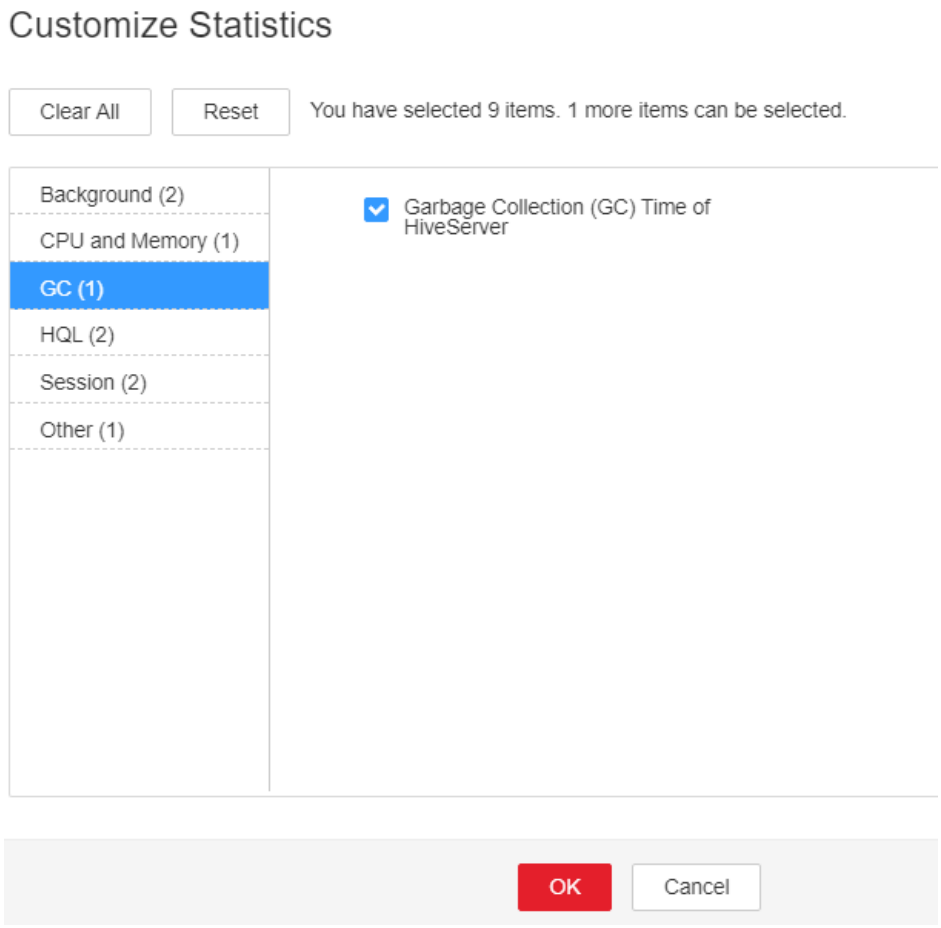

Step 3 On the FusionInsight Manager portal, choose **Cluster** >*Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **GC**, and select **Garbage Collection (GC) Time of MetaStore** and click **OK** to check whether the GC time is longer than 12 seconds.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-34 Garbage Collection (GC) Time of MetaStore

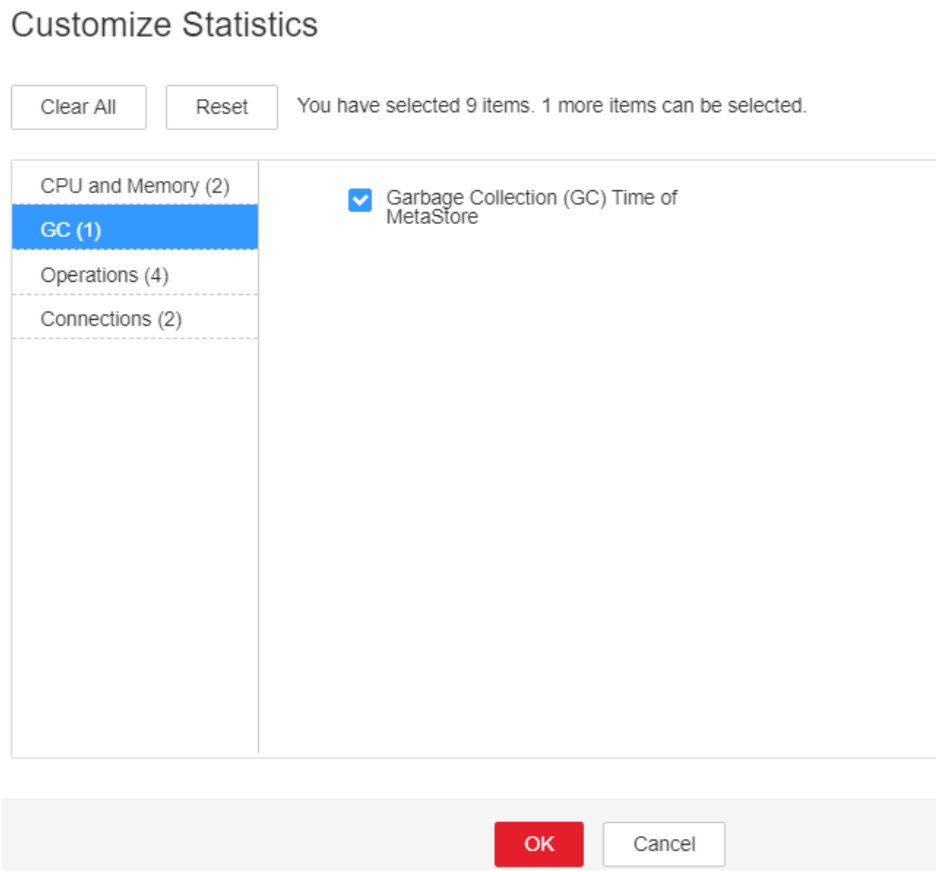

Check the current JVM configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster** >*Name of the desired cluster*> **Services > Hive > Configurations > All Configurations**. Choose **HiveServer/MetaStore** > **JVM**. Adjust the value of **-Xmx** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- When the Hive GC time exceeds the threshold, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2 GB by default, change the value of **-Xmx** to 4 GB.
- You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically.

Suggestions for GC parameter settings for the MetaServer:

- When the Meta GC time exceeds the threshold, change the value of **-Xmx** to twice the default value. For example, if **-Xmx** is set to 2 GB by default, change the value of **-Xmx** to 4 GB.
- You are advised to change the value of **-Xms** to set the ratio of **-Xms** and **-Xmx** to 1:2 to avoid performance problems when JVM dynamically.

Step 5 Click **More > Restart Service** to restart the service.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal of active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 8 In the **Service**, select **Hive** in the required cluster.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.130 ALM-16008 Non-Heap Memory Usage of the Hive Process Exceeds the Threshold

Description

The system checks the Hive service status every 30 seconds. The alarm is generated when the non-heap memory usage of an Hive service exceeds the threshold (95% of the maximum memory).

Users can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Hive** to change the threshold.

The alarm is cleared when the non-heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
16008	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

Impact on the System

When the non-heap memory usage of Hive is overhigh, the performance of Hive task operation is affected. In addition, a memory overflow may occur so that the Hive service is unavailable.

Possible Causes

The non-heap memory of the Hive instance on the node is overused or the non-heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check non-heap memory usage.

Step 1 On the FusionInsight Manager portal, click **O&M > Alarm > Alarms** and select the alarm whose **Alarm ID** is **16008**. Then check the role name in **Location** and confirm the IP adress of the instance.

- If the role for which the alarm is generated is HiveServer, go to **Step 2**.
- If the role for which the alarm is generated is MetaStore, go to **Step 3**.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the HiveServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **HiveServer Memory Usage Statistics** and click **OK**, check whether the used non-heap memory of the HiveServer service reaches the threshold(default value: 95%) of the maximum non-heap memory specified for HiveServer.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-35 HiveServer Memory Usage Statistics

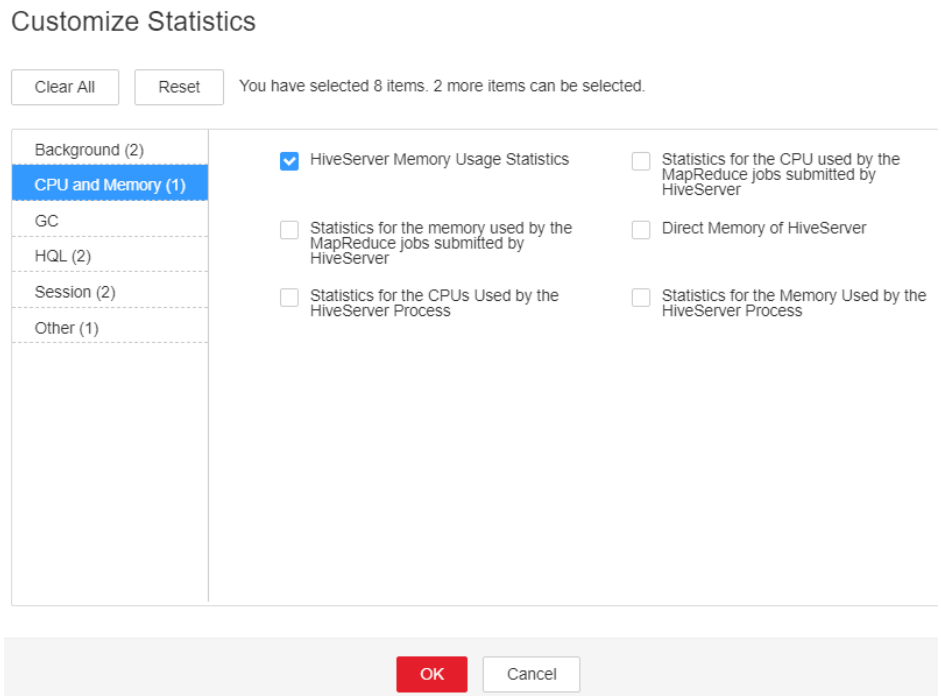

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Instance** and click the MetaStore for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory**, and select **MetaStore Memory Usage Statistics** and click **OK**, check whether the used non-heap memory of the MetaStore service reaches the threshold(default value: 95%) of the maximum non-heap memory specified for MetaStore.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-36 MetaStore Memory Usage Statistics

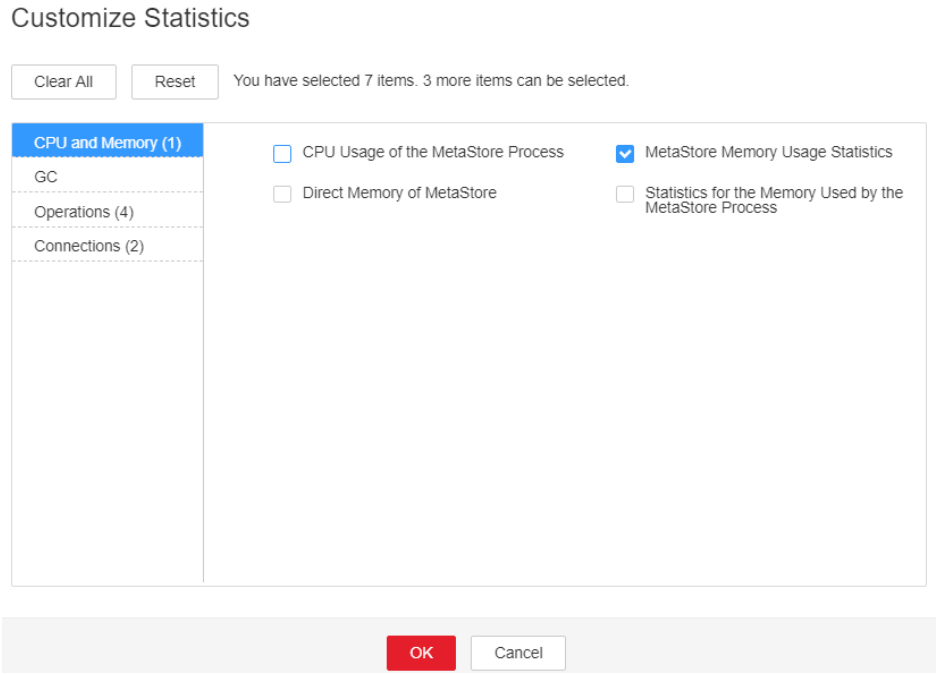

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hive** > **Configurations** > **All Configurations**. Choose **HiveServer/MetaStore** > **JVM**. Adjust the value of **-XX:MaxMetaspaceSize** in **HIVE_GC_OPTS/METASTORE_GC_OPTS** as the following rules. Click **Save**.

NOTE

Suggestions for GC parameter settings for the HiveServer:

- It is recommended that you set the value of **-XX:MaxMetaspaceSize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 2 GB, **-XX:**

  MaxMetaspaceSize is set to 256 MB. If **-Xmx** is set to 4 GB, **-XX:MaxMetaspaceSize** is set to 512 MB.

Suggestions for GC parameter settings for the MetaServer:

- It is recommended that you set the value of **-XX:MaxMetaspaceSize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 2 GB, **-XX:**

  MaxMetaspaceSize is set to 256 MB. If **-Xmx** is set to 4 GB, **-XX:MaxMetaspaceSize** is set to 512 MB

Step 5 Click **More** > **Restart Service** to restart the service.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **Hive** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.131 ALM-16009 Map Number Exceeds the Threshold

Description

The system checks the number of HQL maps in every 30 seconds. This alarm is generated if the number exceeds the threshold. By default, **Trigger Count** is set to **3**, and the threshold is 5000.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16009	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the number of HQL maps executed on Hive is excessively large, the HQL execution speed is slow, and a large number of resources are occupied.

Possible Causes

The HQL statements are not the optimal.

Procedure

Check the number of HQL maps.

Step 1 On FusionInsight Manager portal, choose **Cluster** >*Name of the desired cluster*> **Services** > **Hive** > **Resource**. Check the HQL statements with the excessively large number (5000 or more) of maps in **HQL Map Count**.

Step 2 Locate the corresponding HQL statements, optimize them and execute them again.

Step 3 Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Collect fault information.

Step 4 On the FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 5 Select **Hive** in the required cluster from the **Service**.

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.132 ALM-16045 Hive Data Warehouse Is Deleted

Description

The system checks the Hive data warehouse in every 60 seconds.This alarm is generated when the Hive data warehouse is deleted.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16045	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The default Hive data warehouse is deleted. As a result, creating databases or tables in the default data warehouse fails, and services are affected.

Possible Causes

Hive periodically checks the status of the default data warehouse and finds that the default data warehouse is deleted.

Procedure

Check the default Hive data warehouse.

Step 1 Log in to the node where the client is located as user **root**.

Step 2 Run the following command to check whether the **warehouse** directory exists in **hdfs://hacluster/user/***<username>***/.Trash/Current/**.

hdfs dfs -ls hdfs://hacluster/user/*<username>***/.Trash/Current/**

For example, if **user/hive/warehouse** exists:

```
host01:/opt/client # hdfs dfs -ls hdfs://hacluster/user/test/.Trash/Current/
Found 1 items
drwx------   - test hadoop          0 2019-06-17 19:53 hdfs://hacluster/user/test/.Trash/Current/user
```

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 By default, there is an automatic recovery mechanism for the data warehouse. You can wait for 5 ~10s to check whether the default data warehouse is restored. If

the data warehouse is not recovered, manually run the following command to restore the data warehouse.

hdfs dfs -mv hdfs://hacluster/user/*<username>***/.Trash/Current/user/hive/warehouse /user/hive/warehouse**

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 Collect related information in the **.Trash/Current/** directory on the client background.

Step 6 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.133 ALM-16046 Hive Data Warehouse Permission Is Modified

Description

The system checks the Hive data warehouse permission in every 60 seconds. This alarm is generated if the permission is modified.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16046	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the permission on the Hive default data warehouse is modified, the permission for users or user groups to create databases or tables in the default data warehouse is changed.

Possible Causes

Hive periodically checks the status of the default data warehouse and finds that default data warehouse permission is changed.

Procedure

Check the Hive default data warehouse permission.

Step 1 Log in to the node where the client is located as user **root**.

Step 2 Run the following command to go to the HDFS client installation directory:

cd *Client installation directory*

source bigdata_env

kinit *User who has the supergroup permission* (Skip this step for a common cluster.)

Step 3 Run the following command to restore the default data warehouse permission:

- Security mode: **hdfs dfs -chmod 770 hdfs://hacluster/user/hive/warehouse**
- Non-security mode: **hdfs dfs -chmod 777 hdfs://hacluster/user/hive/warehouse**

Step 4 Check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 5**.

Collect fault information.

Step 5 Collect related information in the **hdfs://hacluster/user/hive/warehouse** directory on the client background.

Step 6 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.134 ALM-16047 HiveServer Has Been Deregistered from ZooKeeper

Alarm Description

The system checks the Hive service every 60 seconds. This alarm is generated when Hive registration information on ZooKeeper is lost or Hive cannot connect to ZooKeeper.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
16047	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

When a Hive client sets up a new connection, it cannot select the HiveServer node that has been deregistered from ZooKeeper. If all HiveServer nodes have been deregistered from ZooKeeper, the HiveServer service will be unavailable.

Possible Causes

- The ZooKeeper instance is abnormal.

- Some Hive configurations are incorrect.

Handling Procedure

Check the ZooKeeper service status.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** and check whether **ALM-12007 Process Fault** exists in the alarm list.

- If yes, go to **Step 2**.
- If no, go to **Step 5**.

Step 2 In **Location** of **ALM-12007 Process Fault**, check whether the service name is **ZooKeeper**.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by following steps provided in **ALM-12007 Process Fault**.

Step 4 In the alarm list, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the Hive configurations are correctly modified.

Step 5 On FusionInsight Manager, choose **Audit**. On the **Audit** page, click **Advanced Search**, click ··· on the right of **Operation Type**, select **Save configuration**, click **OK**, and click **Search**.

Step 6 In the search result, check the historical configurations of Hive- and ZooKeeper-related services in the **Service** column. **Table 10-4** lists some configurations that may affect the connection between Hive and ZooKeeper.

Table 10-4 Configurations related to connection between Hive and ZooKeeper

Service	Parameter	Description
Hive	HIVE_GC_OPTS	HiveServer memory configuration. If the configuration is abnormal, HiveServer may restart repeatedly. In this case, you need to check the health status of the instance processes.
	hive.zookeeper.quorum	IP address of the node accommodating ZooKeeper that is connected to Hive.
	hive.zookeeper.client.port	Port of the ZooKeeper client connected to Hive.
	hive.zookeeper.session.timeout	Timeout interval of the session set up between Hive and ZooKeeper.

Service	Parameter	Description
	hive.zookeeper.connection.timeout	Timeout interval for Hive to connect to ZooKeeper.
	hive.zookeeper.connection.max.retries	Maximum number of retries for Hive to connect to ZooKeeper.
ZooKeeper	clientPort	Port number of the ZooKeeper client.
	ssl.enabled	Whether to enable SSL connections of ZooKeeper.

Restart related instances.

Step 7 Log in to FusionInsight Manager. Choose **O&M** > **Alarm** > **Alarms**, click the drop-down list in the row that contains the alarm, and view the role and the IP address of the node for which the alarm is generated in **Location**.

Step 8 Choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Hive** > **Instance**. On the page that is displayed, select the instance at the IP address for which the alarm is generated, click **More**, and select **Restart Instance**.

Step 9 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, and select **Hive** for the target cluster.

Step 12 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.135 ALM-16048 Tez or Spark Library Path Does Not Exist

Description

The system checks the Tez and Spark library paths every 180 seconds. This alarm is generated when the Tez or Spark library path does not exist.

Attribute

Alarm ID	Alarm Severity	Auto Clear
16048	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The Hive on Tez and Hive on Spark functions are affected.

Possible Causes

The Tez or Spark library path is deleted from the HDFS.

Procedure

Check the default Hive data warehouse.

Step 1 Log in to the node where the client is located as user **root**.

Step 2 Run the following command to check whether the **tezlib** or **sparklib** directory exists in the **hdfs://hacluster/user/{User name}/.Trash/Current/** director:

hdfs dfs -ls hdfs://hacluster/user/*<username>***/.Trash/Current/**

For example, the following information shows that **/user/hive/tezlib/8.1.0.1/** and **/user/hive/sparklib/8.1.0.1/** exist.

```
host01:/opt/client # hdfs dfs -ls hdfs://hacluster/user/test/.Trash/Current/
Found 1 items
drwx------   - test hadoop          0 2019-06-17 19:53 hdfs://hacluster/user/test/.Trash/Current/user
```

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Run the following command to restore **tezlib** and **sparklib**.

hdfs dfs -mv hdfs://hacluster/user/*<username>*/.Trash/Current/user/hive/tezlib/8.1.0.1/tez.tar.gz /user/hive/tezlib/8.1.0.1/tez.tar.gz

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 Collect related information in the **.Trash/Current/** directory on the client background.

Step 6 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.136 ALM-17003 Oozie Service Unavailable

Description

The system checks the Oozie service status in every 5 seconds. This alarm is generated when Oozie or a component on which Oozie depends cannot provide services properly.

This alarm is automatically cleared when the Oozie service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
17003	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Oozie cannot be used to submit jobs.

Possible Causes

- The DBService service is abnormal or the data of Oozie stored in DBService is damaged.
- The HDFS service is abnormal or the data of Oozie stored in HDFS is damaged.
- The Yarn service is abnormal.
- The Nodeagent process is abnormal.

Procedure

Query the Oozie service health status code.

Step 1 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* >**Services** > **Oozie**. Click **oozie** (any one is OK) on the **oozie WebUI**. to go to the Oozie WebUI.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 2 Add **/servicehealth** to the URL in the address box of the browser and access again. The value of **statusCode** is the current Oozie service health status code.

For example, visit **https://10.10.0.117:20026/Oozie/oozie/130/oozie/servicehealth**. The result is as follows:

```
{"beans":[{"name":"serviceStatus","statusCode":0}]}
```

If the health status code cannot be displayed or the browser does not respond, the service may be unavailable due to Oozie process fault. See **Step 13** to rectify the fault.

Step 3 Perform the operations based on the error code. For details, see **Table 10-5**.

Table 10-5 Oozie service health status code

Status Code	Description	Error Cause	Solution
0	The service is running properly.	None	None
18002	The DBService service is abnormal.	Oozie fails to connect to DBService or the data stored in DBService is damaged.	See **Step 4**.
18003	The HDFS service is abnormal.	Oozie fails to connect to HDFS or the data stored in HDFS is damaged.	See **Step 7**.
18005	The MapReduce service is abnormal.	The Yarn service is abnormal.	See **Step 11**.

Check the DBService service.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the DBService service is running properly.

- If yes, go to **Step 6**.
- If no, go to **Step 5**.

Step 5 Resolve the problem of DBService based on the alarm help and check whether the Oozie alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Step 6 Log in to the Oozie database to check whether the data is complete.

1. Log in to the active DBService node as user **root**.

   On the FusionInsight Manager page, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService > Instance** to view the IP address of the active DBservice node.

2. Run the following command to log in to the Oozie database:

   su - omm

   source ${BIGDATA_HOME}/FusionInsight_BASE_8.1.0.1/install/FusionInsight-dbservice-2.7.0/.dbservice_profile

   gsql -U *Username* **-W** *Oozie database password* **-p 20051 -d** *Database name*

3. After the login is successful, enter **\d** to check whether there are 15 data tables.

The Oozie service has 15 data tables by default. If these data tables are deleted or the table structure is modified, the Oozie service may be unavailable. Contact the O&M personnel to back up the data and perform restoration.

Check the HDFS service.

Step 7 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the HDFS service is running properly.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

Step 8 Resolve the problem of HDFS based on the alarm help and check whether the Oozie alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Step 9 Log in to HDFS to check whether the Oozie file directory structure is complete.

1. Download and install an HDFS client..
2. Log in to the client node as user **root** and run the following commands to check whether **/user/oozie/share** exists.

   If the cluster uses the security mode, perform security authentication.

   kinit admin

   hdfs dfs -ls /user/oozie/share

- If yes, go to **Step 18**.
- If no, go to **Step 10**.

Step 10 In the Oozie client installation directory, manually upload the share directory to **/user/oozie** in HDFS, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Check the Yarn and MapReduce service.

Step 11 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**, and check whether the Yarn and MapReduce services are running properly.

- If yes, go to **Step 18**.
- If no, go to **Step 12**.

Step 12 Resolve the problem of Yarn and MapReduce based on the alarm help and check whether the Oozie alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Check the Oozie process.

Step 13 Log in to each node of Oozie as user **root**.

Step 14 Run the **ps -ef | grep oozie** command to check whether the Oozie process exists.

- If yes, go to **Step 15**.
- If no, go to **Step 18**.

Step 15 Collect fault information in **prestartDetail.log**, **oozie.log**, and **catalina.out** in the Oozie log directory **/var/log/Bigdata/oozie**. If the alarm is not caused by manual misoperation, go to **Step 16**.

Check the Nodeagent process.

Step 16 Log in to each node of Oozie as user **root**. Run the **ps -ef | grep nodeagent** command to check whether the Nodeagent process exists.

- If yes, go to **Step 17**.
- If no, go to **Step 18**.

Step 17 Run the **kill -9** *The process ID of nodeagent* command, wait 10 minutes, and check whether alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Step 18 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.137 ALM-17004 Oozie Heap Memory Usage Exceeds the Threshold

Description

The system checks the heap memory usage of the Oozie service every 60 seconds. The alarm is generated when the heap memory usage of a Metadata instance exceeds the threshold (95% of the maximum memory). The alarm is cleared when the heap memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
17004	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The heap memory overflow may cause a service breakdown.

Possible Causes

The heap memory of the Oozie instance is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **Oozie Heap Memory Usage Exceeds the Threshold** > **Location**. Check the IP address of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* > **Services** > **Oozie** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize** > **Memory** > **Oozie Heap Memory Resource Percentage**. Click **OK**.

Figure 10-37 Oozie Heap Memory Resource Percentage

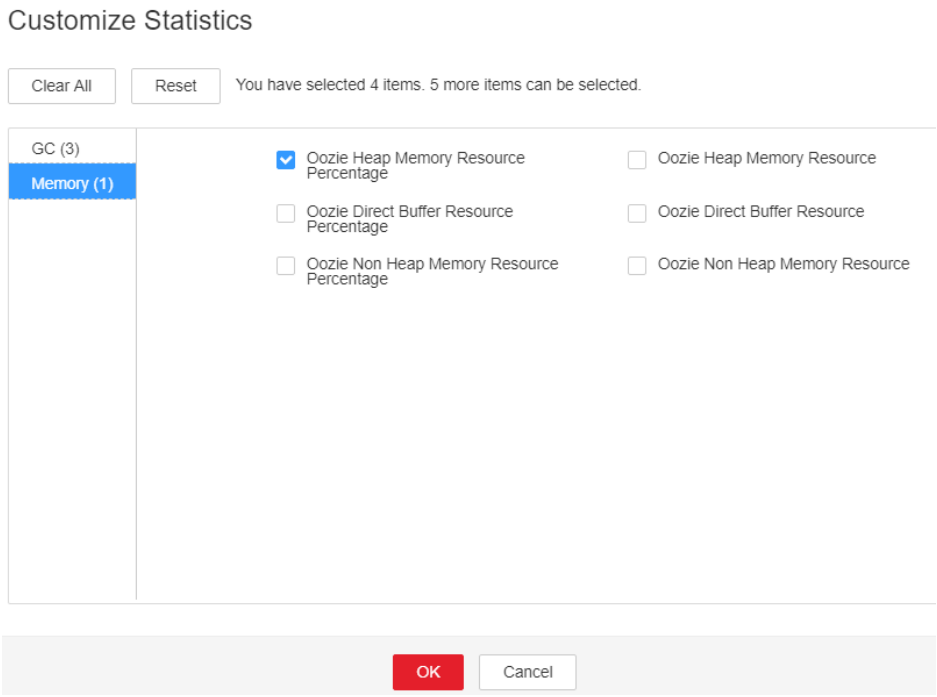

Step 3 Check whether the used heap memory of Oozie reaches the threshold (the default value is 95% of the maximum heap memory) specified for Oozie.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie** > **Configurations** > **All Configurations**. Set Search **GC_OPTS** in the search box. Increase the value of **-Xmx** as required, and click **Save** > **OK**.

NOTE

Suggestions on GC parameter settings for Oozie:

You are advised to set **-Xms** and **-Xmx** to the same value to prevent adverse impact on performance when JVM dynamically adjusts the heap memory size.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Oozie** in the required cluster from the **Service**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.138 ALM-17005 Oozie Non Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the non heap memory usage of Oozie every 30 seconds. This alarm is reported if the non heap memory usage of Oozie exceeds the threshold (80%). This alarm is cleared if the non heap memory usage is lower than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
17005	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Non-heap memory overflow may cause service breakdown.

Possible Causes

The non-heap memory of the Oozie instance is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **Oozie Non Heap Memory Usage Exceeds the Threshold**. On the displayed page, check the location information of the alarm. Check the name of the instance host for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the target cluster* > **Services** > **Oozie** and click the **Instance** tab. On the displayed page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Memory** and select **Oozie Non Heap Memory Resource Percentage**. Click **OK**.

Figure 10-38 Oozie non-heap memory usage

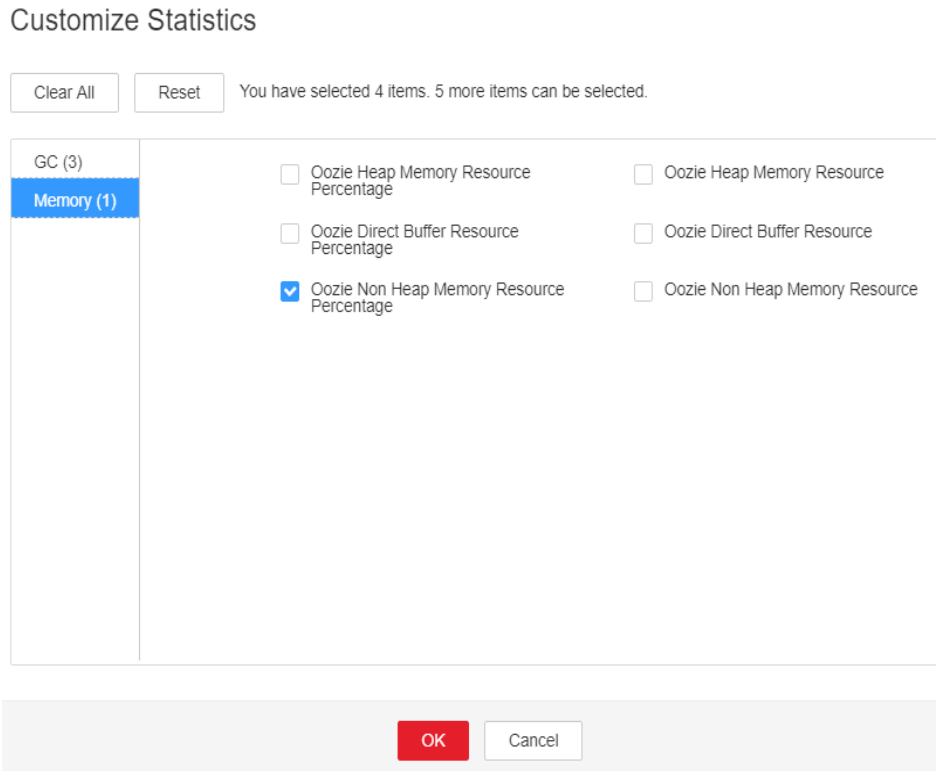

Step 3 Check whether the non-heap memory used by Oozie reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the target cluster* > **Services** > **Oozie** and click the **Configurations** and then **All Configurations**. On

the displayed page, search for the **GC_OPTS** parameter in the search box and check whether it contains **-XX: MaxMetaspaceSize**. If yes, increase the value of **-XX: MaxMetaspaceSize** based on the site requirements. If no, manually add **-XX: MaxMetaspaceSize** and set its value to 1/8 of the value of **-Xmx**. Click **Save**, and then click **OK**

NOTE

JDK1.8 does not support the **MaxPermSize** parameter.

Suggestions on GC parameter settings for Oozie:

Set the value of **-XX:MaxMetaspaceSize** to 1/8 of the value of **-Xmx**. For example, if **-Xmx** is set to 2 GB, **-XX:MaxMetaspaceSize** is set to 256 MB. If **-Xmx** is set to 4 GB, **-XX:MaxMetaspaceSize** is set to 512 MB.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Oozie** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.139 ALM-17006 Oozie Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the Oozie service every 30 seconds. The alarm is generated when the direct memory usage of an Oozie instance exceeds the threshold (80% of the maximum memory). The alarm is cleared when the direct memory usage of Oozie is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
17006	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The direct memory overflow may cause a service breakdown.

Possible Causes

The direct memory of the Oozie instance is overused or the direct memory is inappropriately allocated.

Procedure

Check direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Oozie Direct Memory Usage Exceeds the Threshold** > **Location**. Check the IP address of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* > **Services** > **Oozie** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize** > **Memory** > **Oozie Direct Buffer Resource Percentage**. Click **OK**.

Figure 10-39 Oozie Direct Buffer Resource Percentage

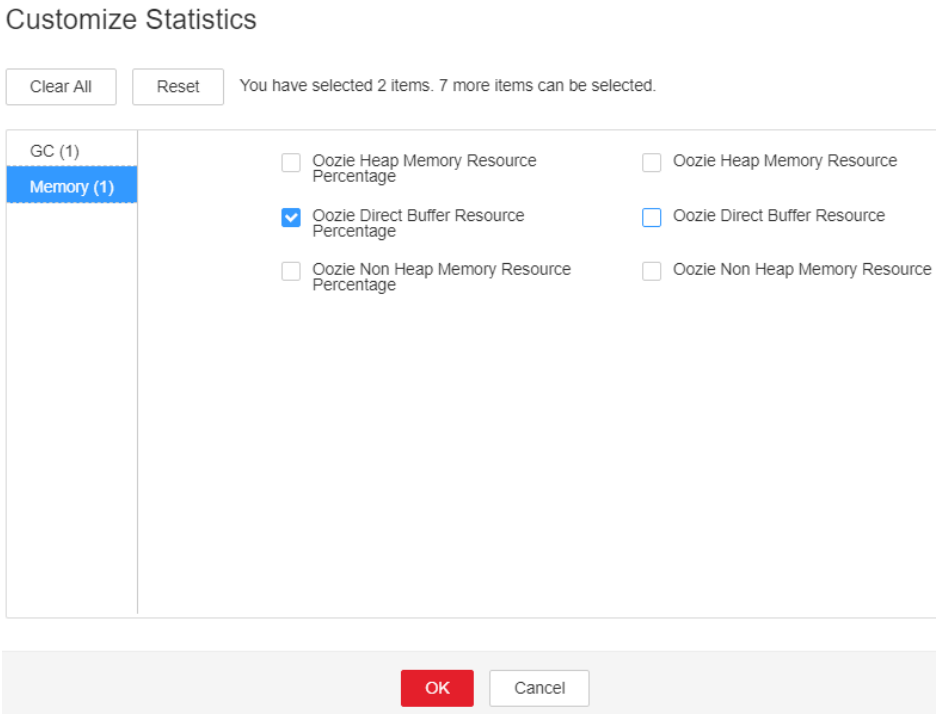

Step 3 Check whether the used direct memory of Oozie reaches the threshold (the default value is 80% of the maximum direct memory) specified for Oozie.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie** > **Configurations**. Click **All Configurations**. Search **GC_OPTS** in the search box. Increase the value of **-XX:MaxDirectMemorySize** as required, and click **Save**. Click **OK**.

NOTE

Suggestions on GC parameter settings for Oozie:

You are advised to set the value of **-XX:MaxDirectMemorySize** to 1/4 of the value of **-Xmx**. For example, if **-Xmx** is set to 4 GB, **-XX:MaxDirectMemorySize** is set to 1024 MB. If **-Xmx** is set to 2 GB, **-XX:MaxDirectMemorySize** is set to 512 MB. It is recommended that the value of **-XX:MaxDirectMemorySize** be greater than or equal to 512 MB.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Oozie** in the required cluster from the **Service** drop-down list.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.140 ALM-17007 Garbage Collection (GC) Time of the Oozie Process Exceeds the Threshold

Description

The system checks GC time of the Oozie process every 60 seconds. The alarm is generated when GC time of the Oozie process exceeds the threshold (default value: **12 seconds**). The alarm is cleared when GC time is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
17007	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Oozie responds slowly when it is used to submit tasks.

Possible Causes

The heap memory of the Oozie instance is overused or the heap memory is inappropriately allocated.

Procedure

Check GC time.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Garbage Collection (GC) Time of the Oozie Process Exceeds the Threshold** > **Location**. Check the IP address of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize** > **GC** > **Garbage Collection (GC) Time of Oozie**. Click **OK**.

Figure 10-40 Garbage Collection (GC) Time of Oozie

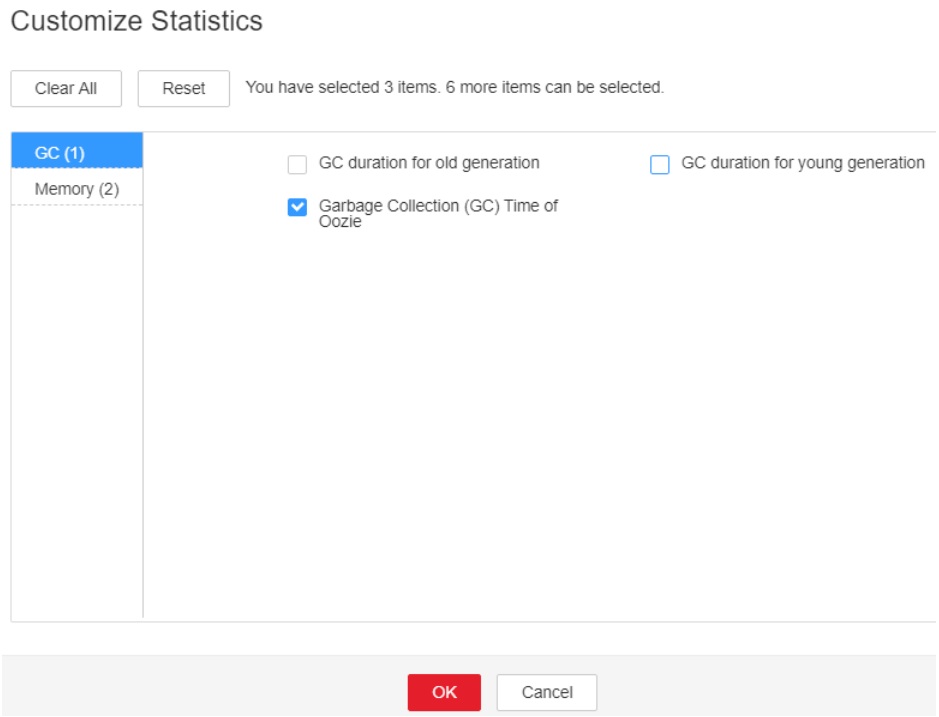

Step 3 Check whether GC time of the Oozie process every second exceeds the threshold (default value: **12 seconds**).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Oozie** > **Configurations**. Click **All Configurations**. Search

GC_OPTS in the search box. Increase the value of **-Xmx** as required, and click **Save**. Click **OK**.

NOTE

Suggestions on GC parameter settings for Oozie:

You are advised to set **-Xms** and **-Xmx** to the same value to prevent adverse impact on performance when JVM dynamically adjusts the heap memory size.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Oozie** in the required cluster from the **Service** drop-down list.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.141 ALM-17008 Abnormal Connection Between Oozie and ZooKeeper

Alarm Description

In HA mode, Oozie depends on ZooKeeper. This alarm is generated when the connection between Oozie and ZooKeeper is abnormal for three consecutive times.

This alarm is cleared when the connection between Oozie and ZooKeeper becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
17008	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

In HA mode, the Oozie service will restart if this alarm is reported.

Possible Causes

- The ZooKeeper service is abnormal.
- Oozie fails to connect to ZooKeeper.

Handling Procedure

Check the ZooKeeper service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of ZooKeeper is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 2**.

Step 2 In the alarm list, check whether **ALM-13000 ZooKeeper Service Unavailable** is reported.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by performing the operations provided for **ALM-13000 ZooKeeper Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and ZooKeeper** is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the connectivity between Oozie and ZooKeeper.

Step 5 Log in to FusionInsight Manager, choose **O&M** > **Log** > **Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement] [ZooKeeper]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and ZooKeeper** is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.142 ALM-17009 Abnormal Connection Between Oozie and DBService

Alarm Description

Oozie depends on DBService. After a task is submitted, the system checks DBService connectivity. This alarm is generated when the service fails the check for 10 consecutive times.

This alarm is cleared when the connection between Oozie and DBService becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
17009	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The Oozie service cannot be used to submit jobs.

Possible Causes

- The DBService service is abnormal.
- Oozie fails to connect to DBService.

Handling Procedure

Check the DBService status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of DBService is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 2**.

Step 2 In the alarm list, check whether **ALM-27001 DBService Service Unavailable** is reported.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by performing the operations provided for **ALM-27001 DBService Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and DBService** is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the connectivity between Oozie and DBService.

Step 5 Log in to FusionInsight Manager, choose **O&M** > **Log** > **Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement][DB Service]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and DBService** is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.143 ALM-17010 Abnormal Connection Between Oozie and HDFS

Alarm Description

Oozie depends on HDFS. After a task is submitted, the system checks HDFS connectivity. This alarm is generated when the service fails the check for 3 consecutive times.

This alarm is cleared when the connection between Oozie and HDFS becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
17010	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The Oozie service is abnormal and Oozie-related jobs cannot be submitted.

Possible Causes

The HDFS service restarts, there is a fault, or the network connectivity is abnormal.

Handling Procedure

Check the HDFS service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of HDFS is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 2**.

Step 2 In the alarm list, check whether the "ALM-14000 HDFS Service Unavailable" alarm is generated.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by performing the operations provided for **ALM-14000 HDFS Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and HDFS** is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the connectivity between Oozie and HDFS.

Step 5 Log in to FusionInsight Manager, choose **O&M** > **Log** > **Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement][HDFS]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and HDFS** is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.144 ALM-17011 Abnormal Connection Between Oozie and Yarn

Alarm Description

Oozie depends on Yarn. After a task is submitted, the system checks Yarn connectivity. This alarm is generated when the service fails the check for 5 consecutive times.

This alarm is cleared when the connection between Oozie and Yarn becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
17011	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The Oozie service is abnormal and tasks cannot be submitted to Yarn.

Possible Causes

- The Yarn service is abnormal.
- The connection between Oozie and Yarn is abnormal.

Handling Procedure

Check the YARN service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of Yarn is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 2**.

Step 2 In the alarm list, check whether **ALM-18000 YARN Service Unavailable** is generated.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by performing the operations provided for **ALM-18000 Yarn Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **Abnormal Connection Between Oozie and Yarn** is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the connectivity between Oozie and Yarn.

Step 5 Log in to FusionInsight Manager, choose **O&M** > **Log** > **Online Search**, select the Oozie service, and search for the keyword **[Oozie Alarm Enhancement][Yarn]** in the log. View the cause in the log, and rectify the fault. In the alarm list, check whether the alarm **Abnormal Connection Between Oozie and Yarn** is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Select **Oozie** for **Service** and click **OK**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.145 ALM-18000 Yarn Service Unavailable

Description

This alarm is generated when the Yarn service is unavailable. The alarm module checks the Yarn service status every 60 seconds.

The alarm is cleared when the Yarn service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18000	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceNam	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The cluster cannot provide Yarn services. Users cannot run new applications. Submitted applications cannot be run.

Possible Causes

- The ZooKeeper service is abnormal.
- The HDFS service is abnormal.
- There is no active ResourceManager instance in the Yarn cluster.
- All the NodeManagers in the Yarn cluster are abnormal.

Procedure

Check ZooKeeper service status.

Step 1 On the FusionInsight Manager, check whether the alarm list contains **ALM-13000 ZooKeeper Service Unavailable**.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 Rectify the fault by following the steps provided in **ALM-13000 ZooKeeper Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Check the HDFS service status.

Step 3 On the FusionInsight Manager, check whether the alarm list contains the HDFS alarms.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Choose **O&M > Alarm > Alarms**, handle HDFS alarms based on the alarm help, and check whether the Yarn alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the ResourceManager status in the Yarn cluster.

Step 5 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn**.

Step 6 In **Dashboard**, check whether there is an active ResourceManager instance in the Yarn cluster.

- If yes, go to **Step 7**.
- If no, go to **Step 10**.

Check the NodeManager node status in the Yarn cluster.

Step 7 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn** > **Instance**.

Step 8 Query NodeManager **Running Status**, and check whether there are unhealthy nodes.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 Rectify the fault by following the steps provided in **ALM-18002 NodeManager Heartbeat Lost** or **ALM-18003 NodeManager Unhealthy**. After the fault is rectified, check whether the Yarn alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On the FusionInsight Manager portal of the active cluster, choose **O&M** > **Log > Download**.

Step 11 Select **Yarn** in the required cluster from the **Service**.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.146 ALM-18002 NodeManager Heartbeat Lost

Description

The system checks the number of lost NodeManager nodes every 30 seconds, and compares the number with the threshold. The Number of Lost Nodes indicator has

a default threshold. The alarm is generated when the value of Number of Lost Nodes exceeds the threshold.

To change the threshold, on FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn**. On the displayed page, choose **Configurations** > **All Configurations**, and change the value of **yarn.nodemanager.lost.alarm.threshold**. You do not need to restart Yarn to make the change take effect.

The default threshold is 0. The alarm is generated when the number of lost nodes exceeds the threshold, and is cleared when the number of lost nodes is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18002	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Lost Host	Specifies the list of hosts with lost nodes.

Impact on the System

- The lost NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

Possible Causes

- NodeManager is forcibly deleted without decommission.
- All the NodeManager instances are stopped or the NodeManager process is faulty.
- The host where the NodeManager node resides is faulty.
- The network between the NodeManager and ResourceManager is disconnected or busy.

Procedure

Check the NodeManager status.

Step 1 On the FusionInsight Manager, and choose **O&M** > **Alarm > Alarms**. Click before the alarm and obtain lost nodes in **Additional Information**.

Step 2 Check whether the lost nodes are hosts that have been manually deleted without decommission.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 After the setting, Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn**. On the displayed page, choose **Configurations** > **All Configurations**. Search for **yarn.nodemanager.lost.alarm.threshold** and change its value to the number of hosts that are not out of service and proactively deleted. After the setting, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Step 4 Manually clear the alarm. Note that decommission must be performed before deleting hosts.

Step 5 On the FusionInsight Manager portal, choose **Cluster > Hosts**, and check whether the nodes obtained in **Step 1** are healthy.

- If yes, go to **Step 7**.
- If no, go to **Step 6**.

Step 6 Rectify the node fault based on **ALM-12006 Node Fault** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check the process status.

Step 7 On the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* **>** **Services** > **Yarn** > **Instance**, and check whether there are NodeManager instances whose status is not **Good**.

- If yes, go to **Step 10**.
- If no, go to **Step 8**.

Step 8 Check whether the NodeManager instance is deleted.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

Step 9 Restart the active and standby ResourceManager instances, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Check the instance status.

Step 10 Select NodeManager instances which running state is not **Normal** and restart them. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Check the network status.

Step 11 Log in to the management node, **ping** the IP address of the lost NodeManager node to check whether the network is disconnected or busy.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 Rectify the network, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collect fault information.

Step 13 On the FusionInsight Manager in the active cluster, choose **O&M** > **Log** > **Download**.

Step 14 Select **Yarn** in the required cluster from the **Service**.

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.147 ALM-18003 NodeManager Unhealthy

Description

The system checks the number of unhealthy NodeManager nodes every 30 seconds, and compares the number with the threshold. The Unhealthy Nodes indicator has a default threshold. This alarm is generated when the value of the Unhealthy Nodes indicator exceeds the threshold.

To change the threshold, on FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn**. On the displayed page, choose **Configurations** > **All Configurations**, and change the value of **yarn.nodemanager.unhealthy.alarm.threshold**. You do not need to restart Yarn to make the change take effect.

The default threshold is 0. The alarm is generated when the number of unhealthy nodes exceeds the threshold, and is cleared when the number of unhealthy nodes is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18003	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Unhealthy Host	Specifies the list of hosts with unhealthy nodes.

Impact on the System

- The faulty NodeManager node cannot provide the Yarn service.
- The number of containers decreases, so the cluster performance deteriorates.

Possible Causes

- The hard disk space of the host where the NodeManager node resides is insufficient.
- User **omm** does not have the permission to access a local directory on the NodeManager node.

Procedure

Check the hard disk space of the host.

Step 1 On the FusionInsight Manager, and choose **O&M** > **Alarm > Alarms**. Click before the alarm and obtain unhealthy nodes in **Additional Information**.

Step 2 Choose **Cluster** > *Name of the desired cluster* **> Services > Yarn** > **Instance**, select the NodeManager instance corresponding to the host, choose **Instance**

Configurations > All Configurations and view disks corresponding to **yarn.nodemanager.local-dirs** and **yarn.nodemanager.log-dirs**.

Step 3 Choose **O&M > Alarm > Alarms**. In the alarm list, check whether the related disk has the alarm **ALM-12017 Insufficient Disk Capacity**.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Rectify the disk fault based on **ALM-12017 Insufficient Disk Capacit** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Step 5 Choose **Hosts** > *Name of the desired host* . On the **Dashboard** page, check the disk usage of the corresponding partition. Check whether the percentage of the used space of the mounted disk exceeds the value of **yarn.nodemanager.disk-health-checker.max-disk-utilization-per-disk-percentage**

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Reduce the disk usage to less than the value of **yarn.nodemanager.disk-health-checker.max-disk-utilization-per-disk-percentage**, wait for 10 to 20 minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check the access permission of the local directory on each NodeManager node.

Step 7 Obtain the NodeManager directory viewed in **Step 2**, log in to each NodeManager node as user **root**, and go to the obtained directory.

Step 8 Run the **ll** command to check whether the permission of the **localdir** and **containerlogs** folders is **755** and whether **User:Group** is **omm:ficommon**.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Run the following command to set the permission to **755** and **User:Group** to **omm:ficommon**:

chmod 755 *<folder_name>*

chown omm:ficommon *<folder_name>*

Step 10 Wait for 10 to 20 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On the FusionInsight Manager in the active cluster, choose **O&M** > **Log > Download**.

Step 12 Select **Yarn** in the required cluster from the **Service**.

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.148 ALM-18008 Heap Memory Usage of ResourceManager Exceeds the Threshold

Description

The system checks the heap memory usage of Yarn ResourceManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Yarn ResourceManager exceeds the threshold (95% of the maximum memory by default).

Users can choose **O&M > Alarm > Thresholds** > *Name of the desired cluster* > **Yarn** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the heap memory usage of Yarn ResourceManager is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the heap memory usage of Yarn ResourceManager is less than or equal to 95% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18008	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the heap memory usage of Yarn ResourceManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

The heap memory of the Yarn ResourceManager instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **Heap Memory Usage of Yarn ResourceManager Exceeds the Threshold** > **Location**. Check the HostName of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* **> Services** > **Yarn** > **Instance** > **ResourceManager** (Indicates the host name of the instance for which the alarm is generated)**.** Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **ResourceManager > Percentage of Used Memory of the ResourceManager**. Check the heap memory usage.

Figure 10-41 Percentage of Used Memory of the ResourceManager

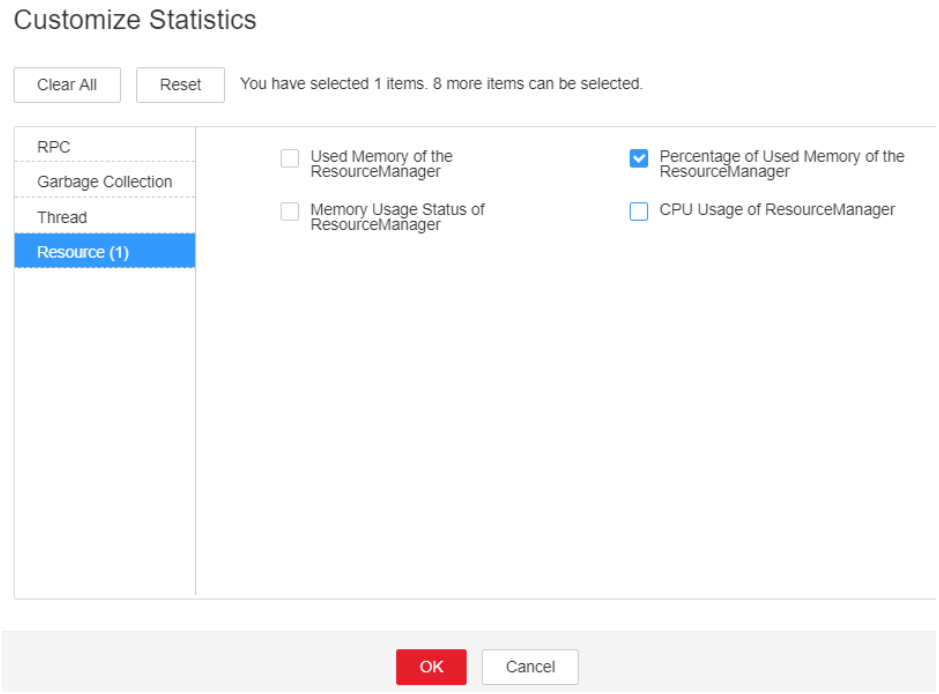

Step 3 Check whether the used heap memory of ResourceManager reaches 95% of the maximum heap memory specified for ResourceManager.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* **> Services** > **Yarn** > **Configurations** > **All Configurations** > **ResourceManager** > **System**. Increase the value of **GC_OPTS** parameter as required, click **Save**. Restart the role instance.

NOTE

The mapping between the number of NodeManager instances in a cluster and the memory size of ResourceManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms10G -Xmx10G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 1000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms20G -Xmx20G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 2000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms40G -Xmx40G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 3000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms60G -Xmx60G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 4000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms80G -Xmx80G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 5000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms100G -Xmx100G -XX:NewSize=3G -XX:MaxNewSize=6G

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.149 ALM-18009 Heap Memory Usage of JobHistoryServer Exceeds the Threshold

Description

The system checks the heap memory usage of Mapreduce JobHistoryServer every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Mapreduce JobHistoryServer exceeds the threshold (95% of the maximum memory by default).

Users can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Mapreduce** to change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the heap memory usage of MapReduce JobHistoryServer is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the heap memory usage of MapReduce JobHistoryServer is less than or equal to 95% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18009	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the heap memory usage of Mapreduce JobHistoryServer is overhigh, the performance of Mapreduce log archiving is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

The heap memory of the Mapreduce JobHistoryServer instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18009 Heap Memory Usage of MapReduce JobHistoryServer Exceeds the Threshold** > **Location**. Check the HostName of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* **> Services** > **Mapreduce** > **Instance** > **JobHistoryServer.** Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **JobHistoryServer heap memory usage statistics**. JobHistoryServer indicates the corresponding HostName of the instance for which the alarm is generated. Check the heap memory usage.

Step 3 Check whether the used heap memory of JobHistoryServer reaches 95% of the maximum heap memory specified for JobHistoryServer.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* **> Services** > **Mapreduce** > **Configurations** > **All Configurations** > **JobHistoryServer** > **System**. Increase the value of **GC_OPTS** parameter as required, click **Save**. Click **OK** and restart the role instance.

NOTE

The mapping between the number of historical tasks (10000) and the memory of JobHistoryServer is as follows:

-Xms30G -Xmx30G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent

- Mapreduce

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.150 ALM-18010 ResourceManager GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the ResourceManager process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18010	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

A long GC duration of the ResourceManager process may interrupt the services.

Possible Causes

The heap memory of the ResourceManager instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18010 ResourceManager GC Time Exceeds the Threshold** > **Location** to check the IP address of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* **> Services** > **Yarn** > **Instance** > **ResourceManager (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Garbage Collection (GC) Time of ResourceManager** to check the GC duration statistics of the Broker process collected every minute.

Figure 10-42 Garbage Collection (GC) Time of ResourceManager

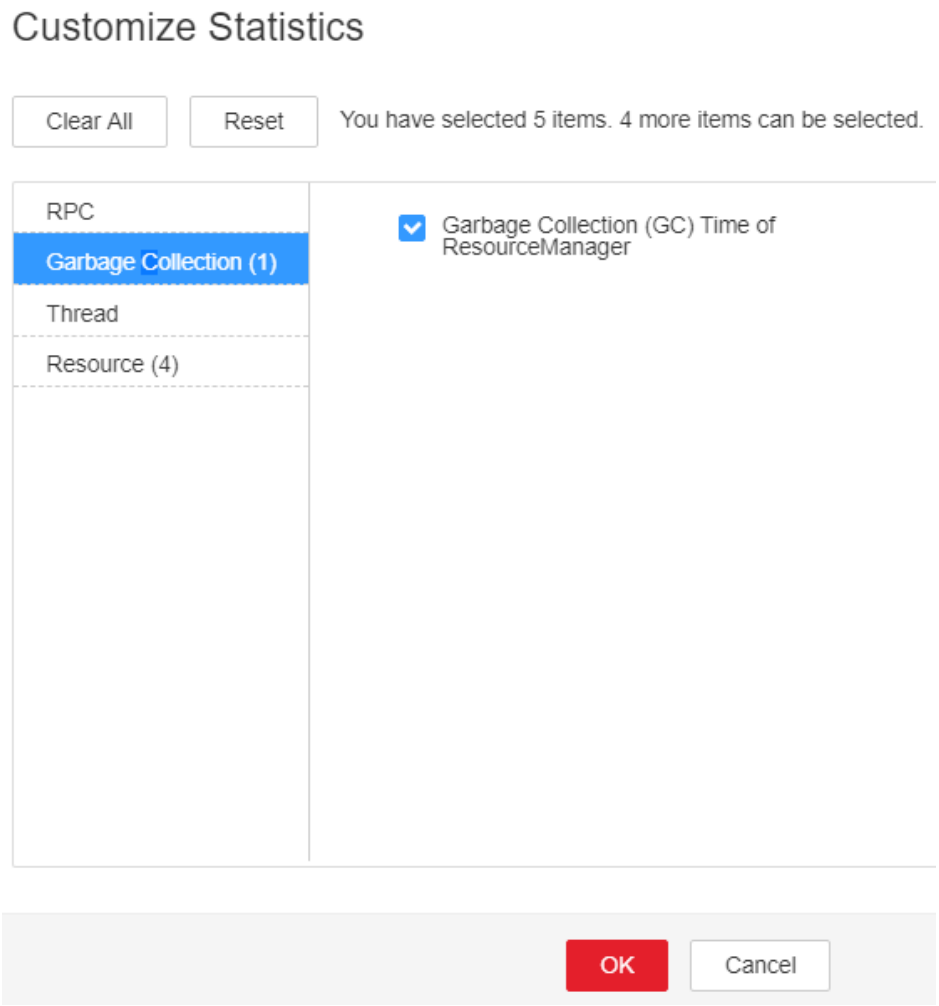

Step 3 Check whether the GC duration of the ResourceManager process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations** > **ResourceManager** > **System** to increase the value of **GC_OPTS** parameter as required.

NOTE

The mapping between the number of NodeManager instances in a cluster and the memory size of ResourceManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms10G -Xmx10G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 1000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms20G -Xmx20G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 2000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms40G -Xmx40G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 3000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms60G -Xmx60G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 4000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms80G -Xmx80G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 5000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms100G -Xmx100G -XX:NewSize=3G -XX:MaxNewSize=6G

Step 5 Save the configuration and restart the ResourceManager instance.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **ResourceManager** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.151 ALM-18011 NodeManager GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the NodeManager process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18011	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

A long GC duration of the NodeManager process may interrupt the services.

Possible Causes

The heap memory of the NodeManager instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18011 NodeManager GC Time Exceeds the Threshold** > **Location** to check the IP address of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn** > **Instance** > **NodeManager (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Garbage Collection (GC) Time of NodeManager** to check the GC duration statistics of the Broker process collected every minute.

Figure 10-43 Garbage Collection (GC) Time of NodeManager

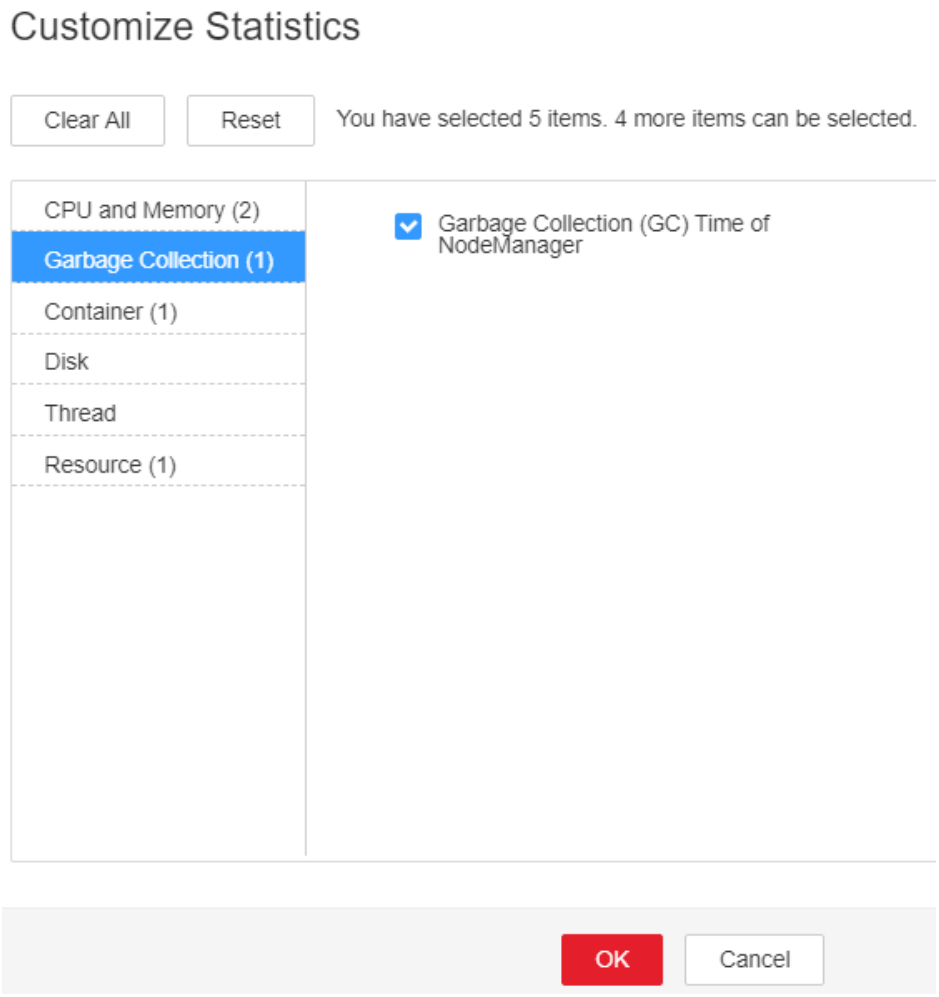

Step 3 Check whether the GC duration of the NodeManager process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn** > **Configurations** > **All Configurations** > **NodeManager** > **System** to increase the value of **GC_OPTS** parameter as required.

NOTE

The mapping between the number of NodeManager instances in a cluster and the memory size of NodeManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters for NodeManager instances are as follows: -Xms2G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters for NodeManager instances are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters for NodeManager instances are as follows: -Xms8G -Xmx8G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Save the configuration and restart the NodeManager instance.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **NodeManager** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.152 ALM-18012 JobHistoryServer GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the JobHistoryServer process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default).

This alarm is cleared when the GC duration is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18012	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

A long GC duration of the JobHistoryServer process may interrupt the services.

Possible Causes

The heap memory of the JobHistoryServer instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18012 JobHistoryServer GC Time Exceeds the Threshold** > **Location** to check the IP address of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **MapReduce** > **Instance** > **JobHistoryServer (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Garbage Collection (GC) Time of the JobHistoryServer** to check the GC duration statistics of the Broker process collected every minute.

Step 3 Check whether the GC duration of the JobHistoryServer process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **MapReduce** > **Configurations** > **All Configurations** > **JobHistoryServer** > **System** to increase the value of **GC_OPTS** parameter as required.

NOTE

The mapping between the number of historical tasks (10000) and the memory of the JobHistoryServer is as follows:

-Xms30G -Xmx30G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Save the configuration and restart the JobHistoryServer instance.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **JobHistoryServer** in the required cluster from the **Service**.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.153 ALM-18013 ResourceManager Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of ResourceManager every 30 seconds. This alarm is generated when the direct memory usage of ResourceManager instances exceeds the threshold (90% of the maximum memory).

This alarm is automatically cleared when the direct memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
18013	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the available direct memory of ResourceManager is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of ResourceManager instances is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-18013 ResourceManager Direct Memory Usage Exceeds the Threshold** > **Location**. View the IP address of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Yarn**. On the page that is displayed, click the **Instances** tab and click the ResourceManager instance for which this alarm is generated. Click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Resource**, and select **Memory Usage Status of ResourceManager** to check the direct memory usage.

Figure 10-44 Customizing ResourceManager memory usage details

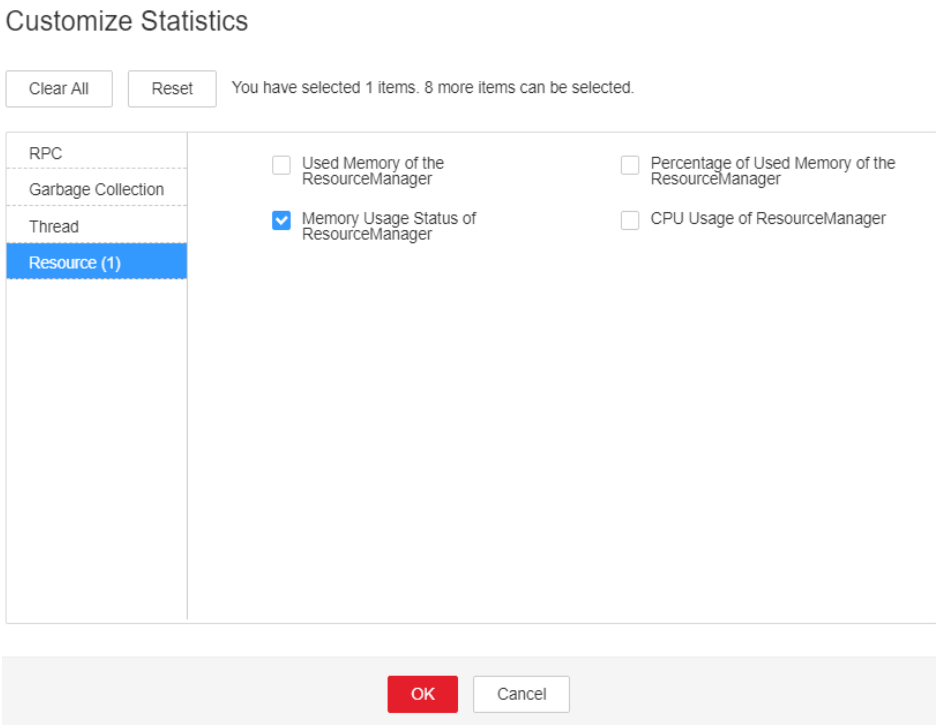

Step 3 Check whether the used direct memory of a ResourceManager instance reaches 90% (default threshold) of the maximum direct memory allocated to it.

- If yes, go to **Step 4**.
- If no, go to **Step 9**.

Step 4 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Yarn** > **Configurations** > **All Configurations** > **ResourceManager** > **System**. Check whether **-XX:MaxDirectMemorySize** exists in the **GC_OPTS** parameter.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Delete the **-XX:MaxDirectMemorySize** parameter from **GC_OPTS** and save the configuration.

NOTE

MaxDirectMemorySize indicates the maximum off-heap memory size. If the **MaxDirectMemorySize** parameter of ResourceManager is not specified, the memory of ResourceManager is not limited. By default, **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter is not set.

Step 6 Perform the following steps to restart the ResourceManager instance:

NOTICE

- Restarting the standby ResourceManager instance does not affect services.
- During the ResourceManager switchover, new jobs cannot be submitted to Yarn, but submitted jobs are not affected.

1. On the Yarn service page, click the **Instances** tab, select the **ResourceManager (Standby)** instance, choose **More**, select **Restart Instance**, and verify the password to restart the instance.
2. After the standby instance is restarted, click the **Dashboard** tab of Yarn, choose **More**, select **Perform ResourceManager Switchover**, and verify the password to perform an active/standby switchover.
3. After the active/standby switchover is complete, click the **Instances** tab on the Yarn service page, select the **ResourceManager (Standby)** instance, choose **More**, select **Restart Instance**, and verify the password to restart the instance. Wait until the instance is restarted.

Step 7 Check whether **ALM-18008 Heap Memory Usage of ResourceManager Exceeds the Threshold** exists.

- If yes, rectify the fault by referring to **ALM-18008 Heap Memory Usage of ResourceManager Exceeds the Threshold**.
- If no, go to **Step 8**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **ResourceManager** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.154 ALM-18014 NodeManager Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the Yarn service every 30 seconds. This alarm is generated when the direct memory usage of a NodeManager instance exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18014	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available direct memory of the Yarn service is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the NodeManager instance is overused or the direct memory is inappropriately allocated.

Procedure

Check the direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **ALM-18014 NodeManager Direct Memory Usage Exceeds the Threshold** > **Location** to check the IP address of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Instance** > **NodeManager (IP address for which the alarm is generated)**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Resource** > **Percentage of Used Memory of the NodeManager** to check the direct memory usage.

Figure 10-45 Percentage of Used Memory of the NodeManager

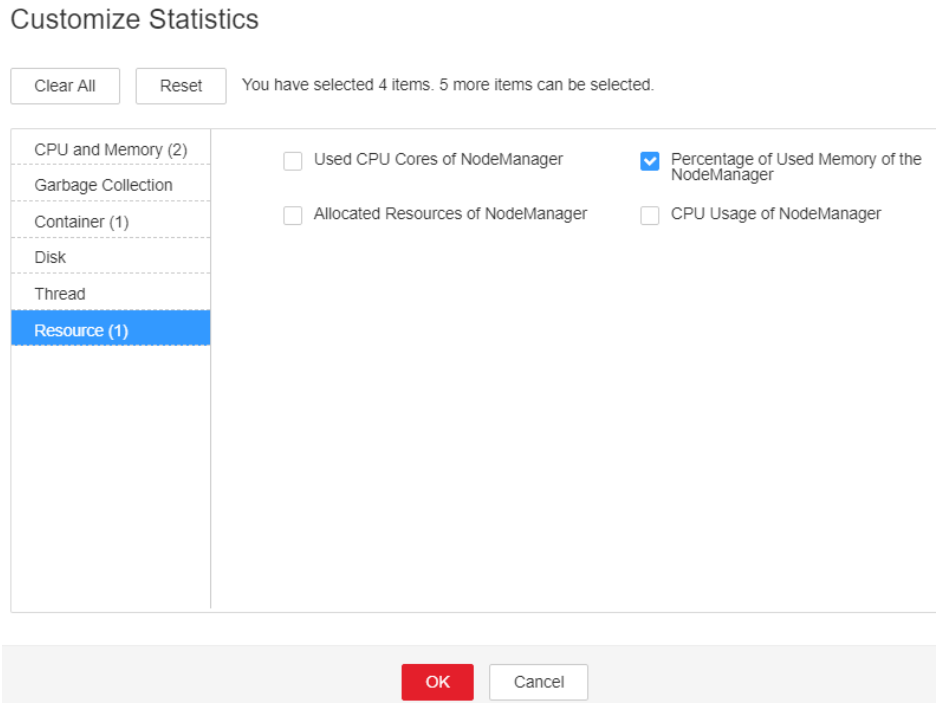

Step 3 Check whether the used direct memory of NodeManager reaches 90% of the maximum direct memory specified for NodeManager by default.

- If yes, go to **Step 4**.
- If no, go to **Step 9**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations**> **NodeManager** > **System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC_OPTS** parameter.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 In the **GC_OPTS** parameter, delete "-XX:MaxDirectMemorySize".

Step 6 Save the configuration and restart the NodeManager instance.

Step 7 Check whether the **ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold** exists.

- If yes, handle the alarm by referring to **ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold**.
- If no, go to **Step 8**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 10 Select **NodeManager** in the required cluster from the **Service**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.155 ALM-18015 JobHistoryServer Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the MapReduce service every 30 seconds. This alarm is generated when the direct memory usage of a JobHistoryServer instance exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18015	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available direct memory of the MapReduce service is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the JobHistoryServer instance is overused or the direct memory is inappropriately allocated.

Procedure

Check the direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18015 JobHistoryServer Direct Memory Usage Exceeds the Threshold** > **Location** to check the IP address of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **MapReduce** > **Instance** > **JobHistoryServer (IP address for which the alarm is generated).** Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Memory Usage Status of JobHistoryServer** to check the direct memory usage.

Step 3 Check whether the used direct memory of JobHistoryServer reaches 90% of the maximum direct memory specified for JobHistoryServer by default.

- If yes, go to **Step 4**.
- If no, go to **Step 9**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **MapReduce** > **Configurations** > **All Configurations** > **JobHistoryServer** > **System** to check whether "-XX:MaxDirectMemorySize" exists in the **GC_OPTS** parameter.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 In the **GC_OPTS** parameter, delete "-XX:MaxDirectMemorySize".

Step 6 Save the configuration and restart the JobHistoryServer instance.

Step 7 Check whether the **ALM-18009 Heap Memory Usage of JobHistoryServer Exceeds the Threshold** exists.

- If yes, handle the alarm by referring to **ALM-18009 Heap Memory Usage of JobHistoryServer Exceeds the Threshold**.
- If no, go to **Step 8**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 10 Select **JobHistoryServer** in the required cluster from the **Service**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.156 ALM-18016 Non Heap Memory Usage of ResourceManager Exceeds the Threshold

Description

The system checks the Non Heap memory usage of Yarn ResourceManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the Non Heap memory usage of Yarn ResourceManager exceeds the threshold (90% of the maximum memory by default).

Users can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** to change the threshold.

The alarm is cleared when the Non Heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18016	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the Non Heap memory usage of Yarn ResourceManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

The Non Heap memory of the Yarn ResourceManager instance on the node is overused or the Non Heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the Non Heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18016 Non Heap Memory Usage of Yarn ResourceManager Exceeds the Threshold** > **Location**. Check the HostName of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn** > **Instance** > **ResourceManager**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Percentage of Used Memory of the ResourceManager**. ResourceManager indicates the

corresponding HostName of the instance for which the alarm is generated. Check the Non Heap memory usage.

Figure 10-46 Percentage of Used Memory of the ResourceManage

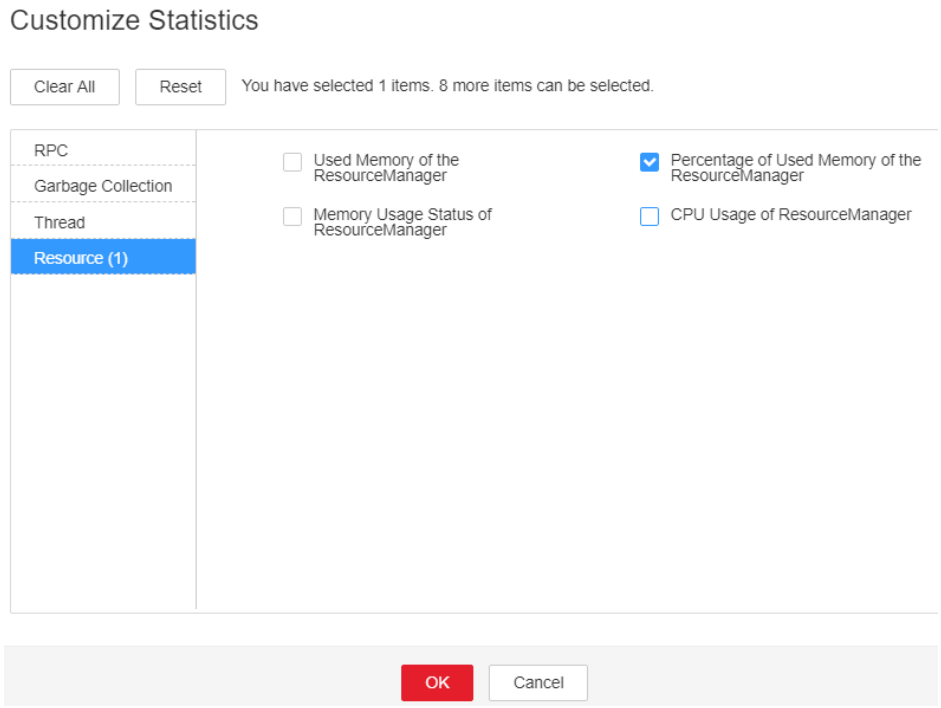

Step 3 Check whether the used Non Heap memory of ResourceManager reaches 90% of the maximum Non Heap memory specified for ResourceManager by default.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations** > **ResourceManager** > **System**. Adjust the **GC_OPTS** memory parameter of ResourceManager. Save the configuration and restart the ResourceManager instance.

NOTE

The mapping between the number of NodeManager instances in a cluster and the memory size of ResourceManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms6G -Xmx6G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms10G -Xmx10G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 1000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms20G -Xmx20G -XX:NewSize=1G -XX:MaxNewSize=2G
- If the number of NodeManager instances in the cluster reaches 2000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms40G -Xmx40G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 3000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms60G -Xmx60G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 4000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms80G -Xmx80G -XX:NewSize=2G -XX:MaxNewSize=4G
- If the number of NodeManager instances in the cluster reaches 5000, the recommended JVM parameters of the ResourceManager instance are as follows: -Xms100G -Xmx100G -XX:NewSize=3G -XX:MaxNewSize=6G

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.157 ALM-18017 Non Heap Memory Usage of NodeManager Exceeds the Threshold

Description

The system checks the Non Heap memory usage of Yarn NodeManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the Non Heap memory usage of Yarn NodeManager exceeds the threshold (90% of the maximum memory by default).

Users can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** to change the threshold.

The alarm is cleared when the Non Heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18017	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the Non Heap memory usage of Yarn NodeManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

The Non Heap memory of the Yarn NodeManager instance on the node is overused or the Non Heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the Non Heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18017 Non Heap Memory Usage of Yarn NodeManager Exceeds the Threshold** > **Location**. Check the HostName of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn** > **Instance** > **NodeManager**. Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **Resource** > **Percentage of Used Memory of the NodeManager**. NodeManager indicates the corresponding HostName of the instance for which the alarm is generated. Check the Non Heap memory usage.

Figure 10-47 Percentage of Used Memory of the NodeManager

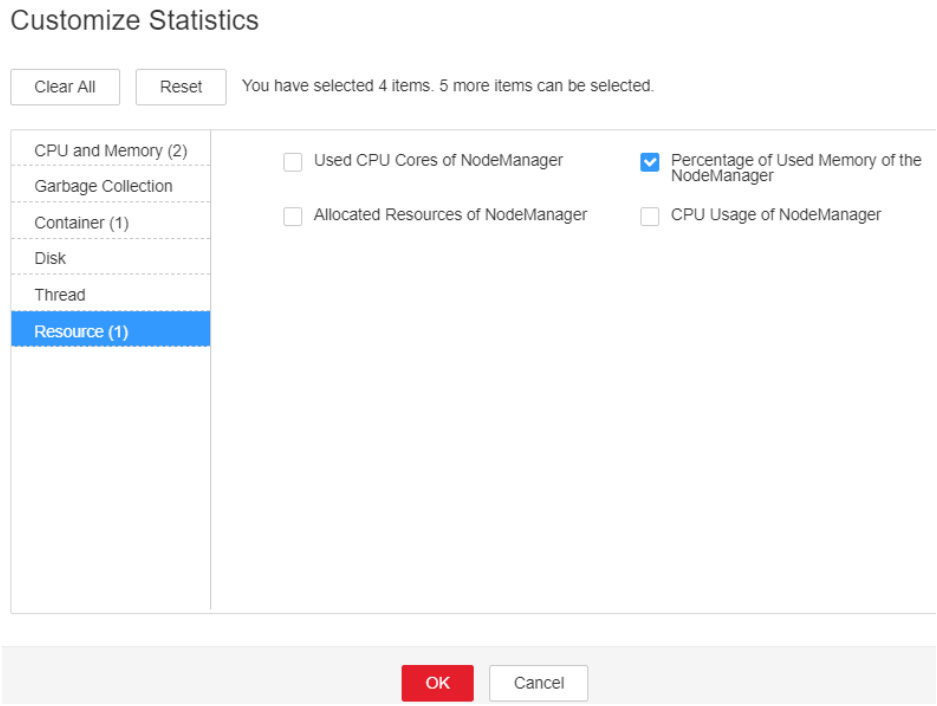

Step 3 Check whether the used Non Heap memory of NodeManager reaches 90% of the maximum Non Heap memory specified for NodeManager by default.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn** > **Configurations** > **All Configurations** > **NodeManager** > **System**. Adjust the **GC_OPTS** memory parameter of NodeManager, click **Save**, and click **OK,** and restart the role instance.

NOTE

The mapping between the number of NodeManager instances in a cluster and the memory size of NodeManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters for NodeManager instances are as follows: -Xms2G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters for NodeManager instances are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters for NodeManager instances are as follows: -Xms8G -Xmx8G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.158 ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold

Description

The system checks the heap memory usage of Yarn NodeManager every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Yarn NodeManager exceeds the threshold (95% of the maximum memory by default).

The alarm is cleared when the heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18018	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the heap memory usage of Yarn NodeManager is overhigh, the performance of Yarn task submission and operation is affected. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

The heap memory of the Yarn NodeManager instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18018 NodeManager Heap Memory Usage Exceeds the Threshold** > **Location**. Check the HostName of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Yarn** > **Instance** > **NodeManager**. Click the drop-down menu

in the upper right corner of **Chart**, choose **Customize** > **Resource** > **Percentage of Used Memory of the NodeManager** to check the heap memory usage.

Figure 10-48 Percentage of Used Memory of the NodeManager

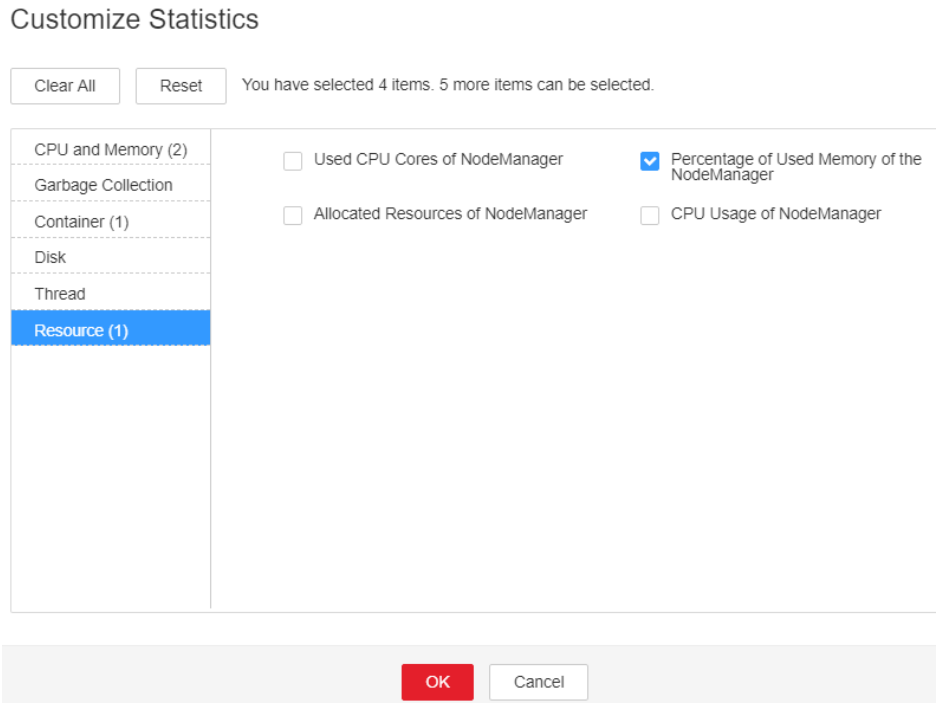

Step 3 Check whether the used heap memory of NodeManager reaches 95% of the maximum heap memory specified for NodeManager.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations** > **NodeManager** > **System**. Increase the value of **GC_OPTS** parameter as required, click **Save**, and click **OK**, and restart the role instance.

NOTE

The mapping between the number of NodeManager instances in a cluster and the memory size of NodeManager is as follows:

- If the number of NodeManager instances in the cluster reaches 100, the recommended JVM parameters for NodeManager instances are as follows: -Xms2G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 200, the recommended JVM parameters for NodeManager instances are as follows: -Xms4G -Xmx4G -XX:NewSize=512M -XX:MaxNewSize=1G
- If the number of NodeManager instances in the cluster reaches 500, the recommended JVM parameters for NodeManager instances are as follows: -Xms8G -Xmx8G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- Yarn

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.159 ALM-18019 Non Heap Memory Usage of JobHistoryServer Exceeds the Threshold

Description

The system checks the Non Heap memory usage of MapReduce JobHistoryServer every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the Non Heap memory usage of MapReduce JobHistoryServer exceeds the threshold (90% of the maximum memory by default).

Users can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **MapReduce** to change the threshold.

The alarm is cleared when the Non Heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18019	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the Non Heap memory usage of MapReduce JobHistoryServer is overhigh, the performance of MapReduce task submission and operation is affected. In addition, a memory overflow may occur so that the MapReduce service is unavailable.

Possible Causes

The Non Heap memory of the MapReduce JobHistoryServer instance on the node is overused or the Non Heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the Non Heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** > **ALM-18019 Non Heap Memory Usage of MapReduce JobHistoryServer Exceeds the Threshold** > **Location**. Check the HostName of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **MapReduce** > **Instance** > **JobHistoryServer.** Click the drop-down menu in the upper right corner of **Chart**, choose **Customize** > **JobHistoryServer Non Heap memory usage statistics**. JobHistoryServer indicates the corresponding HostName of the instance for which the alarm is generated. Check the Non Heap memory usage.

Step 3 Check whether the used Non Heap memory of JobHistoryServer reaches 90% of the maximum Non Heap memory specified for JobHistoryServer.

- If yes, go to **Step 4**.

- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **MapReduce** > **Configurations** > **All Configurations** > **JobHistoryServer** > **System**. Adjust the **GC_OPTS** memory parameter of the NodeManager, click **Save**, and click **OK,** and restart the role instance.

NOTE

The mapping between the number of historical tasks (10000) and the memory of the JobHistoryServer is as follows:

-Xms30G -Xmx30G -XX:NewSize=1G -XX:MaxNewSize=2G

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 7 Select the following node in the required cluster from the **Service**.

- NodeAgent
- MapReduce

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.160 ALM-18020 Yarn Task Execution Timeout

Alarm Description

The system checks MapReduce and Spark tasks (except for permanent JDBC tasks) submitted to Yarn every 15 minutes. This alarm is generated when the task execution time exceeds the timeout duration specified by the user. However, the task can be properly executed. The client timeout parameter of MapReduce is mapreduce.application.timeout.alarm and that of Spark is spark.application.timeout.alarm. The unit is ms.

This alarm is cleared when the task is finished or terminated.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
18020	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
ApplicationName	Specifies the object (application ID) for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The alarm persists after task execution times out. However, the task can still be properly executed, so this alarm does not exert any impact on the system.

Possible Causes

- The specified timeout duration is shorter than the required execution time.
- The queue resources for task running are insufficient.
- Task data skew occurs. As a result, some tasks process a large amount of data and take a long time to execute.

Handling Procedure

Check whether the timeout interval is correctly set.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. The **Alarms** page is displayed.

Step 2 Select the alarm whose ID is **18020**. In the alarm details, view **Location** to obtain the timeout task name and timeout duration.

Step 3 Based on the task name and timeout interval, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **ResourceManager (Active)** to log in to the native Yarn page. Then find the task on the native page, check its **StartTime** and calculate the task execution time based on the current system time. Check whether the task execution time exceeds the timeout duration.

- If yes, go to **Step 4**.
- If no, go to **Step 10**.

Step 4 Evaluate the expected task execution time based on the service and compare it with the task timeout interval. If the timeout interval is too short, set the timeout interval (**mapreduce.application.timeout.alarm** or **spark.application.timeout.alarm**) of the client to the task expected execution time. Run the task again and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the queue resources are sufficient.

Step 5 Find the task on the native page and view the queue name of the task. Click **Scheduler** on the left of the native page. On the **Applications Queues** page, find the corresponding queue name and expand the queue details, as shown in the following figure.

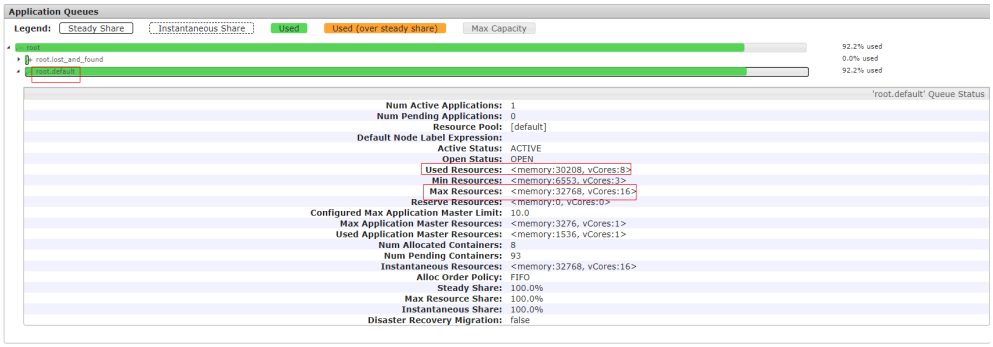

Step 6 Check whether the value of **Used Resources** in the queue details is approximately equal to the value of **Max Resources**, which indicates that the resources in the queue submitted by the task have been used up. If the queue resources are insufficient, choose **Tenant Resources** > **Dynamic Resource Plan** > **Resource Distribution Policy** on FusionInsight Manager and increase the value of **Max Resources** for the queue. Run the task again and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check whether data skew occurs.

Step 7 On the native Yarn page, click *task ID* (for example, **application_1565337919723_0002**) > **Tracking URL:ApplicationMaster** > **job_1565337919723_0002**. The following page is displayed.

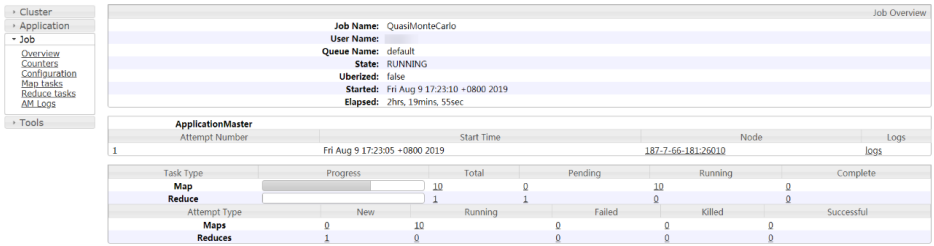

Step 8 Choose **Job** > **Map tasks** or **Job** > **Reduce tasks** on the left and check whether the execution time of each Map or Reduce task differs greatly. If yes, task data skew occurs. In this case, you need to balance the task data.

Step 9 Rectify the fault based on the preceding causes and perform the tasks again. Then, check whether the alarm persists.

- If yes, go to **Step 10**.
- If no, no further action is required.

Collect the fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 12 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.161 ALM-18021 Mapreduce Service Unavailable

Description

The alarm module checks the MapReduce service status every 60 seconds. This alarm is generated when the system detects that the MapReduce service is unavailable.

The alarm is cleared when the MapReduce service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
18021	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The cluster cannot provide the MapReduce service. For example, MapReduce cannot be used to view task logs or the log archive function is unavailable.

Possible Causes

- The JobHistoryServer instance is abnormal.
- The KrbServer service is abnormal.
- The ZooKeeper service abnormal.
- The HDFS service abnormal.
- The Yarn service is abnormal.

Procedure

Check MapReduce service JobHistoryServer instance status.

Step 1 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **MapReduce** > **Instance**.

Step 2 Check whether the running status of JobHistoryServer is **Normal**.

- If yes, go to **Step 11**.
- If no, go to **Step 3**.

Check the KrbServer service status.

Step 3 In the alarm list on FusionInsight Manager, check whether **ALM-25500 KrbServer Service Unavailable** exists.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Rectify the fault by following the steps provided in **ALM-25500 KrbServer Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the ZooKeeper service.

Step 5 In the alarm list on FusionInsight Manager, check whether **ALM-13000 ZooKeeper Service Unavailable** exists.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Rectify the fault by following the steps provided in **ALM-13000 ZooKeeper Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check the HDFS service status.

Step 7 In the alarm list on FusionInsight Manager, check whether **ALM-14000 HDFS Service Unavailable** exists.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Rectify the fault by following the steps provided in **ALM-14000 HDFS Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check the Yarn service status.

Step 9 In the alarm list on FusionInsight Manager, check whether **ALM-18000 Yarn Service Unavailable** exists.

- If yes, go to **Step 10**
- If no, go to **Step 11**.

Step 10 Rectify the fault by following the steps provided in **ALM-18000 Yarn Service Unavailable**, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On the FusionInsight Manager home page of the active cluster, choose **O&M** > **Log > Download.**

Step 12 Select **MapReduce** in the required cluster from the **Service.**

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.162 ALM-18022 Insufficient Yarn Queue Resources

Description

The alarm module checks Yarn queue resources every 60 seconds. This alarm is generated when available resources or ApplicationMaster (AM) resources of a queue are insufficient.

This alarm is cleared when available resources are sufficient.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18022	Minor	Yes

Parameters

Parameter Name	Description
Source	Specifies the cluster for which the alarm is generated.
QueueName	Specifies the queue for which the alarm is generated.
QueueMetric	Specifies the metric of the queue for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

- An application being executed takes longer time.
- An application fails to be executed for a long time after being submitted.

Possible Causes

- NodeManager node resources are insufficient.
- The configured maximum resource capacity of the queue is excessively small.
- The configured maximum AM resource percentage is excessively small.

Procedure

View alarm details.

Step 1 On the FusionInsight Manager, choose **O&M** > **Alarm > Alarms**.

Step 2 View location information of this alarm and check whether **QueueName** is **root** and **QueueMetric** is **Memory** or **QueueName** is **root** and **QueueMetric** is **vCores**.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 The memory or CPU of the Yarn cluster is insufficient. In this case, log in to the node where NodeManager resides and run the **free -g** and **cat /proc/cpuinfo** commands to query the available memory and available CPU of the node, respectively. On FusionInsight Manager, increase the values of **yarn.nodemanager.resource.memory-mb** and **yarn.nodemanager.resource.cpu-vcores** for the Yarn NodeManager based on the query results. Then, restart the NodeManager instance. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Step 4 View location information of this alarm and check whether **QueueName** is **<*Tenant Queue*>** and **QueueMetric** is **Memory**, or **QueueName** is **<*Tenant Queue*>** and **QueueMetric** is **vCores** in **Location**, check whether **available Memory =** or **available vCores =** are included in **Additional Information**.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 The memory or CPU of the tenant queue is insufficient. In this case, choose **Tenant Resources** > **Dynamic Resource Plan > Resource Distribution Policy** and increase the value of **Maximum Capacity**. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Step 6 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations**. Enter the keyword "threshold" and click **ResourceManager**. Adjust the threshold values of the following parameters:

If **Additional Information** contains **available Memory =**, change the value of **yarn.queue.memory.alarm.threshold** to a value smaller than that of **available Memory =** in **Additional Information**.

If **Additional Information** contains **available vCores =**, change the value of **yarn.queue.vcore.alarm.threshold** to a value smaller than that of **available vCores =** in **Additional Information**.

Wait for five minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 7 If **available AmMemory =** or **available AmvCores =** is included in **Additional Information**, ApplicationMaster memory or CPU of the tenant queue is

insufficient. In this case, choose **Tenant Resources** > **Dynamic Resource Plan** > **Queue Configuration** and increase the value of **Maximum Am Resource Percent**. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Step 8 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **Configurations** > **All Configurations**. Enter the keyword "threshold" and click **ResourceManager**. Adjust the threshold values of the following parameters:

If **Additional Information** contains **available AmMemory =**, change the value of **yarn.queue.memory.alarm.threshold** to a value smaller than that of **available AmMemory =** in **Additional Information**.

If **Additional Information** contains **available AmvCores =**, change the value of **yarn.queue.vcore.alarm.threshold** to a value smaller than that of **available AmvCores =** in **Additional Information**.

Wait for five minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 Log in to FusionInsight Manager of the active cluster, and choose **O&M** > **Log** > **Download**.

Step 10 Select **Yarn** in the required cluster from the **Service**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Reference

None

10.163 ALM-18023 Number of Pending Yarn Tasks Exceeds the Threshold

Alarm Description

The alarm module checks the number of pending applications in the Yarn root queue every 60 seconds. The alarm is generated when the number exceeds 60.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
18023	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
QueueName	Specifies the queue for which the alarm was generated.
QueueMetric	Specifies the queue metric for which the alarm was generated.

Impact on the System

- It takes long time to end an application.
- A new application cannot run after submission.

Possible Causes

- NodeManager node resources are insufficient.
- The maximum resource capacity of the queue and the maximum AM resource percentage are too small.
- The monitoring threshold is too small.

Handling Procedure

Check NodeManager resources.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **ResourceManager (Active)** to access the ResourceManager web UI.

Step 2 Click **Scheduler** and check whether the root queue resources are used up in **Application Queues**.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Expand the capacity of the NodeManager instance of the Yarn service. After the capacity expansion, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check the maximum queue resource capacity and the maximum AM resource percentage.

Step 4 Check whether the resources of the queue corresponding to the pending task are used up.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **Tenant Resources** > **Dynamic Resource Plan** and add resources as required. Check whether the alarms are cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Adjust the monitoring thresholds.

Step 6 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** > **Applications** > **Pending Applications**, and increase the thresholds as required.

Step 7 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.164 ALM-18024 Pending Yarn Memory Usage Exceeds the Threshold

Alarm Description

The alarm module checks the pending memory of Yarn every 60 seconds. The alarm is generated when the pending memory exceeds the threshold. Pending

memory indicates the total memory that is not allocated to submitted Yarn applications.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
18024	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
QueueName	Specifies the queue for which the alarm was generated.
QueueMetric	Specifies the queue metric for which the alarm was generated.

Impact on the System

- It takes long time to end an application.
- A new application cannot run after submission.

Possible Causes

- NodeManager node resources are insufficient.
- The maximum resource capacity of the queue and the maximum AM resource percentage are too small.
- The monitoring threshold is too small.

Handling Procedure

Check NodeManager resources.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **ResourceManager (Active)** to access the ResourceManager web UI.

Step 2 Click **Scheduler** and check whether the root queue resources are used up in **Application Queues**.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Expand the capacity of the NodeManager instance of the Yarn service. After the capacity expansion, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check the maximum queue resource capacity and the maximum AM resource percentage.

Step 4 Check whether the resources of the queue corresponding to the pending task are used up.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **Tenant Resources** > **Dynamic Resource Plan** and add resources as required. Check whether the alarms are cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Adjust the monitoring thresholds.

Step 6 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** > **CPU and Memory** > **Pending Memory**, and increase the threshold as required.

Step 7 Check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.165 ALM-18025 Number of Terminated Yarn Tasks Exceeds the Threshold

Description

The alarm module checks the number of terminated applications in the Yarn root queue every 60 seconds. The alarm is generated when the number exceeds 50 for three consecutive times.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18025	Major	Yes

Parameters

Name	Meaning
Cluster Name	Specifies the cluster for which the alarm is generated.
Service Name	Specifies the service for which the alarm is generated.
Role Name	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

A large number of application tasks are forcibly terminated.

Possible Causes

- The user forcibly terminates a large number of tasks.
- The system terminates tasks due to some error.

Procedure

Check the alarm details.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** to go to the alarm page.

Step 2 View **Additional Information** in the alarm details to check whether the alarm threshold is too small.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** > **Other** > **Terminated Applications of root queue** to modify the threshold. Go to **Step 6**.

Step 4 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **ResourceManager(Active)** to access the ResourceManager web UI.

Step 5 Click **KILLED** in **Applications** and click the task on the top. View the description of **Diagnostics** and rectify the fault based on the task termination details (for example, the task is terminated by a user).

Figure 10-49 Click **KILLED**

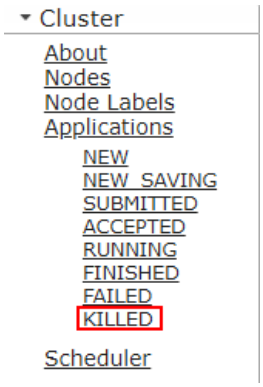

Step 6 Wait for 3 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On the FusionInsight Manager, choose **O&M > Log > Download**.

Step 8 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.166 ALM-18026 Number of Failed Yarn Tasks Exceeds the Threshold

Description

The alarm module checks the number of failed applications in the Yarn root queue every 60 seconds. The alarm is generated when the number exceeds 50 for three consecutive times.

Attribute

Alarm ID	Alarm Severity	Auto Clear
18026	Major	Yes

Parameters

Name	Meaning
Cluster Name	Specifies the cluster for which the alarm is generated.
Service Name	Specifies the service for which the alarm is generated.
Role Name	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

- A large number of application tasks fail to be executed.
- Failed tasks need to be submitted again.

Possible Causes

The task fails to be executed due to some error.

Procedure

Check the alarm details.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** to go to the alarm page.

Step 2 View **Additional Information** in the alarm details to check whether the alarm threshold is too small.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Yarn** > **Other** > **Failed Applications of root queue** to modify the threshold. Go to **Step 6**.

Step 4 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** > **ResourceManager(Active)** to access the ResourceManager web UI.

Step 5 Click **FAILED** in **Applications** and click the task on the top. View the description of **Diagnostics** and rectify the fault based on the task failure causes.

Figure 10-50 Click **FAILED**

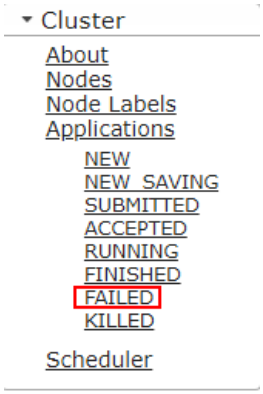

Step 6 Wait for 3 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On the FusionInsight Manager, choose O&M > Log > Download.

Step 8 Expand the **Service** drop-down list, and select **Yarn** for the target cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.167 ALM-19000 HBase Service Unavailable

Alarm Description

The alarm module checks the HBase service status every 120 seconds. This alarm is generated when the HBase service is unavailable.

This alarm is cleared when the HBase service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19000	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Operations such as data read/write and table creation cannot be performed.

Possible Causes

- ZooKeeper is abnormal.
- HDFS is abnormal.
- HBase is abnormal.
- The network connection is abnormal.
- The service configuration value is incorrect.

Handling Procedure

Check the ZooKeeper service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of ZooKeeper is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 2**.

Step 2 In the alarm list, check whether **ALM-13000 ZooKeeper Service Unavailable** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by performing the operations provided for **ALM-13000 ZooKeeper Service Unavailable**.

Step 4 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the HDFS service status.

Step 5 In the alarm list, check whether **ALM-14000 HDFS Service Unavailable** exists.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 Rectify the fault by performing the operations provided for **ALM-14000 HDFS Service Unavailable**.

Step 7 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Step 8 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services** > **HDFS**, and check whether **Safe Mode** of HDFS is **ON**.

- If yes, go to **Step 9**.
- If no, go to **Step 12**.

Step 9 Log in to the HDFS client as user **root**. Run the **cd** command to go to the client installation directory and run the **source bigdata_env** command.

If the cluster uses the security mode, perform security authentication. Obtain the password of user **hdfs** from the MRS cluster administrator, run the **kinit hdfs** command, and enter the password as prompted.

Step 10 Run the following command to manually exit the safe mode:

hdfs dfsadmin -safemode leave

Step 11 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check the HBase service status.

Step 12 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **HBase**.

Step 13 Check whether there is one active HMaster and one standby HMaster.

- If yes, go to **Step 15**.
- If no, go to **Step 14**.

Step 14 Click **Instances** and select the HMaster instance whose status is not **Active**. Click **More** and select **Restart Instance** to restart HMaster. Then check whether there is one active HMaster and one standby HMaster.

- If yes, go to **Step 15**.
- If no, go to **Step 21**.

Step 15 Choose **Cluster**, click the name of the desired cluster, choose **Services** > **HBase**, and click **HMaster(Active)** to access the HMaster web UI.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 16 Check whether at least one RegionServer exists under **Region Servers**.

- If yes, go to **Step 17**.
- If no, go to **Step 21**.

Step 17 Choose **Tables** > **System Tables** and check whether **hbase:meta**, **hbase:namespace**, and **hbase:acl** exist in the **Table Name** column, as shown in **Figure 10-51**.

- If yes, go to **Step 18**.
- If no, go to **Step 19**.

Figure 10-51 HBase system tables

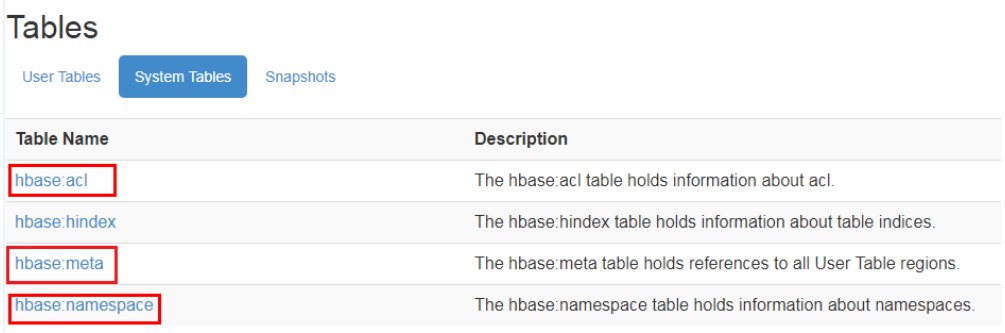

Step 18 Click **hbase:meta**, **hbase:namespace**, and **hbase:acl** to check whether all pages can be opened. If all of them can be opened, the tables are normal.

- If yes, go to **Step 19**.
- If no, go to **Step 25**.

> **NOTE**
>
> In a normal cluster, ACL permission control is disabled for HBase by default. The **hbase:acl** table is generated only after ACL permission control is manually enabled. In this case, you need to check this table.

Step 19 View the HMaster startup status.

On the **Tasks** page shown in **Figure 10-52**, the **RUNNING** value in the **State** column indicates that HMaster is being started and provides how much time HMaster keeps in that state. As shown in **Figure 10-53**, if the state is **COMPLETE**, HMaster has been started.

Check whether HMaster has been in the **RUNNING** state for a long time.

Figure 10-52 HMaster being started

Figure 10-53 HMaster startup completed

- If yes, go to **Step 20**.
- If no, go to **Step 21**.

Step 20 On the HMaster web UI, check whether any **hbase:meta** is in the **Regions in Transition** state for a long time.

Figure 10-54 Regions in Transition

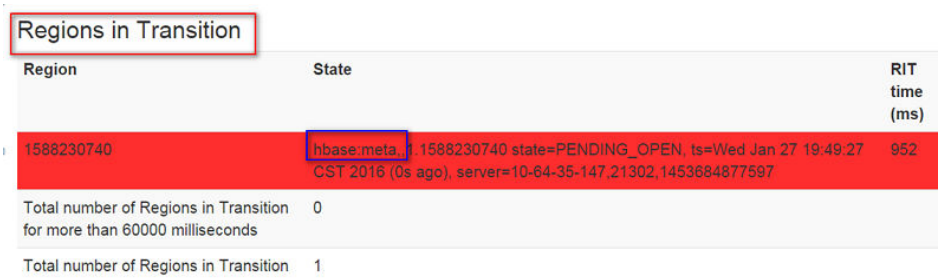

- If yes, go to **Step 21**.
- If no, go to **Step 22**.

Step 21 After ensuring that services are not affected, log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services** > **HBase**, click **More**, and select **Restart Service**. In the dialog box that is displayed, enter the password, and click **OK**.

- If yes, go to **Step 22**.
- If no, go to **Step 25**.

Step 22 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 25**.

Check whether the HBase configurations are correctly modified.

Step 23 On FusionInsight Manager, choose **Audit**. On the **Audit** page, click **Advanced Search**, click ··· on the right of **Operation Type**, select **Save configuration**, click **OK**, and click **Search**.

Step 24 In the search result, check whether the historical configurations of HBase-related services in the **Service** column, such as ZooKeeper, HDFS, and HBase, may affect the HBase service status. **Table 10-6** lists some configurations that may affect the HBase service status.

Table 10-6 Configurations affecting the HBase service status

Parameter	Possible Impact
GC_OPTS	The memory configuration may be improper. You need to check the health status of instance processes.
hbase.rpc.protection	If the HBase service is not restarted offline after the value of this parameter is changed, the connection authentication fails and the HBase service becomes abnormal.
hbase.regionserver.metahandler.count	If there are too many regions in the cluster but this parameter is set to a small value, RIT may occur and regions cannot be brought online for a long time.
hbase.regionserver.thread.compaction.large	If this parameter is set to a large value, the node CPU usage may be too high.
hbase.regionserver.thread.compaction.small	If this parameter is set to a large value, the node CPU usage may be too high.
hbase.coprocessor.master.classes	If a custom coprocessor is used in the configuration, a logic error may cause the service to be unavailable.
hbase.coprocessor.region.classes	If a custom coprocessor is used in the configuration, a logic error may cause the service to be unavailable.
hbase.coprocessor.regionserver.classes	If a custom coprocessor is used in the configuration, a logic error may cause the service to be unavailable.

Parameter	Possible Impact
zookeeper.session.timeout	If this parameter is set to a small value, the connection between HBase and ZooKeeper times out too quickly. As a result, the HMaster instance and RegionServer may restart repeatedly.

Check the network connection between HMaster and dependent components.

Step 25 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **HBase**.

Step 26 Click **Instances**. In the HMaster instance list, record the management IP address of the active HMaster instance.

Step 27 Log in to the active HMaster node as user **omm** through the IP address obtained in **Step 26**.

Step 28 Run the **ping** command to check whether the network connection between the active HMaster node and the host where the dependent components reside is normal. (The dependent components include ZooKeeper, HDFS, and Yarn. The method of obtaining the IP address of the host where the dependent components reside is the same as that of obtaining the IP address of the active HMaster node.)

- If yes, go to **Step 31**.
- If no, go to **Step 29**.

Step 29 Contact the network administrator to restore the network.

Step 30 In the alarm list, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 31**.

Collect fault information.

Step 31 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 32 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select the following services for the target cluster:

- ZooKeeper
- HDFS
- HBase

Step 33 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 34 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.168 ALM-19006 HBase Replication Sync Failed

Description

The alarm module checks the HBase DR data synchronization status every 30 seconds. When disaster recovery (DR) data fails to be synchronized to a standby cluster, the alarm is triggered.

When DR data synchronization succeeds, the alarm is cleared.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
19006	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

HBase data in a cluster fails to be synchronized to the standby cluster, causing data inconsistency between active and standby clusters.

Possible Causes

- The HBase service on the standby cluster is abnormal.
- A network exception occurs.

Procedure

Observe whether the system automatically clears the alarm.

Step 1 On the FusionInsight Manager portal of the active cluster, click **O&M** > **Alarm** > **Alarms.**

Step 2 In the alarm list, click the alarm to obtain alarm generation time from **Generated** of the alarm. Check whether the alarm has existed for five minutes.

- If yes, go to **Step 4**.
- If no, go to **Step 3**.

Step 3 Wait five minutes and check whether the system automatically clears the alarm.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the HBase service status of the standby cluster.

Step 4 Log in to the FusionInsight Manager portal of the active cluster, and click **O&M** > **Alarm** > **Alarms.**

Step 5 In the alarm list, click the alarm to obtain **HostName** from **Location**.

Step 6 Access the node where the HBase client of the active cluster resides as user **omm**.

If the cluster uses a security mode, perform security authentication first and then access the **hbase shell** interface as user **hbase**.

cd /opt/client

source ./bigdata_env

kinit *hbaseuser*

Step 7 Run the **status 'replication', 'source'** command to check the DR synchronization status of the faulty node.

The DR synchronization status of a node is as follows.

```
10-10-10-153:
SOURCE: PeerID=abc, SizeOfLogQueue=0, ShippedBatches=2, ShippedOps=2, ShippedBytes=320,
LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0,
TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=0,
TimeStampsOfLastShippedOp=Mon Jul 18 09:53:28 CST 2016, Replication Lag=0,
FailedReplicationAttempts=0
SOURCE: PeerID=abc1, SizeOfLogQueue=0, ShippedBatches=1, ShippedOps=1, ShippedBytes=160,
LogReadInBytes=1636, LogEditsRead=5, LogEditsFiltered=3, SizeOfLogToReplicate=0,
TimeForLogToReplicate=0, ShippedHFiles=0, SizeOfHFileRefsQueue=0, AgeOfLastShippedOp=16788,
TimeStampsOfLastShippedOp=Sat Jul 16 13:19:00 CST 2016, Replication Lag=16788,
FailedReplicationAttempts=5
```

Step 8 Obtain **PeerID** corresponding to a record whose **FailedReplicationAttempts** value is greater than 0.

In the preceding step, data on the faulty node 10-10-10-153 fails to be synchronized to a standby cluster whose **PeerID** is **abc1**.

Step 9 Run the **list_peers** command to find the cluster and the HBase instance corresponding to the **PeerID** value.

```
PEER_ID CLUSTER_KEY STATE TABLE_CFS
abc1 10.10.10.110,10.10.10.119,10.10.10.133:2181:/hbase2 ENABLED
abc 10.10.10.110,10.10.10.119,10.10.10.133:2181:/hbase ENABLED
```

In the preceding information, **/hbase2** indicates that data is synchronized to the HBase2 instance of the standby cluster.

Step 10 In the service list of FusionInsight Manager of the standby cluster, check whether the running status of the HBase instance obtained by using **Step 9** is **Normal**.

- If yes, go to **Step 14**.
- If no, go to **Step 11**.

Step 11 In the alarm list, check whether the **ALM-19000 HBase Service Unavailable** alarm is generated.

- If yes, go to **Step 12**.
- If no, go to **Step 14**.

Step 12 Follow troubleshooting procedures in **ALM-19000 HBase Service Unavailable** to rectify the fault.

Step 13 Wait for a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Check network connections between RegionServers on active and standby clusters.

Step 14 Log in to the FusionInsight Manager portal of the active cluster, and click **O&M** > **Alarm** > **Alarms.**

Step 15 In the alarm list, click the alarm to obtain **HostName** from **Location**.

Step 16 Use the IP address obtained in **Step 15** to log in to a faulty RegionServer node as user **omm**.

Step 17 Run the **ping** command to check whether network connections between the faulty RegionServer node and the host where RegionServer of the standby cluster resides are in the normal state.

- If yes, go to **Step 20**.
- If no, go to **Step 18**.

Step 18 Contact the network administrator to restore the network.

Step 19 After the network is running properly, check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 20**.

Collect fault information.

Step 20 On the FusionInsight Manager interface of active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 21 In the **Service** drop-down list box, select **HBase** in the required cluster.

Step 22 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 23 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.169 ALM-19007 HBase GC Time Exceeds the Threshold

Description

The system checks the old generation garbage collection (GC) time of the HBase service every 60 seconds. This alarm is generated when the detected old generation GC time exceeds the threshold (exceeds 5 seconds for three consecutive checks by default). To change the threshold, on the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **HBase > GC > GC time for old generation**. This alarm is cleared when the old generation GC time of the HBase service is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
19007	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.

Name	Meaning
HostName	Specifies the object (host ID) for which the alarm is generated.

Impact on the System

If the old generation GC time exceeds the threshold, HBase data read and write are affected.

Possible Causes

The memory of HBase instances is overused, the heap memory is inappropriately allocated, or a large number of I/O operations exist in HBase. As a result, GCs occur frequently.

Procedure

Check the GC time.

Step 1 On the FusionInsight Manager portal, click **O&M** > **Alarm** > **Alarms** and select the alarm whose **ID** is **19007**. Then check the role name in **Location** and confirm the IP adress of the instance.

- If the role for which the alarm is generated is HMaster, go to **Step 2**.
- If the role for which the alarm is generated is RegionServer, go to **Step 3**.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the HMaster for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **GC** > **Garbage Collection (GC) Time of HMaster** and click **OK** to check whether the value of **GC time for old generation** is greater than the threshold (exceeds 5 seconds for three consecutive checks periods by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Figure 10-55 Garbage Collection (GC) Time of HMaster

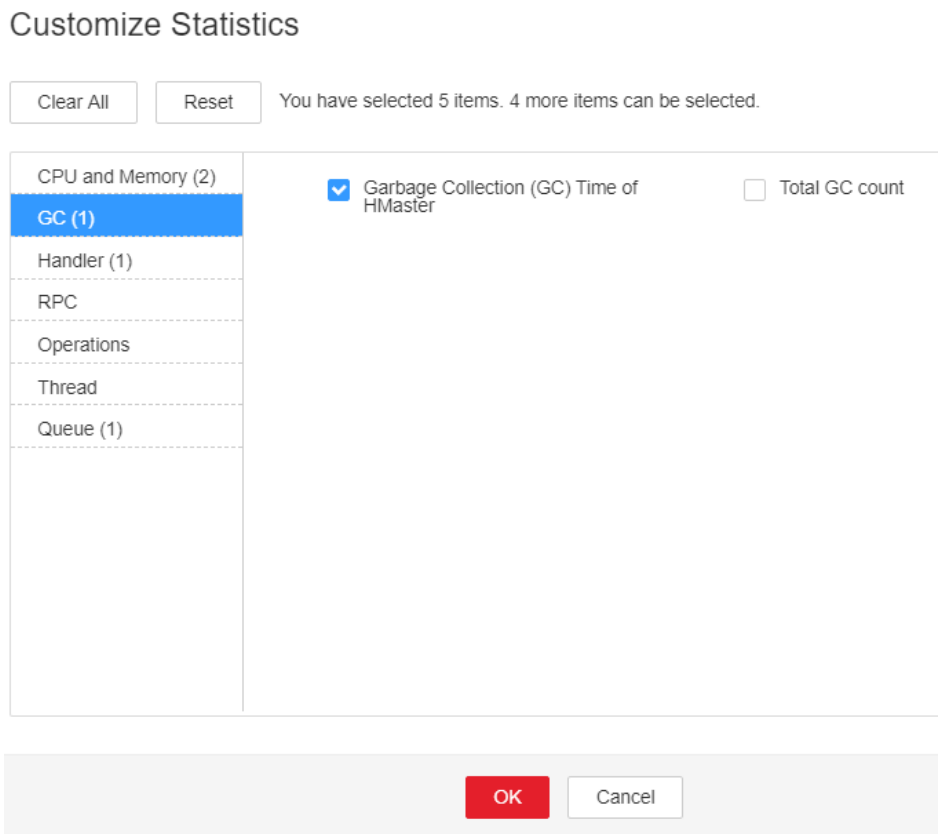

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the RegionServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **GC** > **Garbage Collection (GC) Time of RegionServer** and click **OK** to check whether the value of **GC time for old generation** is greater than the threshold (exceeds 5 seconds for three consecutive checks periods by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Figure 10-56 Garbage Collection (GC) Time of RegionServer

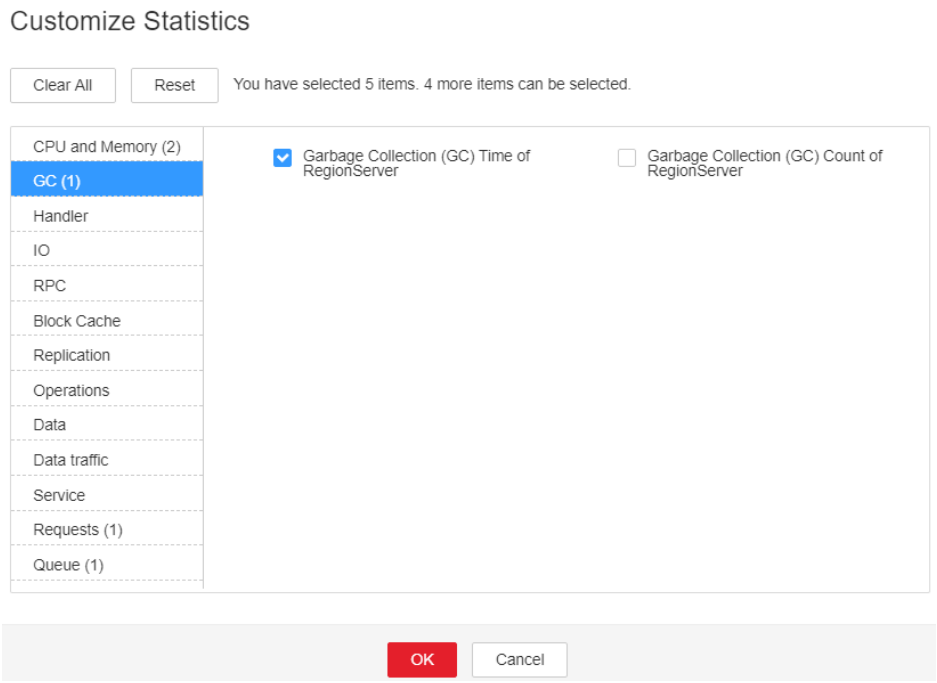

Check the current JVM configuration.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. In Search, enter **GC_OPTS** to check the **GC_OPTS** memory parameter of role HMaster(HBase->HMaster), RegionServer(HBase->RegionServer). Adjust the values of **-Xmx** and **-XX:CMSInitiatingOccupancyFraction** of the GC_OPTS parameter by referring to the Note.

NOTE

1. Suggestions on GC parameter configurations for HMaster
   - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
   - Set **-XX:NewSize** to the value of **-XX:MaxNewSize**, which is one eighth of **-Xmx**.
   - For large-scale HBase clusters with a large number of regions, increase values of **GC_OPTS** parameters for HMaster. Specifically, set **-Xmx** to 4 GB if the number of regions is less than 100,000. If the number of regions is more than 100,000, set -Xmx to be greater than or equal to 6 GB. For each increased 35,000 regions, increase the value of **-Xmx** by 2 GB. The maximum value of **-Xmx** is 32 GB.
2. Suggestions on GC parameter configurations for RegionServer
   - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
   - Set **-XX:NewSize** to one eighth of **-Xmx**.
   - Set the memory for RegionServer to be greater than that for HMaster. If sufficient memory is available, increase the heap memory.
   - Set **-Xmx** based on the machine memory size. Specifically, set **-Xmx** to 32 GB if the machine memory is greater than 200 GB, to 16 GB if the machine memory is greater than 128 GB and less than 200 GB, and to 8 GB if the machine memory is less than 128 GB. When **-Xmx** is set to 32 GB, a RegionServer node supports 2000 regions and 200 hotspot regions.
   - **XX:CMSInitiatingOccupancyFraction** to be less than and equal to **85**, and it is calculated as follows: 100 x (hfile.block.cache.size + hbase.regionserver.global.memstore.size)

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager interface of active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 7 In the **Service** drop-down list box, select **HBase** in the required cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.170 ALM-19008 Heap Memory Usage of the HBase Process Exceeds the Threshold

Description

The system checks the HBase service status every 30 seconds. The alarm is generated when the heap memory usage of an HBase service exceeds the threshold (90% of the maximum memory).

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
19008	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

Impact on the System

If the available HBase heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the HBase service is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, click **O&M** > **Alarm** > **Alarms** and select the alarm whose **ID** is **19008**. Then check the role name in **Location** and confirm the IP adress of the instance.

- If the role for which the alarm is generated is HMaster, go to **Step 2**.
- If the role for which the alarm is generated is RegionServer, go to **Step 3**.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the HMaster for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory** > **HMaster Heap Memory Usage and Direct Memory Usage Statistics** and click **OK**, check whether the used heap memory of the HBase service reaches 90% of the maximum heap memory specified for HBase.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Figure 10-57 HMaster Heap Memory Usage and Direct Memory Usage Statistics

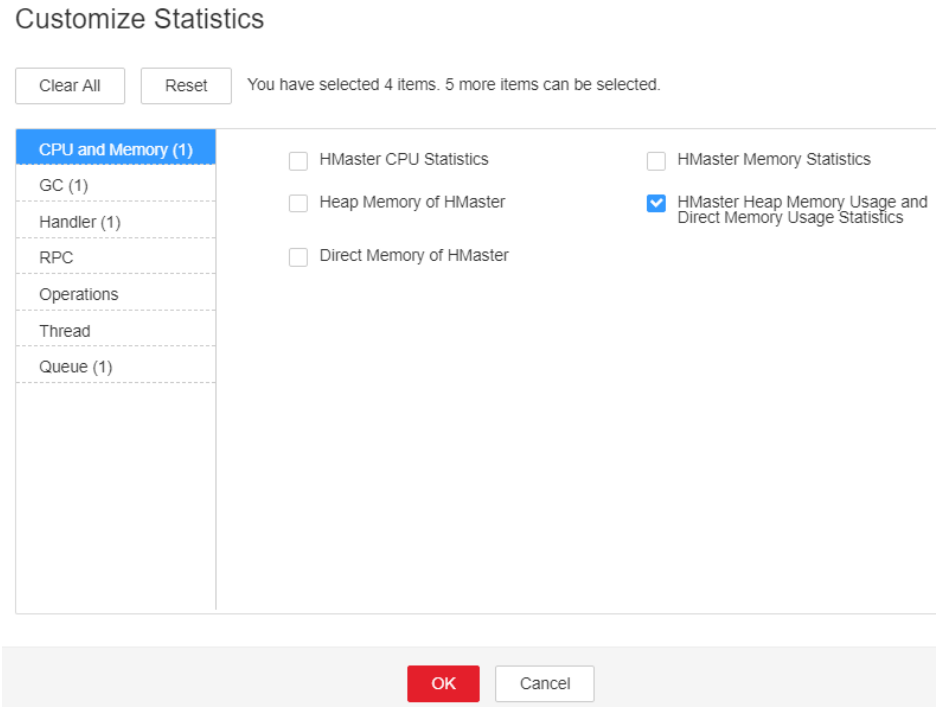

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the RegionServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory** > **RegionServer Heap Memory Usage and Direct Memory Usage Statistics** and click **OK**, check whether the used heap memory of the HBase service reaches 90% of the maximum heap memory specified for HBase.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Figure 10-58 RegionServer Heap Memory Usage and Direct Memory Usage Statistics

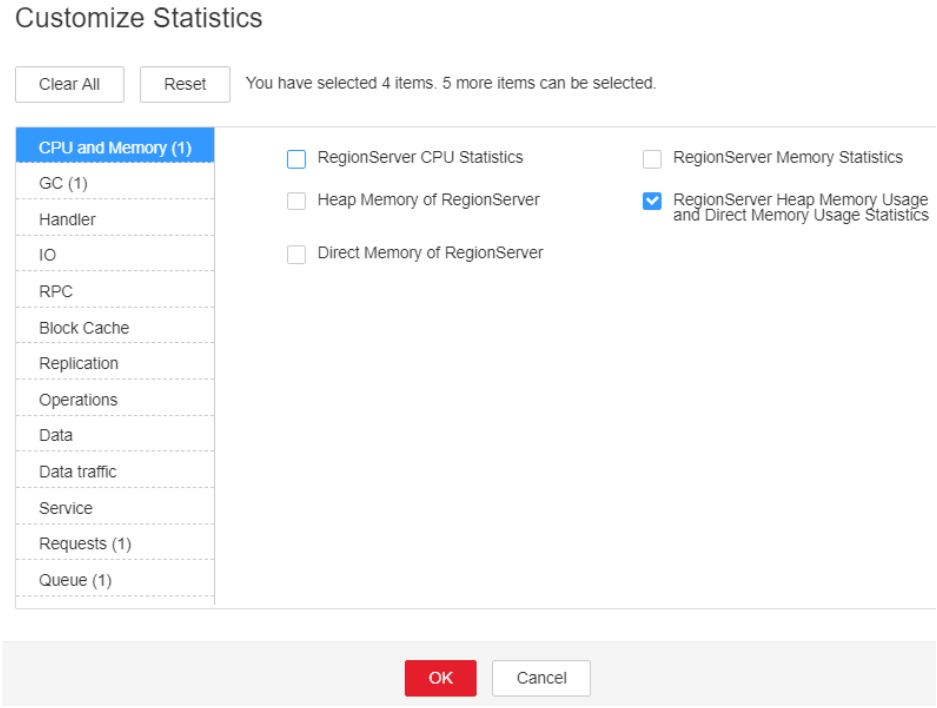

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. Choose **HMaster/RegionServer** > **System**. Increase the value of **-Xmx** in **GC_OPTS** by referring to the Note.

NOTE

1. Suggestions on GC parameter configurations for HMaster
   - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
   - Set **-XX:NewSize** to the value of **-XX:MaxNewSize**, which is one eighth of **-Xmx**.
   - For large-scale HBase clusters with a large number of regions, increase values of **GC_OPTS** parameters for HMaster. Specifically, set **-Xmx** to 4 GB if the number of regions is less than 100,000. If the number of regions is more than 100,000, set -Xmx to be greater than or equal to 6 GB. For each increased 35,000 regions, increase the value of **-Xmx** by 2 GB. The maximum value of **-Xmx** is 32 GB.
2. Suggestions on GC parameter configurations for RegionServer
   - Set **-Xms** and **-Xmx** to the same value to prevent JVM from dynamically adjusting the heap memory size and affecting performance.
   - Set **-XX:NewSize** to one eighth of **-Xmx**.
   - Set the memory for RegionServer to be greater than that for HMaster. If sufficient memory is available, increase the heap memory.
   - Set **-Xmx** based on the machine memory size. Specifically, set **-Xmx** to 32 GB if the machine memory is greater than 200 GB, to 16 GB if the machine memory is greater than 128 GB and less than 200 GB, and to 8 GB if the machine memory is less than 128 GB. When **-Xmx** is set to 32 GB, a RegionServer node supports 2000 regions and 200 hotspot regions.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **HBase** in the required cluster from the **Service** drop-down list.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.171 ALM-19009 Direct Memory Usage of the HBase Process Exceeds the Threshold

Description

The system checks the HBase service status every 30 seconds. The alarm is generated when the direct memory usage of an HBase service exceeds the threshold (90% of the maximum memory).

The alarm is cleared when the direct memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
19009	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

Impact on the System

If the available HBase direct memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the HBase service is overused or the direct memory is inappropriately allocated.

Procedure

Check direct memory usage.

Step 1 On the FusionInsight Manager portal, click **O&M** > **Alarm** > **Alarms** and select the alarm whose **ID** is **19009**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

- If the role for which the alarm is generated is HMaster, go to **Step 2**.
- If the role for which the alarm is generated is RegionServer, go to **Step 3**.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the HMaster for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory** > **HMaster Heap Memory Usage and Direct Memory Usage Statistics** and click **OK** to check whether the used direct memory of the HBase service reaches 90% of the maximum direct memory specified for HBase.

- If yes, go to **Step 4**.
- If no, go to **Step 8**.

Figure 10-59 HMaster Heap Memory Usage and Direct Memory Usage Statistics

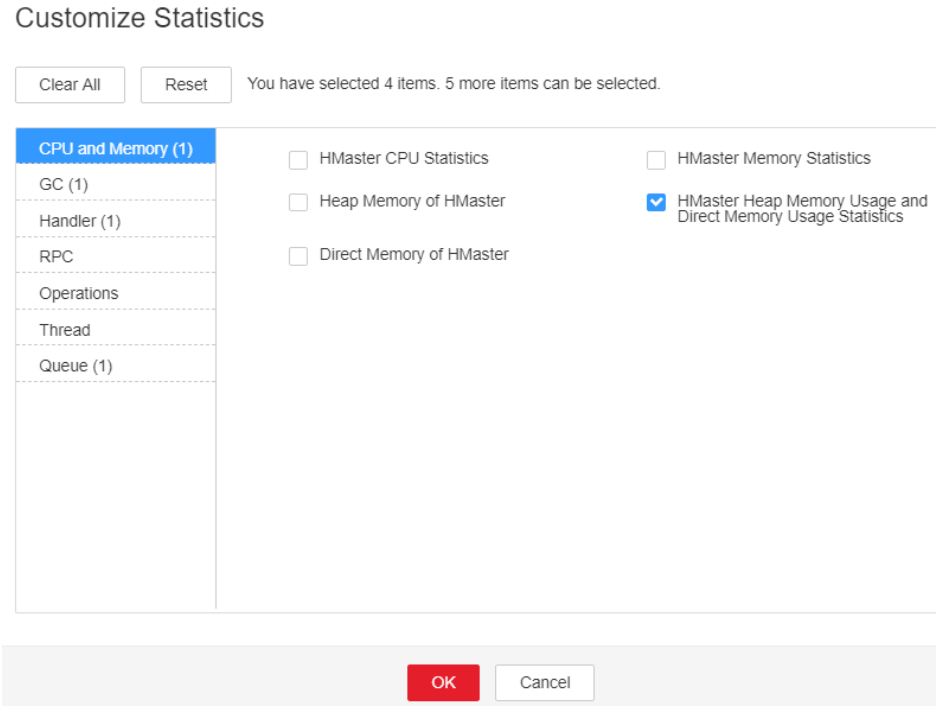

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Instance** and click the RegionServer for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the **Chart** area and choose **Customize** > **CPU and Memory** > **RegionServer Heap Memory Usage and Direct Memory Usage Statistics** and click **OK** to check whether the used direct memory of the HBase service reaches 90% of the maximum direct memory specified for HBase.

- If yes, go to **Step 4**.
- If no, go to **Step 8**.

Figure 10-60 RegionServer Heap Memory Usage and Direct Memory Usage Statistics

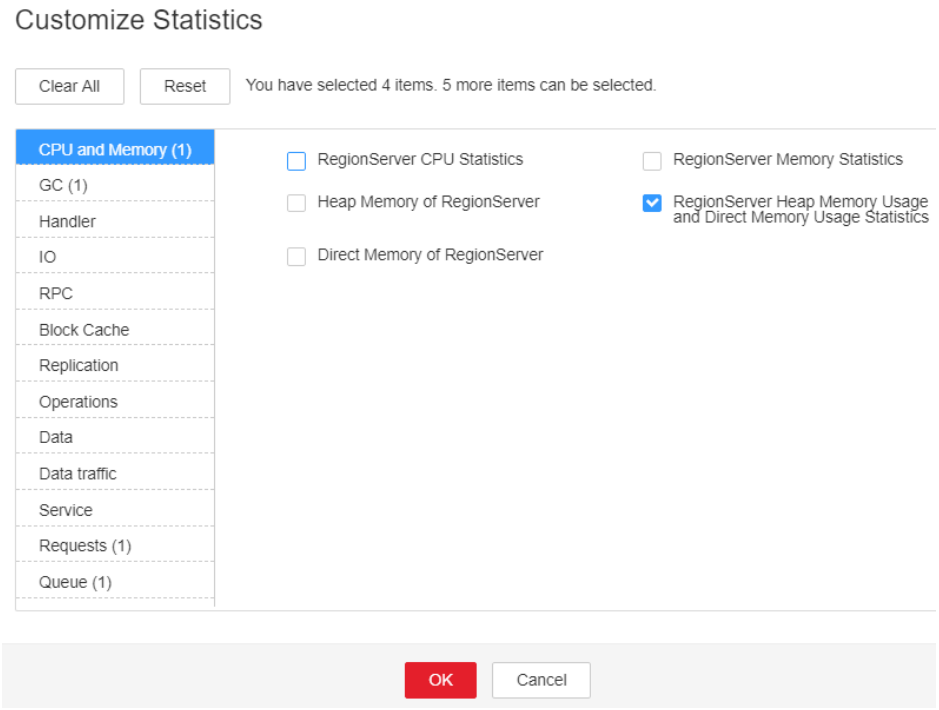

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. Choose **HMaster/RegionServer** > **System** and check whether **XX:MaxDirectMemorySize** exists in **GC_OPTS**.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On the FusionInsight Manager portal, choose **Cluster** > *Nameof the desired cluster* > **Services** > **HBase** > **Configurations**, and click **All Configurations**. Choose **HMaster/RegionServer** > **System** and delete **XX:MaxDirectMemorySize** from **GC_OPTS**.

Step 6 Check whether the **ALM-19008 Heap Memory Usage of the HBase Process Exceeds the Threshold** alarm is generated.

If yes, handle the alarm by referring to **ALM-19008 Heap Memory Usage of the HBase Process Exceeds the Threshold**.

If no, go to **Step 8**.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager interface of active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 9 In the **Service** in the required cluster drop-down list box, select **HBase**.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.172 ALM-19011 RegionServer Region Number Exceeds the Threshold

Description

The system checks the number of regions on each RegionServer in each HBase service instance every 30 seconds. The region number is displayed on the HBase service monitoring page and RegionServer role monitoring page. This alarm is generated when the number of regions on a RegionServer exceeds the threshold (default value: 2000) for 20 consecutive times. The threshold can be changed by choosing **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **HBase**. This alarm is cleared when the number of regions is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
19011	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The data read/write performance of HBase is affected when the number of regions on a RegionServer exceeds the threshold.

Possible Causes

- The RegionServer region distribution is unbalanced.
- The HBase cluster scale is too small.

Procedure

View alarm location information.

Step 1 On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Alarms**, select this alarm, and view the service instance and host name in **Location**.

Step 2 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services**, click the HBase service instance for which the alarm is generated, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, check whether the region distribution on the RegionServer is balanced.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

- If yes, go to **Step 9**.
- If no, go to **Step 3**.

Figure 10-61 WebUI of HBase instance

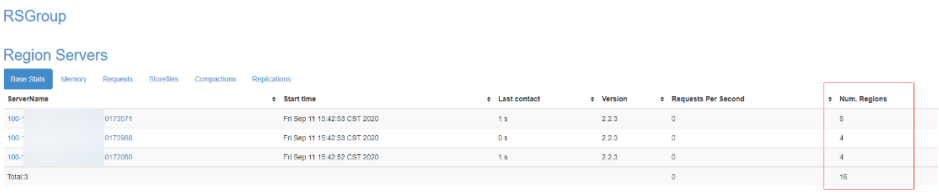

Enable load balancing.

Step 3 Log in to the node where the HBase client is located as user **root**. Go to the client installation directory, and set environment variables.

cd *client installation directory*

source bigdata_env

If the cluster adopts the security mode, perform security authentication. Specifically, run the **kinit hbase** command and enter the password as prompted (obtain the password from the administrator).

Step 4 Run the following commands to go to the HBase shell command window and check whether the load balancing function is enabled.

hbase shell

balancer_enabled

- If yes, go to **Step 6**.
- If no, go to **Step 5**.

Step 5 On the HBase shell command window, run the following commands to enable the load balancing function and check whether the function is enabled.

balance_switch true

balancer_enabled

Step 6 On the HBase shell command window, run the **balancer** command to manually trigger the load balancing function.

NOTE

You are advised to enable and manually trigger the load balancing function during off-peak hours.

Step 7 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, refresh the page and check whether the region distribution is balanced.

- If yes, go to **Step 8**.
- If no, go to **Step 21**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Delete unwanted HBase tables.

NOTE

Exercise caution when deleting data to ensure data is deleted correctly.

Step 9 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, view tables stored in the HBase service instance and record unwanted tables that can be deleted.

Step 10 On the HBase shell command window, run the **disable** command and **drop** command to delete the table to decrease the number of regions.

disable '*name of the table to be deleted***'**

drop '*name of the table to be deleted***'**

Step 11 On the HBase shell command window, run the following command to check whether the load balancing function is enabled.

balancer_enabled

- If yes, go to **Step 13**.
- If no, go to **Step 12**.

Step 12 On the HBase shell command window, run the following commands to enable the load balancing function and confirm that the function is enabled.

balance_switch true

balancer_enabled

Step 13 On the HBase shell command window, run the **balancer** command to manually trigger the load balancing function.

Step 14 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, refresh the page and check whether the region distribution is balanced.

- If yes, go to **Step 15**.
- If no, go to **Step 21**.

Step 15 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Adjust the threshold.

Step 16 On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **HBase** > **Regions(RegionServer)**, select the applied rule, and click **Modify** to check whether the threshold is proper.

- If it is excessively small, increase the threshold as required and go to **Step 17**.
- If it is proper, go to **Step 18**.

Figure 10-62 Regions(RegionServer_*1*)

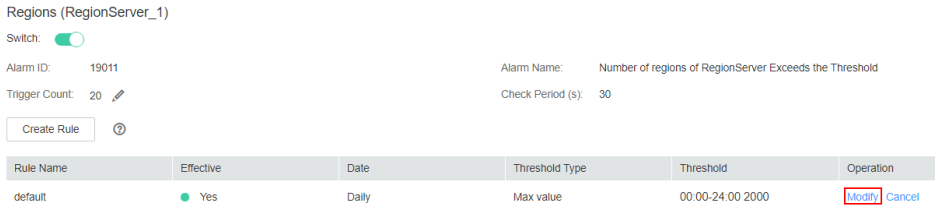

Step 17 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 18**.

Perform system capacity expansion.

Step 18 Add nodes to the HBase cluster and add RegionServer instances to the nodes. Then enable and manually trigger the load balancing function.

Step 19 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services**, click the HBase service instance for which the alarm is generated, and click **HMaster(Active)**. On the displayed WebUI of the HBase instance, refresh the page and check whether the region distribution is balanced.

- If yes, go to **Step 20**.
- If no, go to **Step 21**.

Step 20 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 21**.

Collect fault information.

Step 21 On the FusionInsight Manager home page of the active and standby clusters, choose **O&M**> **Log** > **Download**.

Step 22 Select **HBase** in the required cluster from the **Service**.

Step 23 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 24 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.173 ALM-19012 HBase System Table Directory or File Lost

Description

The system checks whether HBase directories and files exist on the HDFS every 120 seconds. This alarm is generated when the system detects that the files or directories do not exist. This alarm is cleared when the files or directories are restored.

The HBase directories and files are as follows:

- Directory of the namespace **hbase** on the HDFS
- **hbase.version** file
- Directory of the table **hbase:meta** on the HDFS, .tableinfo file, and .regioninfo file
- Directory of the table **hbase:namespace** on the HDFS, .tableinfo file, and .regioninfo file

- Directory of the table **hbase:hindex** on the HDFS, .tableinfo file, and .regioninfo file
- Directory of the **hbase:acl** table on the HDFS, .tableinfo, and .regioninfo file (This table does not exist in the common mode cluster by default.)

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
19012	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The HBase service fails to restart or start.

Possible Causes

Files or directories on the HDFS are missing.

Procedure

Locate the alarm cause.

Step 1 On the FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Click this alarm and check whether **Alarm Cause** indicates unknown errors.

- If yes, go to **Step 4**.
- If no, go to **Step 2**

Step 2 On the FusionInsight Manager home page, choose **O&M** > **Backup and Restoration** > **Backup Management**. Check whether there are success records of the backup task named **default** or other HBase metadata backup tasks that have been successfully executed.

- If yes, go to **Step 3**.

- If no, go to **Step 4**.

Step 3 Use the latest backup metadata to restore the metadata of the HBase service.

Collect fault information.

Step 4 On the FusionInsight Manager page of the active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 5 In the **Service** area, select faulty HBase services in the required cluster.

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.174 ALM-19013 Duration of Regions in transaction State Exceeds the Threshold

Description

The system checks the number of regions in transaction state on HBase every 300 seconds. This alarm is generated when the system detects that the duration of regions in transaction state exceeds the threshold for two consecutive times. This alarm is cleared when all timeout regions are restored.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
19013	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Some data in the table gets lost or becomes unavailable.

Possible Causes

- Compaction is permanently blocked.
- The HDFS files are abnormal.

Procedure

Locate the alarm cause.

Step 1 On the FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, select this alarm, and view the **HostName** and **RoleName** in **Location**.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services > HBase**, Click the drop-down menu in the chartarea and choose **Customize > Service** >

Region in transaction count to view **Region in transaction count over threshold**. Check whether the monitoring item detects a value in three consecutive detection periods. (The default threshold is 60 seconds.)

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Figure 10-63 Region in transaction count

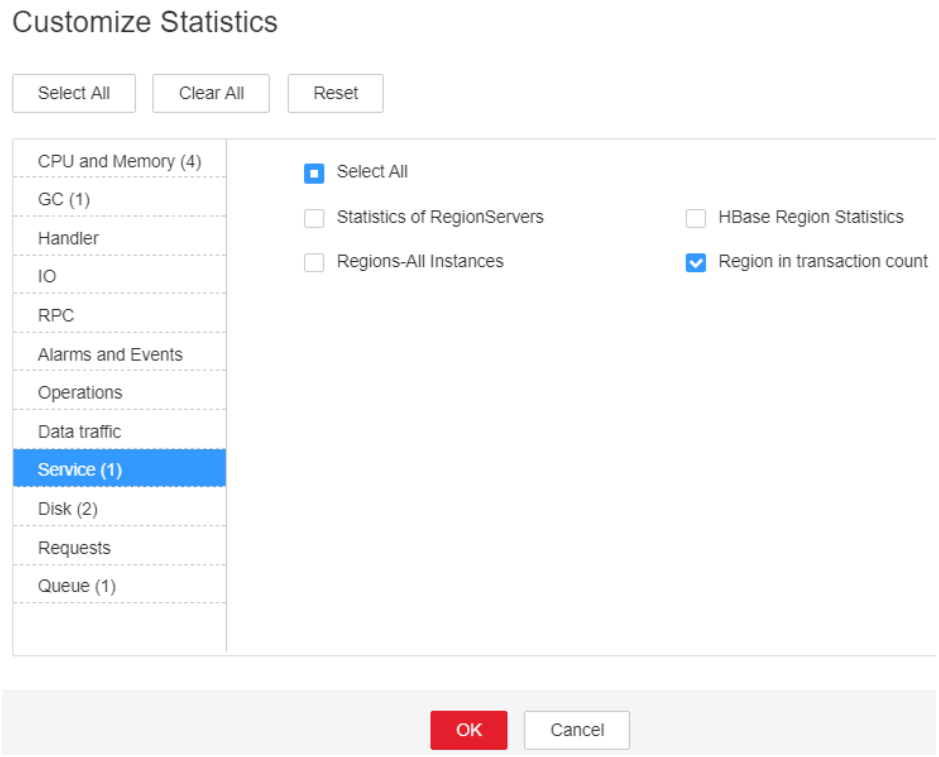

Step 3 Choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase** > **HMaster (Active)** > **Tables** to check whether the regions of only one table transaction status time out.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Figure 10-64 Tables

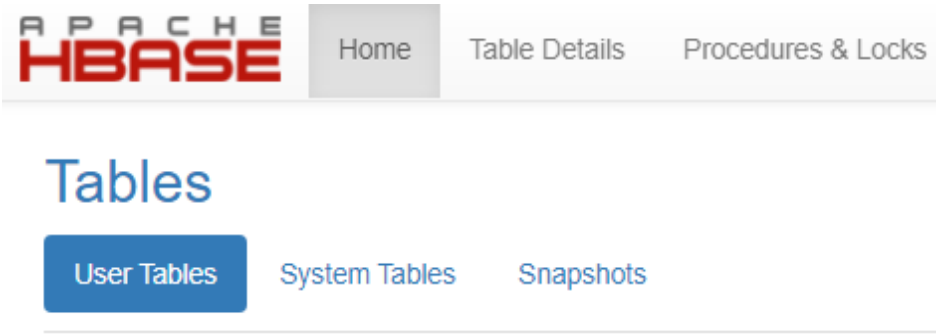

Step 4 Run the **hbase hbck** command on the client and check whether the error message "No table descriptor file under hdfs://hacluster/hbase/data/default/table" is displayed.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Log in to the client as user **root**. Run the following command:

cd *client installation directory*

source bigdata_env

If the cluster is in security mode, run the **kinit hbase** command

Log in to the HMaster WebUI, choose **Procedure & Locks** in the navigation tree, and check whether any process ID is in the **Waiting** state in **Procedures**. If yes, run the following command to release the procedure lock:

hbase hbck -j *client installation directory*/**HBase/hbase/tools/hbase-hbck2-*.jar bypass -o** *pid*

Check whether the state is in the **Bypass** state. If the procedure on the UI is always in **RUNNABLE(Bypass)** state, perform an active/standby switchover. Run the **assigns** command to bring the region online again.

hbase hbck -j *client installation directory*/**HBase/hbase/tools/hbase-hbck2-*.jar assigns -o** *regionName*

Step 6 Repeat **Step 4**. Run the **hbase hbck** command on the client and check whether the error message "No table descriptor file under hdfs://hacluster/hbase/data/default/table" is displayed.

- If yes, go to **Step 7**.
- If no, no further action is required.

Collect fault information.

Step 7 On the FusionInsight Manager page of the active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 8 In the **Service** area, select faulty HBase services in the required cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.175 ALM-19014 Capacity Quota Usage on ZooKeeper Exceeds the Threshold Severely

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the ZNode capacity usage of the HBase service exceeds the critical alarm threshold (90% by default).

This alarm is cleared when the ZNode capacity usage is less than the critical alarm threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19014	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Threshold	Specifies the threshold for generating the alarm.

Impact on the System

This alarm indicates that the capacity usage of the ZNode of HBase has exceeded the threshold severely. As a result, the write request of the HBase service fails.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the capacity configuration and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, select the alarm whose ID is **19014**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

cd *Client installation directory*

Run the following command to set environment variables:

source bigdata_env

If the cluster uses the security mode, run the following command to perform security authentication:

kinit hbase

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 3 Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode capacity quota of the HBase service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the capacity configuration of the root ZNode of the HBase service.

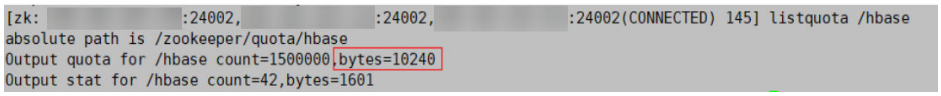

Step 4 Run the **getusage /hbase/splitWAL** command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether the alarm whose ID is **12007**, **19000**, or **19013** and the **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.

Step 6 Run the **getusage /hbase/replication** command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether the alarm is cleared five minutes later.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **HBase** for the target cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.176 ALM-19015 Quantity Quota Usage on ZooKeeper Exceeds the Threshold

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the system detects that the ZNode quantity usage of the HBase service exceeds the alarm threshold (75% by default).

This alarm is cleared when the ZNode quantity usage is less than the alarm threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19015	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Threshold	Specifies the threshold for generating the alarm.

Impact on the System

This alarm indicates that the ZNode quantity usage in the HBase service has exceeded the threshold. If this alarm is not handled in a timely manner, the problem severity may be escalated to **Critical**, affecting data writing.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the quantity quota and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, select the alarm whose ID is **19015**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

cd *Client installation directory*

Run the following command to set environment variables:

source bigdata_env

If the cluster uses the security mode, run the following command to perform security authentication:

kinit hbase

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 3 Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode quantity quota of the HBase service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the quantity quota configuration of the root ZNode of the HBase service.

```
[zk:                :24002,                :24002,                :24002(CONNECTED) 7] listquota /hbase
absolute path is /zookeeper/quota/hbase
Output quota for /hbase count=1500000,bytes=10240
Output stat for /hbase count=59,bytes=1902
```


Step 4 Run the **getusage /hbase/splitWAL** command to check the ZNode quantity usage and check whether the ratio of **Node count** in the command output to the ZNode quantity quota is close to the alarm threshold.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether the alarm whose ID is **12007**, **19000**, or **19013** and the **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.

Step 6 Run the **getusage /hbase/replication** command to check the ZNode quantity usage and check whether the ratio of **Node count** in the command output to the ZNode quantity quota is close to the alarm threshold.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.177 ALM-19016 Quantity Quota Usage on ZooKeeper Exceeds the Threshold Severely

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the znode usage of the HBase service exceeds the critical alarm threshold (90% by default).

This alarm is cleared when the quantity usage of the ZNode is less than the critical alarm threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19016	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Threshold	Specifies the threshold for generating the alarm.

Impact on the System

This alarm indicates that the quantity usage of the ZNode of HBase has exceeded the threshold severely. As a result, the write request of the HBase service fails.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the quantity quota and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, select the alarm whose ID is **19016**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

cd *Client installation directory*

Run the following command to set environment variables:

source bigdata_env

If the cluster uses the security mode, run the following command to perform security authentication:

kinit hbase

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 3 Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode quantity quota of the HBase service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the quantity configuration of the root ZNode of the HBase service.

```
[zk:                :24002,                :24002,                :24002(CONNECTED) 7] listquota /hbase
absolute path is /zookeeper/quota/hbase
Output quota for /hbase count=1500000,bytes=10240
Output stat for /hbase count=59,bytes=1902
```

Step 4 Run the **getusage /hbase/splitWAL** command to check the ZNode usage and check whether the ratio of **Node count** in the command output to the znode quantity quota is close to the alarm threshold.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether the alarm whose ID is **12007**, **19000**, or **19013** and the **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.

Step 6 Run the **getusage /hbase/replication** command to check the ZNode usage and check whether the ratio of **Node count** in the command output to the ZNode quantity quota is close to the alarm threshold.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.178 ALM-19017 Capacity Quota Usage on ZooKeeper Exceeds the Threshold

Alarm Description

The system checks the ZNode usage of the HBase service every 120 seconds. This alarm is generated when the system detects that the ZNodes capacity usage of the HBase service exceeds the alarm threshold (75% by default).

This alarm is cleared when the capacity usage of the ZNode capacity is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19017	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Threshold	Specifies the threshold for generating the alarm.

Impact on the System

This alarm indicates that the ZNodes capacity usage in the HBase service has exceeded the threshold. If this alarm is not handled in a timely manner, the problem severity may be escalated to **Critical**, affecting data writing.

Possible Causes

- DR is configured for HBase, and data synchronization fails or is slow in DR.
- A large number of WAL files are being split in the HBase cluster.

Handling Procedure

Check the capacity configuration and usage of ZNodes.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, select the alarm whose ID is **19017**, and view the threshold in **Additional Information**.

Step 2 Log in to the HBase client as user **root**. Run the following command to go to the client installation directory:

cd *Client installation directory*

Run the following command to set environment variables:

source bigdata_env

If the cluster uses the security mode, run the following command to perform security authentication:

kinit hbase

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 3 Run the **hbase zkcli** command to log in to the ZooKeeper client and run the **listquota /hbase** command to check the ZNode quantity quota of the HBase

service. The ZNode root directory in the command is specified by the **zookeeper.znode.parent** parameter of the HBase service. The marked area in the following figure shows the quantity configuration of the root ZNode of the HBase service.

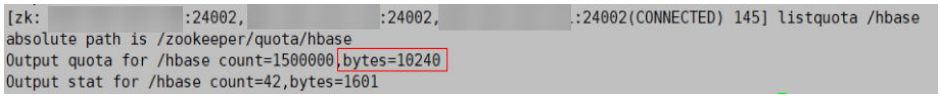

Step 4 Run the **getusage /hbase/splitWAL** command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 On FusionInsight Manager, check whether the alarm whose ID is **12007**, **19000**, or **19013** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 7**.

Step 6 Run the **getusage /hbase/replication** command to check the capacity usage of the ZNode. Check whether the ratio of **Data size** to the ZNode capacity quota is close to the alarm threshold.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether the alarm whose ID is **19006** and **ServiceName** in **Location** is the current HBase service exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.179 ALM-19018 HBase Compaction Queue Size Exceeds the Threshold

Alarm Description

The system checks the HBase compaction queue size every 30 seconds. This alarm is generated when the compaction queue size exceeds the alarm threshold (**100** by default) for three consecutive times. This alarm is cleared when the compaction queue size is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19018	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The cluster performance may deteriorate, affecting data read and write.

Possible Causes

- The number of HBase RegionServers is too small.
- There are excessive regions on a single RegionServer of HBase.

- The HBase RegionServer heap size is small.
- Resources are insufficient.
- Related parameters are not configured properly.

Handling Procedure

Check whether related parameters are properly configured.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, check whether the alarm whose **Alarm ID** is **19008** or **19011** exists.

- If yes, click **View Help** next to the alarm and rectify the fault by referring to the help document. Then, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **HBase**. On the page that is displayed, click the **Configurations** tab then the **All Configurations** sub-tab, search for **hbase.hstore.compaction.min**, **hbase.hstore.compaction.max**, **hbase.regionserver.thread.compaction.small**, and **hbase.regionserver.thread.compaction.throttle**, and set them to larger values.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect the fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.180 ALM-19019 Number of HBase HFiles to Be Synchronized Exceeds the Threshold

Alarm Description

The system checks the number of HFiles to be synchronized by the RegionServer of each HBase service instance every 30 seconds. This indicator can be viewed on the RegionServer role monitoring page. This alarm is generated when the number of HFiles to be synchronized on a RegionServer exceeds the threshold (exceeding 128 for 20 consecutive times by default). To change the threshold, choose **O&M** > **Alarm** > **Threshold Configuration** > *Name of the desired cluster* > **HBase** . This alarm is cleared when the number of HFiles to be synchronized is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19019	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the number of HFiles to be synchronized by a RegionServer exceeds the threshold, the number of ZNodes used by HBase exceeds the threshold, affecting the HBase service status.

Possible Causes

- The network is abnormal.
- The RegionServer region distribution is unbalanced.
- The HBase service scale of the standby cluster is too small.

Handling Procedure

View alarm location information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19019**, and view the service instance and host name in **Location**.

Check the network connection between RegionServers on active and standby clusters.

Step 2 Run the **ping** command to check whether the network connection between the faulty RegionServer node and the host where RegionServer of the standby cluster resides is normal.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Contact the network administrator to restore the network.

Step 4 After the network recovers, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the RegionServer region distribution in the active cluster.

Step 5 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance and check whether regions are evenly distributed on the Region Server.

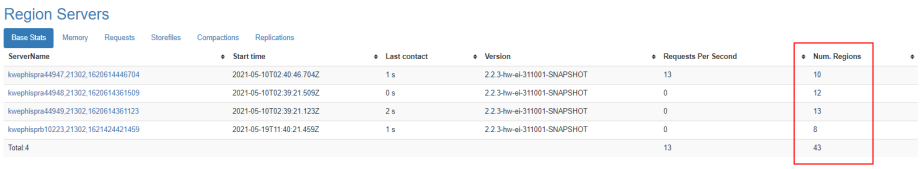

Step 6 Log in to the faulty RegionServer node as user **omm**.

Step 7 Run the following commands to go to the client installation directory and set the environment variable:

cd *Client installation directory*

source bigdata_env

If the cluster uses the security mode, perform security authentication. Run the **kinit hbase** command and enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 8 Run the following commands to check whether the load balancing function is enabled.

hbase shell

balancer_enabled

- If yes, go to **Step 10**.
- If no, go to **Step 9**.

Step 9 Run the following commands in HBase Shell to enable the load balancing function and check whether the function is enabled.

balance_switch true

balancer_enabled

Step 10 Run the **balancer** command to manually trigger the load balancing function.

NOTE

You are advised to enable and manually trigger the load balancing function during off-peak hours.

Step 11 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check the HBase service scale of the standby cluster.

Step 12 Expand the HBase cluster, add a node, and add a RegionServer instance on the node. Then, perform **Step 6** to **Step 10** to enable the load balancing function and manually trigger it.

Step 13 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance, refresh the page, and check whether regions are evenly distributed.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Collect the fault information.

Step 15 On FusionInsight Manager of the standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 16 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 17 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.181 ALM-19020 Number of HBase WAL Files to Be Synchronized Exceeds the Threshold

Alarm Description

The system checks the number of WAL files to be synchronized by the RegionServer of each HBase service instance every 30 seconds. This indicator can be viewed on the RegionServer role monitoring page. This alarm is generated when the number of WAL files to be synchronized on a RegionServer exceeds the threshold (exceeding 128 for 20 consecutive times by default). To change the threshold, choose **O&M** > **Alarm** > **Threshold Configuration** > *Name of the desired cluster* > **HBase** . This alarm is cleared when the number of WAL files to be synchronized is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19020	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the number of WAL files to be synchronized by a RegionServer exceeds the threshold, the number of ZNodes used by HBase exceeds the threshold, affecting the HBase service status.

Possible Causes

- The network is abnormal.
- The RegionServer region distribution is unbalanced.
- The HBase service scale of the standby cluster is too small.

Handling Procedure

View alarm location information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19020**, and view the service instance and host name in **Location**.

Check the network connection between RegionServers on active and standby clusters.

Step 2 Run the **ping** command to check whether the network connection between the faulty RegionServer node and the host where RegionServer of the standby cluster resides is normal.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Contact the network administrator to restore the network.

Step 4 After the network recovers, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the RegionServer region distribution in the active cluster.

Step 5 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance and check whether regions are evenly distributed on the Region Server.

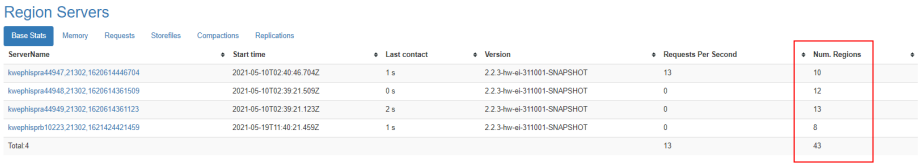

Step 6 Log in to the faulty RegionServer node as user **omm**.

Step 7 Run the following commands to go to the client installation directory and set the environment variable:

cd *Client installation directory*

source bigdata_env

If the cluster uses the security mode, perform security authentication. Run the **kinit hbase** command and enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 8 Run the following commands to check whether the load balancing function is enabled.

hbase shell

balancer_enabled

- If yes, go to **Step 10**.
- If no, go to **Step 9**.

Step 9 Run the following commands in HBase Shell to enable the load balancing function and check whether the function is enabled.

balance_switch true

balancer_enabled

Step 10 Run the **balancer** command to manually trigger the load balancing function.

NOTE

You are advised to enable and manually trigger the load balancing function during off-peak hours.

Step 11 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check the HBase service scale of the standby cluster.

Step 12 Expand the HBase cluster, add a node, and add a RegionServer instance on the node. Then, perform **Step 6** to **Step 10** to enable the load balancing function and manually trigger it.

Step 13 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HBase**. Click **HMaster(Active)** to go to the web UI of the HBase instance, refresh the page, and check whether regions are evenly distributed.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Collect the fault information.

Step 15 On FusionInsight Manager of the standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 16 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 17 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.182 ALM-19021 Handler Usage of RegionServer Exceeds the Threshold

Alarm Description

The system checks the RegionServer handler usage of each HBase service instance every 30 seconds. This alarm is generated when the handler usage of a RegionServer exceeds the threshold (90% for five consecutive times by default). This alarm is cleared if the handler usage is lower than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19021	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

RegionServers or HBase cannot provide services properly.

Possible Causes

- The value of a handler is too small.
- Hotspotting occurs.

Handling Procedure

View alarm location information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**, locate the row that contains the alarm whose **Alarm ID** is **19021**, and view the service instance and host name in **Location**.

Check the handler configuration.

Step 2 Choose **Cluster** > **Services** > **HBase** and click the **Configurations** tab. In the upper right corner of the page, search for **hbase.regionserver.handler.count** and check whether its value is too small. The default value is **200**.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Change the value of this parameter to a larger value and save the configuration. Choose **Cluster** > **Services** > **HBase**, click the **Instance** tab, select the affected RegionServer instances, and choose **More** > **Instance Rolling Restart**. In the displayed dialog box, enter the username and password. In the **Instance Rolling Restart** dialog box, click **OK** and wait until the rolling restart is complete.

Step 4 After the configuration takes effect, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether cluster hotspotting occurs.

Step 5 On FusionInsight Manager, choose **Cluster** > **Services** > **HBase** and click **HMaster(Active)** following **HMaster WebUI** to go to the web UI of the HBase instance. In the **Region Servers** area of the **Home** page, click **Requests** and check whether the requests in the **Filtered Read Request Count** and **Write Request Count** columns are evenly distributed.

- If yes, go to **Step 13**.
- If no, go to **Step 6**.

Step 6 Check whether regions are evenly distributed.

On FusionInsight Manager, choose **Cluster** > **Services** > **HBase** and click **HMaster(Active)** following **HMaster WebUI** to go to the web UI of the HBase

instance. In the **Region Servers** area of the **Home** page, click **Base Stats** and check whether the regions in the **Num.Regions** column are evenly distributed.

- If yes, go to **Step 13**.
- If no, go to **Step 7**.

Step 7 Log in to the faulty RegionServer node as user **omm**.

Step 8 Run the following commands to go to the client installation directory and set the environment variable:

cd *Client installation directory*

source bigdata_env

If the cluster uses the security mode, run the following command to perform security authentication:

kinit hbase

Enter the password as prompted (obtain the password from the MRS cluster administrator).

Step 9 Run the following commands to check whether the load balancing function is enabled. If the command output is **true**, the load balancing function is enabled.

hbase shell

balancer_enabled

```
hbase:004:0> balancer_enabled
true
Took 0.0165 seconds
=> true
```

- If yes, go to **Step 13**.
- If no, go to **Step 10**.

Step 10 Run the following commands in HBase Shell to enable the load balancing function and check whether the function is enabled.

balance_switch true

balancer_enabled

NOTE

You are advised to enable and manually trigger the load balancing function during off-peak hours.

Step 11 Run the following command to manually trigger' the load balancing function:

balancer

Step 12 After the load balancing is complete, log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collect the fault information.

Step 13 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 14 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 16 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.183 ALM-19022 HBase Hotspot Detection Is Unavailable

Alarm Description

When the MetricController instance is installed for HBase, the alarm module checks the health status of the active HBase MetricController instance every 120 seconds. This alarm is generated when the active HBase MetricController instance does not exist or is unavailable and the hotspot detection function is unavailable.

This alarm is cleared when the active HBase MetricController instance recovers.

NOTE

This alarm applies only to MRS 3.3.0 or later.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19022	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The HBase hotspot detection function is unavailable.

Possible Causes

- The ZooKeeper service is abnormal.
- The HBase service is abnormal.
- In the current HBase service, the MetricController instance on the same node as the active HMaster instance is not started.
- The network is abnormal.

Handling Procedure

Check the ZooKeeper service status.

Step 1 In the service list on FusionInsight Manager, check whether **Running Status** of ZooKeeper is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 2**.

Step 2 In the alarm list, check whether **ALM-13000 ZooKeeper Service Unavailable** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by performing the operations provided for **ALM-13000 ZooKeeper Service Unavailable**.

Step 4 Wait for several minutes and check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the HBase service status.

Step 5 In the service list on FusionInsight Manager, check whether **Running Status** of HBase is **Normal**.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

Step 6 In the alarm list, check whether the alarm ALM-19000 HBase Service Unavailable exists.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Rectify the fault by following the steps provided for **ALM-19000 HBase Service Unavailable**.

Step 8 Wait for several minutes and check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether the MetricController instance deployed on the same node as the active HMaster instance is started.

Step 9 On FusionInsight Manager, choose **Cluster** > **Service** > **HBase**, and click **Instances** to check whether the **MetricController(Active)** instance exists.

- If yes, go to **Step 12**.
- If no, go to **Step 10**.

Step 10 Select the MetricController instance whose management IP address is the same as that of the active HMaster instance, and click **Start Instance**.

Step 11 After the MetricController instance is restarted, check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check the network connectivity between the started MetricController instances and the active HMaster node.

Step 12 Log in to the node where the active HMaser instance is deployed and run **ping** *IP address of the node where the standby MetricController instance is deployed* to check whether the network connection between the started MetricController instances and the host where the active HMaster instance is deployed is normal.

- If yes, go to **Step 15**.
- If no, go to **Step 13**.

Step 13 Contact the network administrator to restore the network.

Step 14 After the network recovers, check whether the alarm **HBase Hotspot Detection Is Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Collect fault information.

Step 15 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 16 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 17 In the **Host** area, select the host where the HMaster instance is deployed.

Step 18 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.184 ALM-19023 Region Traffic Restriction for HBase

Alarm Description

When the MetricController instance is installed for the HBase service, self-healing from hotspotting is automatically enabled. The alarm module checks whether there are regions whose request traffic is restricted due to hotspot issues in HBase every 120 seconds. This alarm is generated when the region where hotspot traffic is restricted is detected in HBase.

This alarm is cleared when the region is no longer a hotspot.

This alarm applies only to MRS 3.3.0 or later.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19023	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the traffic of a hotspot region is restricted, the number of handlers for processing the requests in the region is limited. As a result, services requesting the region may slow down or retry upon failure.

Possible Causes

Too many requests are directed to a single region when the HBase service is accessed.

Handling Procedure

Check whether there are too many requests in a single region of HBase.

Step 1 Log in to FusionInsight Manager, and Choose **O&M** > **Alarm** > **Alarms**.

Step 2 In **Additional Information** of **Region Traffic Restriction for HBase**, view the reported table name and region information.

Step 3 On FusionInsight Manager, choose **Cluster** > **Service** > **HBase** and click the hyperlink on the right of HMaster web UI.

Step 4 Click **Table Details** and adjust service configurations in the region where the table in **Step 2** is deployed.

Step 5 Wait a moment and then check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 8 In the **Host** area, select the host where the HMaster instance is deployed.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm will be automatically cleared.

Related Information

None.

10.185 ALM-19024 RPC Requests P99 Latency on RegionServer Exceeds the Threshold

Alarm Description

The system checks P99 latency for RPC requests on each RegionServer instance of the HBase service every 30 seconds. This alarm is generated when P99 latency for RPC requests on a RegionServer exceeds the threshold for 10 consecutive times.

This alarm is cleared when P99 latency for RPC requests on a RegionServer instance is less than or equal to the threshold.

This alarm applies only to MRS 3.3.0 or later.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19024	• **Critical**: The default threshold is 10 seconds. • **Major**: The default threshold is 5 seconds.	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If RPC requests P99 latency exceeds the threshold, the RegionServer cannot deliver normal service performance externally. If RPC requests P99 latency on most RegionServers in the cluster exceeds the threshold, HBase may fail to provide services for external systems.

Possible Causes

- RegionServer GC duration is too long.
- The HDFS RPC response is too slow.
- RegionServer request concurrency is too high.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19024**, and view the service instance and host name in **Location**.

Check the GC duration of RegionServer.

Step 2 In the alarm list on FusionInsight Manager, check whether the "HBase GC Duration Exceeds the Threshold" alarm is generated for the service instance in **Step 1**.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Rectify the fault by following the handling procedure of "ALM-19007 HBase GC Duration Exceeds the Threshold".

Step 4 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check HDFS RPC response time.

Step 5 In the alarm list on FusionInsight Manager, check whether alarm "Average NameNode RPC Processing Time Exceeds the Threshold" is generated for the HDFS service on which the HBase service depends.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 Rectify the fault by following the handling procedure of "ALM-14021 Average NameNode RPC Processing Time Exceeds the Threshold".

Step 7 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check the number of concurrent processes on a RegionServer.

Step 8 In the alarm list on FusionInsight Manager, check whether the "Handler Usage of RegionServer Exceeds the Threshold" alarm is generated for the service instance in **Step 1**.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

Step 9 Rectify the fault by following the handling procedure of "ALM-19021 Handler Usage of RegionServer Exceeds the Threshold".

Step 10 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 12 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 13 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.186 ALM-19025 Damaged StoreFile in HBase

Alarm Description

The system checks the **hdfs://hacluster/hbase/autocorrupt** and **hdfs://hacluster/hbase/MasterData/autocorrupt** directories on HDFS of each HBase service every 120 seconds. This alarm is generated when there are files in the directories.

This alarm is cleared when the **hdfs://hacluster/hbase/autocorrupt** and **hdfs://hacluster/hbase/MasterData/autocorrupt** directories do not exist or are empty.

This alarm applies only to MRS 3.3.0 or later.

NOTE

hdfs://hacluster indicates the name of the file system used by HBase, and **/hbase** indicates the root directory of HBase in the file system. You can log in to FusionInsight Manager, choose **Cluster** > **Services** > **HBase** and click **Configuration**. Search for **fs.defaultFS** and **hbase.data.rootdir**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19025	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

There are damaged StoreFile files in HBase, which may cause data loss.

Possible Causes

The StoreFile files are damaged.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19025**, and view the service in **Location**.

Step 2 Log in to the node where the HDFS and HBase clients are installed as the client installation user and run the following commands:

cd *Client installation directory*

source bigdata_env

kinit *Component service user* (If Kerberos authentication is disabled for the cluster (the cluster is in normal mode), skip this step.)

Step 3 Check the damaged StoreFile file.

- Run the following command to check whether the **/hbase/autocorrupt** directory of HDFS is empty. If it is not, go to **Step 4**.

  hdfs dfs -ls -R hdfs://hacluster/hbase/autocorrupt

- Run the following command to check whether the **/hbase/MasterData/autocorrupt** directory of HDFS is empty. If it is not, go to **Step 9**.

  hdfs dfs -ls -R hdfs://hacluster/hbase/MasterData/autocorrupt

Step 4 Run the following command to restore the StoreFile files in the **hdfs://hacluster/hbase/autocorrupt** directory:

hdfs debug recoverLease -path hdfs://hacluster/hbase/autocorrupt/*Name space*/Table/Region/Column family/*StoreFile files*

Step 5 Check whether the damaged StoreFile files are restored. If the following information is displayed, the restoration is successful:

```
recoverLease SUCCEEDED on hdfs://hacluster/hbase/autocorrupt/
default/h1/865665fe32db62dadada68b644359809/cf1/95f210f931ad44c99e4028470be7d292
```

If yes, go to **Step 6**.

If no, go to **Step 9**.

Step 6 Run the following command to move the files back to the **hdfs://hacluster/hbase/data** directory:

hdfs dfs -mv hdfs://hacluster/hbase/autocorrupt/*Name space*/Table/Region/Column family/*StoreFile files***hdfs://hacluster/hbase/data/***Name space*/Table/Region/Column family/*StoreFile files*

Step 7 Run the following command on HBase Shell to bring the region online again:

hbase shell

unassign'*Region*'

assign'*Region*'

Step 8 Wait several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.187 ALM-19026 Damaged WAL Files in HBase

Alarm Description

The system checks the **hdfs://hacluster/hbase/corrupt** directory on the HDFS of each HBase service every 120 seconds. This alarm is generated when there are WAL files in the **/hbase/corrupt** directory.

This alarm is cleared when the **/hbase/corrupt** directory does not exist or does not contain WAL files.

This alarm applies only to MRS 3.3.0 or later.

NOTE

hdfs://hacluster indicates the name of the file system used by HBase, and **/hbase** indicates the root directory of HBase in the file system. You can log in to FusionInsight Manager, choose **Cluster** > **Services** > **HBase** and click **Configuration**. Search for **fs.defaultFS** and **hbase.data.rootdir**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
19026	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

There are damaged WAL files in HBase, which may cause data loss.

Possible Causes

The WAL files are damaged.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **19026**, and view the service in **Location**.

Step 2 Log in to the node where the HDFS clients are installed as the client installation user and run the following commands:

cd *Client installation directory*

source bigdata_env

kinit *Component service user* (If Kerberos authentication is disabled for the cluster (the cluster is in normal mode), skip this step.)

Step 3 Run the following command to check the damaged WAL files and go to **Step 4**:

hdfs dfs -ls hdfs://hacluster/hbase/corrupt/*%2C*

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **HBase** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.188 ALM-20002 Hue Service Unavailable

Description

This alarm is generated when the Hue service is unavailable. The system checks the Hue service status every 60 seconds.

This alarm is cleared when the Hue service is normal.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
20002	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The system cannot provide data loading, query, and extraction services.

Possible Causes

- The internal KrbServer service on which the Hue service depends is abnormal.
- The internal DBService service on which the Hue service depends is abnormal.
- The network connection to the DBService is abnormal.

Procedure

Check whether the KrbServer is abnormal.

Step 1 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services**. In the service list, check whether the **KrbServer** running status is **Normal**.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Restart the KrbServer service.

Step 3 Wait several minutes, and check whether **Hue Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the DBService is abnormal.

Step 4 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 5 In the service list, check whether the **DBService** running status is **Normal**.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 Restart the DBService.

NOTE

To restart the service, enter the FusionInsight Manager administrator password.

Step 7 Wait several minutes, and check whether **Hue Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check whether the network connection to the DBService is normal.

Step 8 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Hue** > **Instance**, record the IP address of the active Hue.

Step 9 Log in to the active Hue.

Step 10 Run the **ping** command to check whether communication between the host that runs the active Hue and the hosts that run the DBService is normal. (Obtain the IP addresses of the hosts that run the DBService in the same way as that for obtaining the IP address of the active Hue.)

- If yes, go to **Step 13**.
- If no, go to **Step 11**.

Step 11 Contact the administrator to restore the network.

Step 12 Wait several minutes, and check whether **Hue Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collect fault information.

Step 13 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 14 Select the following nodes in the required cluster from the **Service** drop-down list:

- Hue
- Controller

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 On the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Hue**.

Step 17 Choose **More** > **Restart Service**, and click **OK**.

Step 18 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 19**.

Step 19 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.189 ALM-23001 Loader Service Unavailable

Description

The system checks the Loader service availability every 60 seconds. This alarm is generated when the system detects that the Loader service is unavailable. This alarm is cleared when the Loader service is available.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23001	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.

Impact on the System

When the Loader service is unavailable, the data loading, import, and conversion functions are unavailable.

Possible Causes

- The internal service on which the Loader service depends is abnormal.
  - The ZooKeeper service is abnormal.
  - The HDFS service is abnormal.
  - The DBService service is abnormal.
  - The Yarn service is abnormal.
  - The Mapreduce service is abnormal.
- Environment fault: The network is abnormal, which the Loader service cannot communicate with the depended internal services and cannot provide services.
- Software fault: The Loader service cannot run properly.

Procedure

Check the ZooKeeper service status.

Step 1 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **ZooKeeper** to check whether the ZooKeeper running status is **Normal**.

- If yes, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 Choose **More** > **Restart Service** to restart the ZooKeeper service. In the alarm list, check whether **LoaderService Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 On the FusionInsight Manager, check whether the alarm list contains **Process Fault**.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 In the **Location** area of **Process Fault**, check whether **ServiceName** is **ZooKeeper**.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Rectify the fault by following the steps provided in **ALM-12007 Process Fault**.

Step 6 In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check the HDFS service status.

Step 7 On the FusionInsight Manager, check whether the alarm list contains **HDFS Service Unavailable**.

- If yes, go to **Step 8**.
- If no, go to **Step 10**.

Step 8 Rectify the fault by following the steps provided in **ALM-14000 HDFS Service Unavailable**.

Step 9 In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check the DBService status.

Step 10 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService** to check whether the DBService running status is **Normal**.

- If yes, go to **Step 12**.
- If no, go to **Step 11**.

Step 11 Choose **More** > **Restart Service** to restart the DBService service. In the alarm list, check whether **LoaderService Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check the Mapreduce status.

Step 12 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Mapreduce** to check whether the Mapreduce running status is **Normal**.

- If yes, go to **Step 16**.
- If no, go to **Step 13**.

Step 13 Choose **More** > **Restart Service** to restart the Mapreduce service. In the alarm list, check whether **LoaderService Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Check the Yarn status.

Step 14 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Yarn** to check whether the Yarn running status is **Normal**.

- If yes, go to **Step 16**.
- If no, go to **Step 15**.

Step 15 Choose **More** > **Restart Service** to restart the Yarn service. In the alarm list, check whether **LoaderService Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Step 16 On the FusionInsight Manager, check whether the alarm list contains **Yarn Service Unavailable**.

- If yes, go to **Step 17**.
- If no, go to **Step 19**.

Step 17 Rectify the fault by following the steps provided in **ALM-18000 Yarn Service Unavailable**.

Step 18 In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 19**.

Check the network connection between Loader and dependent components.

Step 19 On the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.

Step 20 Click **Instance** and the LoaderServer instance list is displayed.

Step 21 Record the **Management IP Address** in the row of **LoaderServer(Active)**.

Step 22 Log in to the host where the active LoaderServer runs as **omm** user using the IP address obtained in **Step 21**.

Step 23 Run the **ping** command to check whether communication between the host that runs the active LoaderServer and the hosts that run the dependent components. (The dependent components include ZooKeeper, DBService, HDFS, Mapreduce and Yarn. Obtain the IP addresses of the hosts that run these services in the same way as that for obtaining the IP address of the active LoaderServer.)

- If yes, go to **Step 26**.
- If no, go to **Step 24**.

Step 24 Contact the administrator to restore the network.

Step 25 In the alarm list, check whether **Loader Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 26**.

Collect fault information.

Step 26 On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 27 Select the following nodes in the required cluster from the **Service** drop-down list:

- ZooKeeper
- HDFS
- DBService
- Yarn
- Mapreduce
- Loader

Step 28 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 29 On the FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader**.

Step 30 Choose **More** > **Restart Service**, and click **OK**.

Step 31 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 32**.

Step 32 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.190 ALM-23003 Loader Task Execution Failure

Description

This alarm is generated immediately when the system detects that the Loader job fails. This alarm is cleared when the failed job is manually handled by a user. This alarm must be manually cleared.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23003	Minor	No

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.
JobID	Specifies the ID of failed Loader job.
JobName	Specifies the failed Loader job.
UserName	Specifies the name of the user who submits the Loader job.
Details	Supplementary information for which the alarm is generated.

Impact on the System

An exception occurs during the execution of the submitted Loader job, and the execution fails. No execution result can be obtained. Execute the job again after rectifying the fault.

Possible Causes

- Task parameters are incorrectly configured.
- Exceptions occur when Yarn is executing a job.

Procedure

Check whether task parameters are incorrectly configured.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** and click the alarm drop-down list from the alarm list, obtain the **Alarm Cause**.

Step 2 If the alarm cause is "Failure to submit job", view error details in **Additional Information**, and go to the Loader WebUI to view the execution history of the job.

NOTE

By default, the **admin** user does not have the permissions to manage other components. If the page cannot be opened or the displayed content is incomplete when you access the native UI of a component due to insufficient permissions, you can manually create a user with the permissions to manage that component.

Step 3 Submit the task again.

Step 4 Check whether the task executed successfully.

- If yes, go to **Step 9**.
- If no, go to **Step 5**.

Check whether exceptions occur when Yarn is executing a job.

Step 5 On FusionInsight Manager, click the alarm drop-down list from the alarm list, obtain the **Alarm Cause**.

Step 6 Check whether the Yarn activity is executed properly in the **Alarm Cause**. If the alarm cause is "Yarn execution failed", the Yarn activity is abnormal.

- If yes, go to **Step 7**.
- If no, go to **Step 10**.

Step 7 Submit the task again.

Step 8 Please check whether the task executed successfully.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 In the alarm list, click **Clear** from **Operation** to manually clear the alarm. No further action is required.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 11 Select the following nodes in the required cluster from the **Service** drop-down list:

- DBService
- HDFS
- Loader
- Mapreduce
- Yarn
- ZooKeeper

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system does not automatically clear this alarm, and you need to manually clear the alarm.

Related Information

None

10.191 ALM-23004 Loader Heap Memory Usage Exceeds the Threshold

Description

The system checks the heap memory usage of the Loader service every 60 seconds. The alarm is generated when the heap memory usage of a Loader instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. The alarm is cleared when the heap memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23004	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The heap memory overflow may cause a service breakdown.

Possible Causes

The heap memory of the Loader instance is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Loader Heap Memory Usage Exceeds the Threshold** > **Location.** Check the host name of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize** > **Memory** > **Loader Heap Memory Resource Percentage**. Click **OK**.

Figure 10-65 Loader Heap Memory Resource Percentage

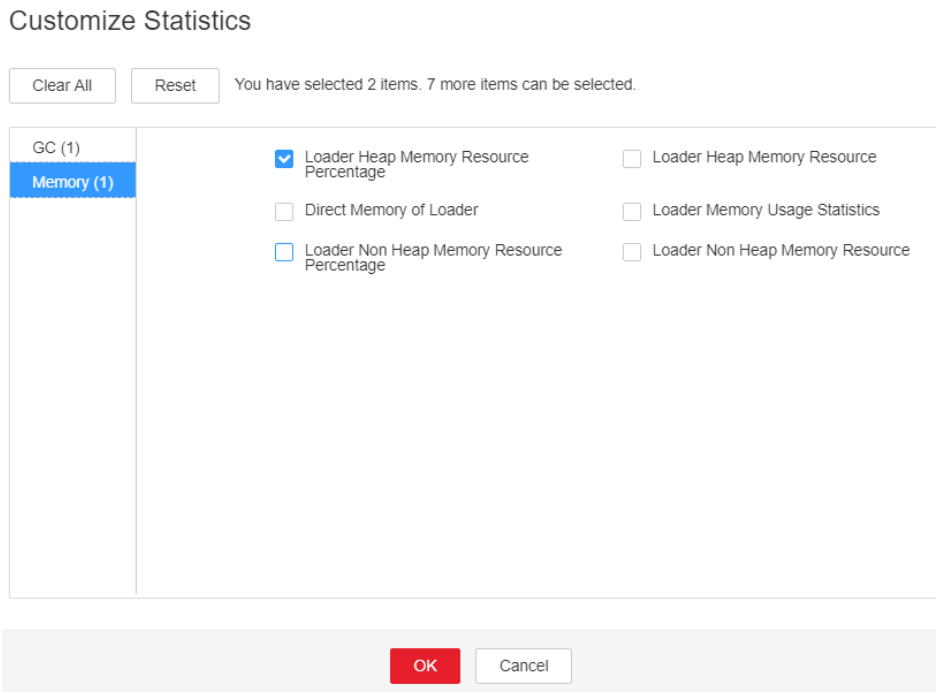

Step 3 Check whether the used heap memory of Loader reaches the threshold (the default value is 95% of the maximum heap memory) specified for Loader.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Configurations**. Click **All Configurations**. Increase the value of **-Xmx** in **GC_OPTS** as required, and click **Save**. Click **OK**.

NOTE

- If this alarm is generated, the heap memory configured for the current Loader instance is insufficient for data transmission. You are advised to open the instance monitoring page, display the Loader heap memory resource status monitoring chart, and observe the change trend of the heap memory used by Loader in the monitoring chart. Then change the value of **-Xmx** to twice the current heap memory usage or to another value to meet site requirements.
- When setting the heap memory, you can set **-Xms** and **-Xmx** to similar values to avoid performance deterioration caused by heap size adjustment after each GC.
- Ensure that the sum of **-Xmx** and **XX:MaxPermSize** is not greater than the physical memory of the node server.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Loader** in the required cluster from the **Service** drop-down list.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.192 ALM-23005 Loader Non-Heap Memory Usage Exceeds the Threshold

Description

The system checks the non-heap memory usage of the Loader service every 30 seconds. The alarm is generated when the non-heap memory usage of a Loader instance exceeds the threshold (80% of the maximum memory) for 5 consecutive times. The alarm is cleared when the non-heap memory usage is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23005	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The non-heap memory overflow may cause a service breakdown.

Possible Causes

The non-heap memory of the Loader instance is overused or the non-heap memory is inappropriately allocated.

Procedure

Check non-heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Loader Non-Heap Memory Usage Exceeds the Threshold** > **Location.** Check the host name of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize** > **Memory** > **Loader Non Heap Memory Resource Percentage**. Click **OK**.

Figure 10-66 Loader Non Heap Memory Resource Percentage

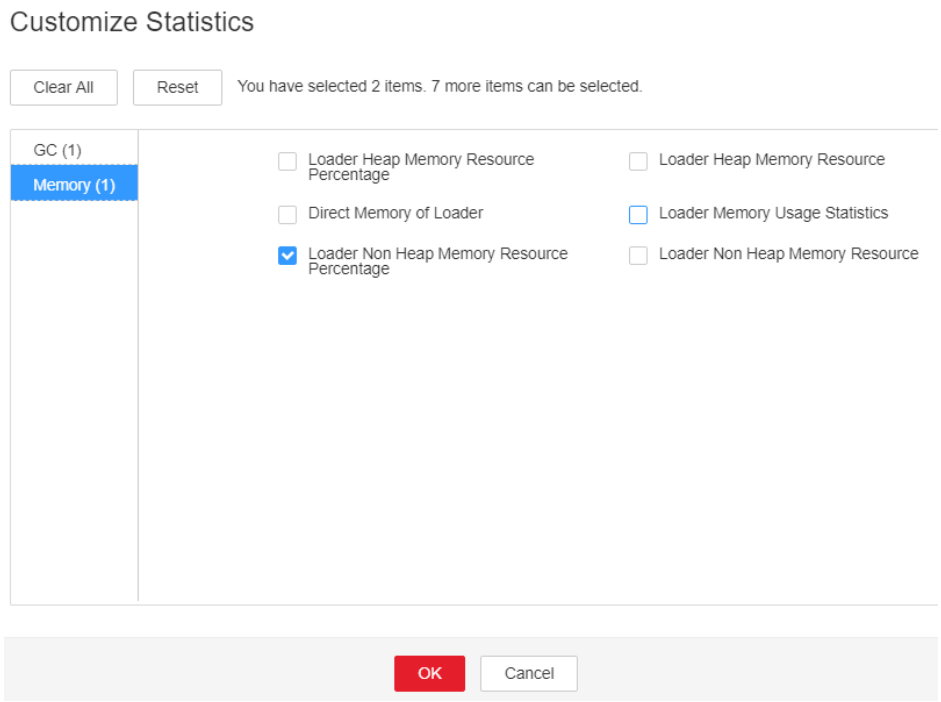

Step 3 Check whether the used non-heap memory of Loader reaches the threshold (the default value is 80% of the maximum non-heap memory) specified for Loader.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Configurations**. Click **All Configurations** Search **LOADER_GC_OPTS** in the search box. If the **-XX: MaxPermSize** parameter is not configured, set the initial value to **-XX: MaxPermSize=256M** for the first time. (If the alarm persists after the first adjustment, perform the second adjustment by referring to the following note.) And click **Save**. Click **OK**.

NOTE

If this alarm is generated, the non-heap memory configured for the current Loader instance is insufficient for the service scenario. You are advised to open the instance monitoring page, open the Loader non-heap memory resource status monitoring chart, and observe the change trend of the non-heap memory used by Loader in the monitoring chart. Then change the value of **-XX:MaxPermSize** to twice the current non-heap memory usage or to another value to meet site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Loader** in the required cluster from the **Service** drop-down list.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.193 ALM-23006 Loader Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the Loader service every 30 seconds. The alarm is generated when the direct memory usage of a Loader

instance exceeds the threshold (80% of the maximum memory) for 5 consecutive times. The alarm is cleared when the direct memory usage of Loader is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23006	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The direct memory overflow may cause a service breakdown.

Possible Causes

The direct memory of the Loader instance is overused or the direct memory is inappropriately allocated.

Procedure

Check direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Loader Direct Memory Usage Exceeds the Threshold** > **Location**. Check the host name of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the

chart area and choose **Customize** > **Memory** > **Loader Memory Usage Statistics**. Click **OK**.

Figure 10-67 Loader Memory Usage Statistics

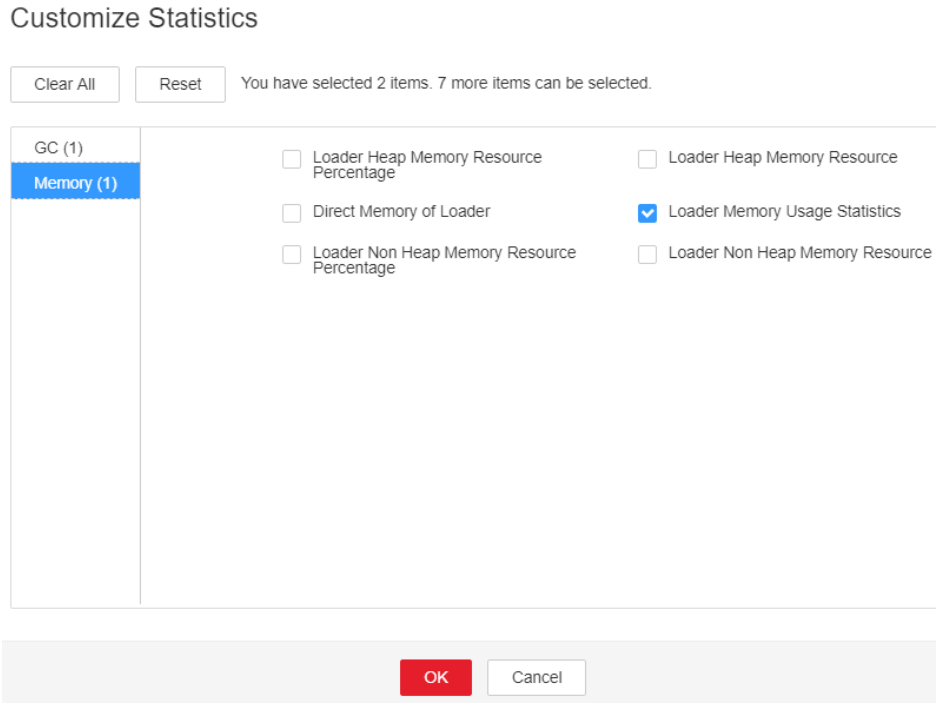

Step 3 Check whether the used direct memory of Loader reaches the threshold (the default value is 80% of the maximum direct memory) specified for Loader.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Configurations**. Click **All Configurations**. Search **LOADER_GC_OPTS** in the search box. Increase the value of **-XX:MaxDirectMemorySize** as required, and click **Save**. Click **OK**.

NOTE

If this alarm is generated, the direct memory configured for the current Loader instance is insufficient for the service scenario. You are advised to open the instance monitoring page, display the Loader direct memory resource status monitoring chart, and observe the change trend of the direct memory used by Loader in the monitoring chart. Then change the value of **-XX:MaxDirectMemorySize** to twice the current direct memory usage or to another value to meet site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Loader** in the required cluster from the **Service** drop-down list.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.194 ALM-23007 Garbage Collection (GC) Time of the Loader Process Exceeds the Threshold

Description

The system checks GC time of the Loader process every 60 seconds. The alarm is generated when GC time of the Loader process exceeds the threshold (default value: **12 seconds**) for 5 consecutive times. The alarm is cleared when GC time is less than the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
23007	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Loader service response is slow.

Possible Causes

The heap memory of the Loader instance is overused or the heap memory is inappropriately allocated.

Procedure

Check GC time.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Garbage Collection (GC) Time of the Loader Process Exceeds the Threshold** > **Location**. Check the host name of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize** > **GC** > **Garbage Collection (GC) Time of Loader**. Click **OK**.

Figure 10-68 Garbage Collection (GC) Time of Loader

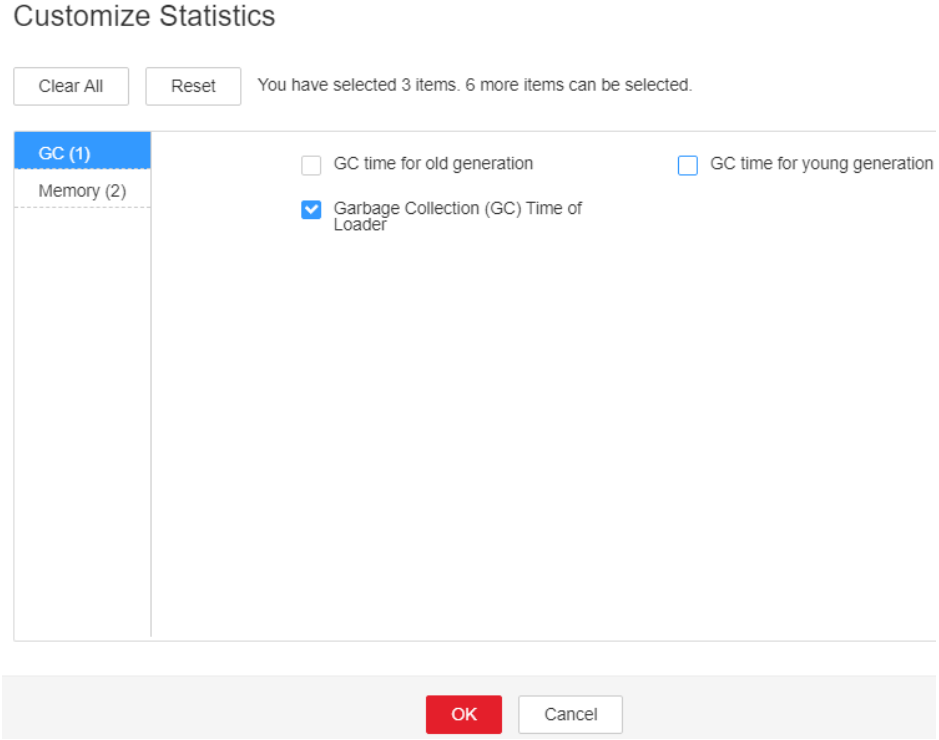

Step 3 Check whether GC time of the Loader process every second exceeds the threshold (default value: **12 seconds**).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Loader** > **Configurations**. Click **All Configurations**. Search **LOADER_GC_OPTS** in the search box. Increase the value of **-Xmx** as required, and click **Save**. Click **OK**.

NOTE

If this alarm is generated, the heap memory configured for the current Loader instance cannot meet the heap memory required for data transmission. You are advised to handle the problem by referring to **Step 4** in section **ALM-23004 Loader Heap Memory Usage Exceeds the Threshold**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Loader** in the required cluster from the **Service** drop-down list.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.195 ALM-24000 Flume Service Unavailable

Alarm Description

The alarm module checks the Flume service status every 180 seconds. This alarm is generated if the Flume service is abnormal.

This alarm is automatically cleared after the Flume service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24000	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Flume cannot work and data transmission is interrupted.

Possible Causes

All Flume instances are faulty.

Handling Procedure

Step 1 Log in to a Flume node as user **omm** and run the **ps -ef|grep "flume.role=server"** command to check whether the Flume process exists on the node.

- If yes, go to **Step 3**.
- If no, restart the faulty Flume node or Flume service and go to **Step 2**.

Step 2 In the alarm list, check whether alarm "Flume Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Collect the fault information.

Step 3 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 4 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 5 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 6 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.196 ALM-24001 Flume Agent Exception

Alarm Description

The Flume agent monitoring module monitors the Flume agent status. This alarm is generated when the Flume agent process is faulty (checked every 5 seconds) or the Flume agent fails to start (an alarm is reported immediately).

This alarm is cleared when the Flume agent process recovers, Flume agent starts successfully, and the alarm handling is completed.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24001	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
AgentId	Specifies the ID of the agent for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The Flume agent instance for which the alarm is generated cannot provide services properly, and the data transmission tasks of the instance are temporarily interrupted. Real-time data is lost during real-time data transmission.

Possible Causes

- The **JAVA_HOME** directory does not exist, or the Java permission is incorrect.
- The Flume agent directory permission is incorrect.
- The Flume agent fails to start.

Handling Procedure

Check whether the JAVA_HOME directory exists or whether the JAVA permission is correct.

Step 1 Log in to the host for which the alarm is generated as user **root**.

Step 2 Obtain the installation directory of the Flume client for which the alarm is generated. (The value of **AgentId** can be obtained from **Location** of the alarm.)

ps -ef|grep *AgentId* **| grep -v grep | awk -F 'conf-file ' '{print $2}' | awk -F 'fusioninsight' '{print $1}'**

Step 3 Run the **su -** *Flume installation user* command to switch to the Flume installation user and run the **cd** *Flume client installation directory*/**fusioninsight-flume-1.9.0/conf/** command to go to the Flume configuration directory.

Step 4 Run the **cat ENV_VARS | grep JAVA_HOME** command.

Step 5 Check whether the **JAVA_HOME** directory exists. If both the command output in **Step 4** and **ll $JAVA_HOME/** are not empty, the **JAVA_HOME** directory exists.

- If yes, go to **Step 7**.
- If no, go to **Step 6**.

Step 6 Specify a correct **JAVA_HOME** directory, for example, **export JAVA_HOME=${BIGDATA_HOME}/common/runtime0/jdk***Version number*.

Step 7 Run the **$JAVA_HOME/bin/java -version** command to check whether the Flume agent running user has the Java execution permission. If the Java version is displayed in the command output, the Java permission meets the requirement. Otherwise, the Java permission does not meet the requirement.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

NOTE

JAVA_HOME is the environment variable exported during Flume client installation. You can also go to *Flume client installation directory*/**fusioninsight-flume-1.9.0/conf** and run the **cat ENV_VARS | grep JAVA_HOME** command to view the variable value.

Step 8 Run the **chmod 750 $JAVA_HOME/bin/java** command to grant the Java execution permission to the Flume agent running user.

Check the directory permission of the Flume agent.

Step 9 Log in to the host for which the alarm is generated as user **root**.

Step 10 Run the following command to switch to the Flume agent installation directory:

cd *Flume client installation directory*/**fusioninsight-flume-1.9.0/conf/**

Step 11 Run the **ls -al * -R** command to check whether any file owner is the user running the Flume agent.

- If yes, go to **Step 12**.
- If no, run the **chown** command to change the file owner to the user who runs the Flume agent.

Check the Flume agent configuration.

Step 12 Run the **cat properties.properties | grep spooldir** and **cat properties.properties | grep TAILDIR** commands to check whether the Flume source type is spoolDir or tailDir. If any command output is displayed, the Flume source type is spoolDir or tailDir.

- If yes, go to **Step 13**.
- If no, go to **Step 17**.

Step 13 Check whether the data monitoring directory exists.

- If yes, go to **Step 15**.
- If no, go to **Step 14**.

NOTE

Run the **cat properties.properties | grep spoolDir** command to view the spoolDir monitoring directory.

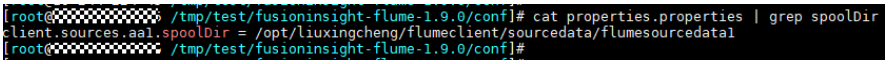

Run the **cat properties.properties | grep parentDir** command to view the tailDir monitoring directory.

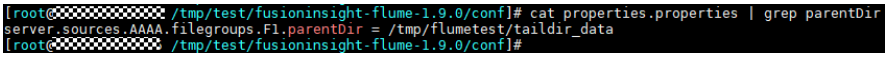

Step 14 Specify a correct data monitoring directory.

Step 15 Check whether the Flume agent user has the read, write, and execute permissions on the monitoring directory specified in **Step 13**.

- If yes, go to **Step 17**.
- If no, go to **Step 16**.

NOTE

Go to the monitoring directory as the Flume running user. If files can be created, the Flume running user has the read, write, and execute permissions on the monitoring directory.

Step 16 Run the **chmod 777** *Flume monitoring directory* command to grant the Flume agent running user the read, write, and execute permissions on the monitoring directory specified in **Step 13**.

Step 17 Check whether the components connected to the Flume sink are in safe mode.

- If yes, go to **Step 18**.

- If no, go to **Step 23**.

  NOTE

  If the sinks in the **properties.properties** configuration file are the HDFS sink and HBase sink, and the configuration file contains a keytab file, the components connected to the Flume sink are in safe mode.

  If the sink in the **properties.properties** configuration file is the Kafka sink and ***.security.protocol** is set to **SASL_PLAINTEXT** or **SASL_SSL**, Kafka connected to the Flume sink is in safe mode.

Step 18 Run the **ll** *ketab path* command to check whether the keytab authentication path specified by the ***.kerberosKeytab** parameter in the configuration file exists.

- If yes, go to **Step 20**.
- If no, go to **Step 19**.

  NOTE

  To view the ketab path, run the **cat properties.properties | grep keytab** command.

Step 19 Change the value of **kerberosKeytab** in **Step 18** to the custom keytab path and go to **Step 21**.

Step 20 Go to **Step 18** and check whether the Flume agent running user has the permission to access the keytab authentication file. If the keytab path is returned, the user has the permission. Otherwise, the user does not have the permission.

- If yes, go to **Step 22**.
- If no, go to **Step 21**.

Step 21 Run the **chmod 755** *ketab file* command to grant the read permission on the keytab file specified in **Step 19**, and restart the Flume process.

Step 22 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 23**.

Collect fault information.

Step 23 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 24 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 25 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 26 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.197 ALM-24003 Flume Client Connection Interrupted

Alarm Description

The alarm module monitors the port connection status on the Flume server. This alarm is generated if the Flume server fails to receive a connection message from the Flume client in three consecutive minutes.

This alarm is cleared after the Flume server receives a connection message from the Flume client.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24003	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
Client IP Address	Specifies the IP address of the Flume client.
Client Name	Specifies the agent name of the Flume client.
Sink Name	Specifies the sink name of Flume Agent.

Impact on the System

The communication between the Flume client and the server fails. The Flume client cannot send data to the Flume server.

Possible Causes

- The network connection between the Flume client and the server is faulty.
- The Flume client's process is abnormal.
- The Flume client is incorrectly configured.

Handling Procedure

Check the network connection between the Flume client and the server.

Step 1 Log in to the host whose IP address is specified by **Flume ClientIP** in the alarm information as user **root**.

Step 2 Run the **ping** *Flume server IP address* command to check whether the network connection between the Flume client and the server is normal.

- If yes, go to **Step 3**.
- If no, go to **Step 11**.

Check whether the Flume client's process is normal.

Step 3 Log in to the host whose IP address is specified by **Flume ClientIP** in the alarm information as user **root**.

Step 4 Run the **ps -ef|grep flume |grep client** command to check whether the Flume client process exists.

- If yes, go to **Step 5**.
- If no, go to **Step 11**.

Check the Flume client configuration.

Step 5 Log in to the host whose IP address is specified by **Flume ClientIP** in the alarm information as user **root**.

Step 6 Run the **cd** *Flume client installation directory*/**fusioninsight-flume-1.9.0/conf/** command to go to Flume's configuration directory.

Step 7 Run the **cat properties.properties** command to query the current configuration file of the Flume client.

Step 8 Check whether the **properties.properties** file is correctly configured according to the configuration description of the Flume agent.

- If yes, go to **Step 9**.
- If no, go to **Step 11**.

Step 9 Modify the **properties.properties** configuration file.

Check whether the alarm is cleared.

Step 10 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect the fault information.

Step 11 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 12 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 13 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Collect logs in the **/var/log/Bigdata/flume-client** directory on the Flume client using a transmission tool.

Step 15 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.198 ALM-24004 Exception Occurs When Flume Reads Data

Alarm Description

The alarm module monitors the Flume source status. This alarm is generated when the duration in which the source cannot read data exceeds the threshold.

The default threshold is **0**, indicating that this function is disabled. You can change the threshold by modifying the **properties.properties** file in the **conf** directory. Specifically, modify the **NoDatatime** parameter of required the source.

The alarm is cleared when the source reads the data and the alarm handling is complete.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24004	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
AgentId	Specifies the ID of the agent for which the alarm was generated.
ComponentType	Specifies the type of the component for which the alarm was generated.

Parameter	Description
ComponentName	Specifies the name of the component for which the alarm was generated.

Impact on the System

If data is found in the data source but the Flume source continuously fails to read data, the collection is stopped.

Possible Causes

- The Flume source is faulty, so data cannot be sent.
- The network is faulty, so the data cannot be sent.

Handling Procedure

Check whether the Flume source is faulty.

Step 1 Open the **properties.properties** configuration file on the local PC, search for keyword **type = spooldir** in the file, and check whether the Flume source type is spoolDir.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 View the spoolDir monitoring directory to check whether all files are already transferred.

- If yes, no further action is required.
- If no, go to **Step 5**.

NOTE

The monitoring directory of spoolDir is specified by the **.spoolDir** parameter in the **properties.properties** configuration file. If all files in the monitoring directory have been transferred, the file name extension of all files in the monitoring directory is **.COMPLETED**.

Step 3 Open the **properties.properties** configuration file on the local PC, search for **org.apache.flume.source.kafka.KafkaSource** in the file, and check whether the Flume source type is Kafka.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Step 4 Check whether the topic data configured by the Kafka source has been used up.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Flume** > **Instances**.

Step 6 Go to the Flume instance page of the faulty node to check whether the **Source Speed Metrics** in the alarm is **0**.

- If yes, go to **Step 11**.
- If no, go to **Step 7**.

Check the network connectivity between the node with the IP address configured for the Flume source and the faulty node.

Step 7 Open the **properties.properties** configuration file on the local PC, search for **type = avro** in the file, and check whether the Flume source type is Avro.

- If yes, go to **Step 8**.
- If no, go to **Step 11**.

Step 8 Log in to the faulty node as user **root**, and run the **ping** *IP address of the Flume source* command to check whether the peer host can be pinged successfully.

- If yes, go to **Step 11**.
- If no, go to **Step 9**.

Step 9 Contact the network administrator to restore the network.

Step 10 Wait for a while and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 12 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 13 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.199 ALM-24005 Exception Occurs When Flume Transmits Data

Alarm Description

The alarm module monitors the capacity status of Flume Channel. The alarm is generated immediately when the duration that Channel is fully occupied exceeds

the threshold or the number of times that Source fails to send data to Channel exceeds the threshold.

The default threshold is **10**. You can change the threshold by modifying the **channelfullcount** parameter of the related channel in the **properties.properties** configuration file in the **conf** directory.

The alarm is cleared when the space of Flume Channel is released and the alarm handling is complete.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24005	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
AgentId	Specifies the ID of the agent for which the alarm was generated.
ComponentType	Specifies the type of the component for which the alarm was generated.
ComponentName	Specifies the name of the component for which the alarm was generated.

Impact on the System

If the disk usage of Flume Channel increases continuously, the time required for importing data to a specified destination prolongs. When the disk usage of Flume Channel reaches 100%, the Flume agent process pauses.

Possible Causes

- Flume Sink is faulty, so the data cannot be sent.
- The network is faulty, so the data cannot be sent.

Handling Procedure

Check whether Flume Sink is faulty.

Step 1 Open the **properties.properties** configuration file on the local PC, search for **type = hdfs** in the file, and check whether the Flume sink type is HDFS.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 On FusionInsight Manager, check whether **HDFS Service Unavailable** alarm is generated in the alarm list and whether the HDFS service is stopped in the service list.

- If the alarm is reported, clear it according to the handling suggestions of ALM-14000 HDFS Service Unavailable; if the HDFS service is stopped, start it. Then, go to **Step 7**.
- If no, go to **Step 7**.

Step 3 Open the **properties.properties** configuration file on the local PC, search for **type = hbase** in the file, and check whether the Flume sink type is HBase.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 On FusionInsight Manager, check whether **HBase Service Unavailable** alarm is generated in the alarm list and whether the HBase service is stopped in the service list.

- If the alarm is reported, clear it according to the handling suggestions of ALM-19000 HBase Service Unavailable; if the HBase service is stopped, start it. Then, go to **Step 7**.
- If no, go to **Step 7**.

Step 5 Open the **properties.properties** configuration file on the local PC, search for **org.apache.flume.sink.kafka.KafkaSink** in the file, and check whether the Flume sink type is Kafka.

- If yes, go to **Step 6**.
- If no, go to **Step 9**.

Step 6 On FusionInsight Manager, check whether **Kafka Service Unavailable** alarm is generated in the alarm list and whether the Kafka service is stopped in the service list.

- If the alarm is reported, clear it according to the handling suggestions of ALM-38000 Kafka Service Unavailable; if the Kafka service is stopped, start it. Then, go to **Step 7**.
- If no, go to **Step 7**.

Step 7 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Flume** > **Instance**.

Step 8 Go to the Flume instance page of the faulty node to check whether the indicator **Sink Speed Metrics** is 0.

- If yes, go to **Step 13**.
- If no, go to **Step 9**.

Check the network connection between the faulty node and the node that corresponds to the Flume Sink IP address.

Step 9 Open the **properties.properties** configuration file on the local PC, search for **type = avro** in the file, and check whether the Flume sink type is Avro.

- If yes, go to **Step 10**.
- If no, go to **Step 13**.

Step 10 Log in to the faulty node as user **root**, and run the **ping** *IP address of the Flume sink* command to check whether the peer host can be pinged successfully.

- If yes, go to **Step 13**.
- If no, go to **Step 11**.

Step 11 Contact the network administrator to restore the network.

Step 12 In the alarm list, check whether the alarm is cleared after a period.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collect the fault information.

Step 13 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 14 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 15 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.200 ALM-24006 Heap Memory Usage of Flume Server Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the Flume service every 60 seconds. This alarm is generated when the heap memory usage of the Flume instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24006	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Heap memory overflow may cause service breakdown.

Possible Causes

The heap memory of the Flume instance is overused or the heap memory is inappropriately allocated.

Handling Procedure

Check the heap memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **Flume Heap Memory Usage Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the target cluster* > **Services** > **Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Flume Heap Memory Resource Percentage**. Then, click **OK**.

Figure 10-69 Flume Heap Memory Resource Percentage

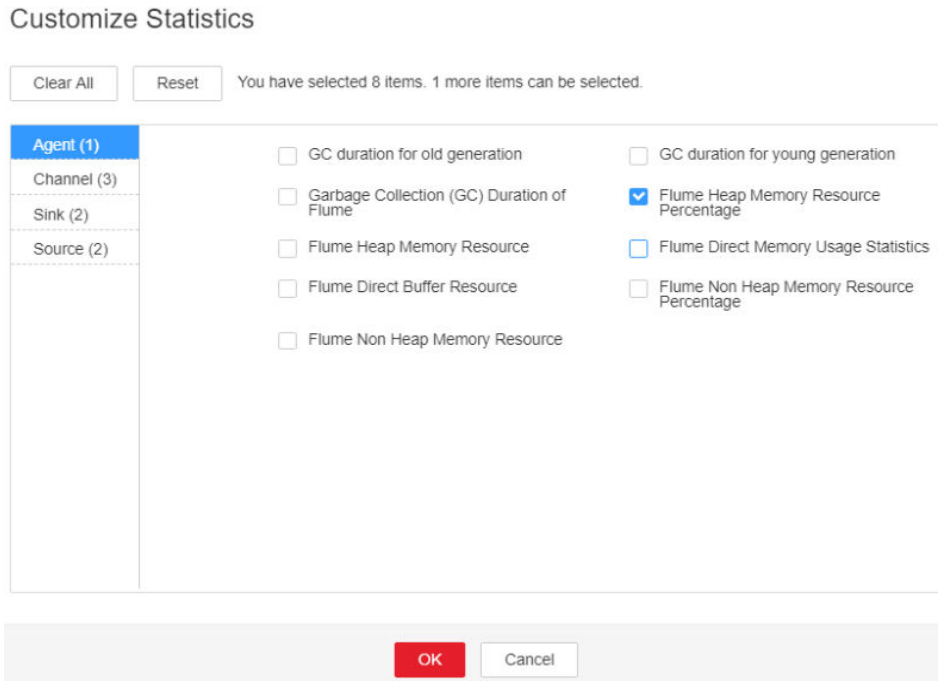

Step 3 Check whether the heap memory used by Flume reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume** > **System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for the Flume server is insufficient for data transmission. You are advised to change the heap memory to: Channel capacity x Maximum size of a single data record x Number of channels. Note that the value of **xmx** cannot exceed the remaining memory of the node.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

NOTICE

- During the restart, the Flume service is interrupted.
- During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.201 ALM-24007 Flume Server Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the Flume service every 60 seconds. This alarm is generated when the direct memory usage of the Flume instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the Flume direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24007	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service breakdown.

Possible Causes

The direct memory of the Flume process is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **Flume Direct Memory Usage Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the target cluster* > **Services** > **Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Flume Direct Memory Resource Percentage**. Then, click **OK**.

Figure 10-70 Flume Direct Memory Usage Statistics

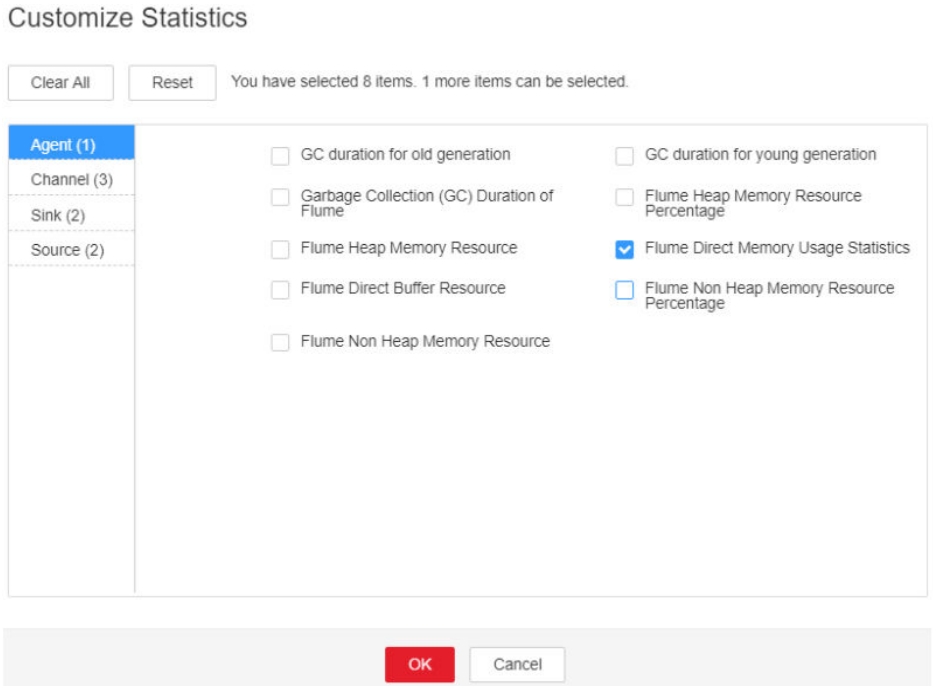

Step 3 Check whether the direct memory used by Flume reaches the threshold (80% of the maximum direct memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume** > **System**. Set **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the direct memory size configured for the Flume server instance cannot meet service requirements. You are advised to change the value of **-XX:MaxDirectMemorySize** to twice the current direct memory size or change the value based on site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

NOTICE

- During the restart, the Flume service is interrupted.
- During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.202 ALM-24008 Flume Server Non Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the non-heap memory usage of the Flume service every 60 seconds. This alarm is generated when the non-heap memory usage of the Flume instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24008	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Non-heap memory overflow may cause service breakdown.

Possible Causes

The non-heap memory of the Flume instance is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **Flume Non Heap Memory Usage Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the target cluster* > **Services** > **Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Flume Non Heap Memory Resource Percentage**. Then, click **OK**.

Figure 10-71 Flume Non-Heap Memory Resource Percentage

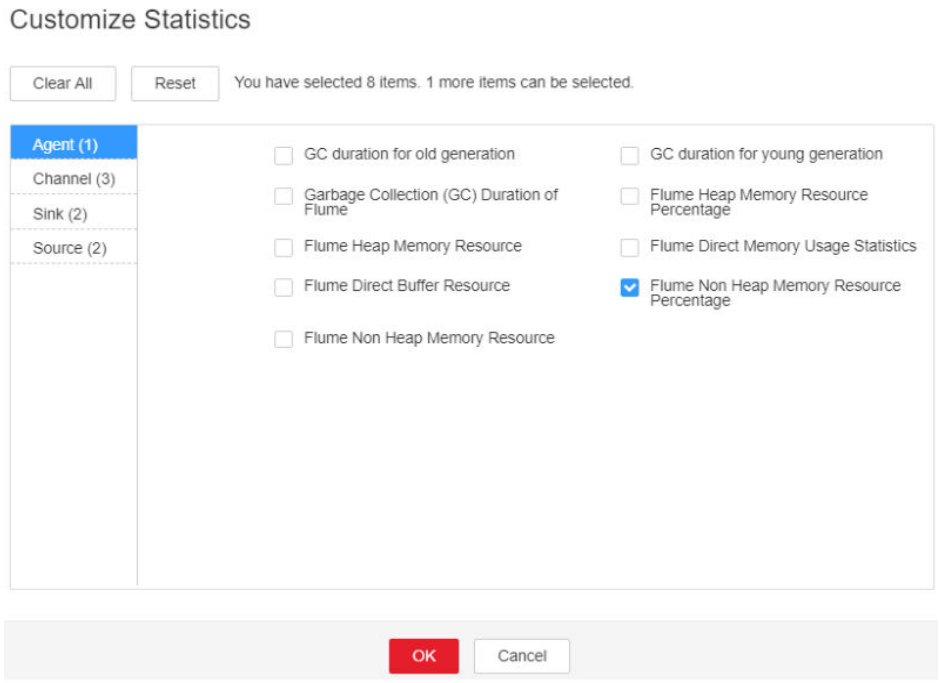

Step 3 Check whether the non-heap memory used by Flume reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume** > **System**. Set **-XX: MaxPermSize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the non-heap memory size configured for the Flume server instance cannot meet service requirements. You are advised to change the value of **-XX:MaxPermSize** to twice the current non-heap memory size or change the value based on site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

NOTICE

- During the restart, the Flume service is interrupted.
- During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.203 ALM-24009 Flume Server Garbage Collection (GC) Time Exceeds the Threshold

Alarm Description

The system checks the GC duration of the Flume process every 60 seconds. This alarm is generated when the GC duration of the Flume process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24009	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Flume data transmission efficiency decreases.

Possible Causes

The heap memory of the Flume process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **GC Duration Exceeds the Threshold**, and view the **Location** information. Check the name of the host for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the target cluster* > **Services** > **Flume**. On the page that is displayed, click the **Instance** tab. On the displayed tab page, select the role corresponding to the host name for which the alarm is generated and select **Customize** from the drop-down list in the upper right corner of the chart area. Choose **Agent** and select **Garbage Collection (GC) Duration of Flume**. Then, click **OK**.

Figure 10-72 Garbage Collection (GC) Duration of Flume

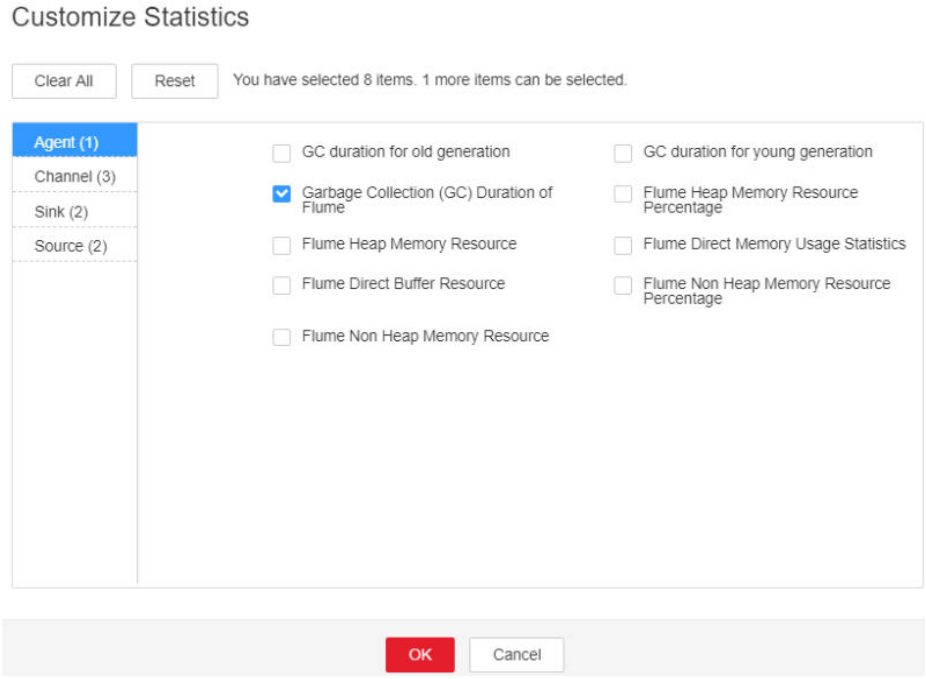

Step 3 Check whether the GC duration of the Flume process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Configuration**. On the page that is displayed, click **All Configurations** and choose **Flume** > **System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for the Flume server is insufficient for data transmission. You are advised to change the heap memory to: Channel capacity x Maximum size of a single data record x Number of channels. Note that the value of **xmx** cannot exceed the remaining memory of the node.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

NOTICE

- During the restart, the Flume service is interrupted.
- During the instance restart, if the failover mode of SinkGroup is configured and at least one instance is running properly, the Flume service is not interrupted. Otherwise, the Flume service is interrupted.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.204 ALM-24010 Flume Certificate File Is Invalid or Damaged

This section applies to MRS 3.2.0 or later.

Alarm Description

Flume checks whether the Flume certificate file is valid (whether the certificate exists and whether the certificate format is correct) every hour. This alarm is generated when the certificate file is invalid or damaged. This alarm is automatically cleared when the certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24010	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The Flume client cannot access the Flume server.

Possible Causes

The Flume certificate file is invalid or damaged.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row

containing **ALM-24010 Flume Certificate File Is Invalid or Damaged**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

Check whether the certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command to go to the Flume service certificate directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf

Step 4 Run the **ls -l** command to check whether the **flume_sChat.crt** file exists.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Run the **openssl x509 -in flume_sChat.crt -text -noout** command to check whether certificate details are displayed properly.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

Step 6 Run the following command to go to the Flume script directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin

Step 7 Run the following command to generate a new certificate file. Then check whether the alarm is automatically cleared one hour later.

sh geneJKS.sh -f *Custom certificate password of the Flume role on the server* **-g** *Custom certificate password of the Flume role on the client*

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

Step 8 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.205 ALM-24011 Flume Certificate File Is About to Expire

This section applies to MRS 3.2.0 or later.

Alarm Description

Flume checks whether the Flume certificate file is about to expire every hour. This alarm is generated when the remaining validity period is at most 30 days. This alarm is automatically cleared when the remaining validity period is greater than 30 days.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24011	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Currently, there is no impact on the system.

Possible Causes

The Flume certificate file is about to expire.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **ALM-24011 Flume Certificate Is About to Expire**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

Check whether the certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command to go to the Flume service certificate directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf

Step 4 Run the following command to check the effective time and expiration time of the Flume user certificate:

openssl x509 -noout -text -in flume_sChat.crt

Step 5 Perform **Step 6** to **Step 7** during off-peak hours to update the certificate file as needed.

Step 6 Run the following command to go to the Flume script directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin

Step 7 Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

sh geneJKS.sh -f *Custom certificate password of the Flume role on the server* **-g** *Custom certificate password of the Flume role on the client*

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

Step 8 Log in to the Flume node for which the alarm is generated as user **omm** and repeat **Step 6** to **Step 7**. Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 10**.
- If no, no further action is required.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 12 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.206 ALM-24012 Flume Certificate File Has Expired

This section applies to MRS 3.2.0 or later.

Alarm Description

Flume checks whether its certificate file in the system has expired every hour. This alarm is generated when the server certificate has expired. This alarm is automatically cleared when the certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24012	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The Flume client cannot access the Flume server.

Possible Causes

The Flume certificate file has expired.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **ALM-24012 Flume Certificate Has Expired**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

Check whether the certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command to go to the Flume service certificate directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf

Step 4 Run the following command to check the effective time and expiration time of the HA user certificate to determine whether the certificate file is still in the validity period:

openssl x509 -noout -text -in flume_sChat.crt

- If yes, go to **Step 9**.
- If no, go to **Step 5**.

Step 5 Run the following command to go to the Flume script directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin

Step 6 Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

sh geneJKS.sh -f *Custom certificate password of the Flume role on the server* **-g** *Custom certificate password of the Flume role on the client*

- If yes, go to **Step 8**.
- If no, go to **Step 7**.

NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

Step 7 Log in to the Flume node for which the alarm is generated as user **omm** and repeat **Step 5** to **Step 6**. Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Flume** for the target cluster.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.207 ALM-24013 Flume MonitorServer Certificate File Is Invalid or Damaged

This section applies to MRS 3.2.0 or later.

Alarm Description

MonitorServer checks whether its certificate file is valid (whether the certificate exists and whether the certificate format is correct) every hour. This alarm is generated when the certificate file is invalid or damaged. This alarm is automatically cleared when the certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24013	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The Flume client cannot access the Flume server.

Possible Causes

The MonitorServer certificate file is invalid or damaged.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row

containing **ALM-24013 MonitorServer Certificate File Is Invalid or Damaged**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

Check whether the certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command to go to the MonitorServer certificate file directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf

Step 4 Run the **ls -l** command to check whether the **ms_sChat.crt** file exists:

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Run the **openssl x509 -in ms_sChat.crt -text -noout** command to check whether certificate details are displayed.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

Step 6 Run the following command to go to the Flume script directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin

Step 7 Run the following command to generate a new certificate file. Then check whether the alarm is automatically cleared one hour later.

sh geneJKS.sh -m *Custom password of the MonitorServer certificate on the server* **-n** *Custom password of the MonitorServer certificate on the client*

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

  NOTE

  The custom certificate passwords must meet the following complexity requirements:

  - Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
  - Contain 8 to 64 characters.
  - Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

Step 8 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Select **MonitorServer** in the required cluster for **Service**.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.208 ALM-24014 Flume MonitorServer Certificate Is About to Expire

This section applies to MRS 3.2.0 or later.

Alarm Description

MonitorServer checks whether its certificate file is about to expire every hour. This alarm is generated when the remaining validity period is at most 30 days. This alarm is automatically cleared when the remaining validity period is greater than 30 days.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24014	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Currently, there is no impact on the system.

Possible Causes

The MonitorServer certificate file is about to expire.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **ALM-24014 MonitorServer Certificate Is About to Expire**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

Check whether the certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command to go to the MonitorServer certificate file directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf

Step 4 Run the following command to check the effective time and expiration time of the MonitorServer user certificate:

openssl x509 -noout -text -in ms_sChat.crt

Step 5 Perform **Step 6** to **Step 7** during off-peak hours to update the certificate file as needed.

Step 6 Run the following command to go to the Flume script directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin

Step 7 Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

sh geneJKS.sh -m *Custom password of the MonitorServer certificate on the server* **-n** *Custom password of the MonitorServer certificate on the client*

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

Step 8 Log in to the Flume node for which the alarm is generated as user **omm** and repeat **Step 6** to **Step 7**. Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 10**.
- If no, no further action is required.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Select **MonitorServer** in the required cluster for **Service**.

Step 12 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.209 ALM-24015 Flume MonitorServer Certificate File Has Expired

This section applies to MRS 3.2.0 or later.

Alarm Description

MonitorServer checks whether its certificate file in the system has expired every hour. This alarm is generated when the server certificate has expired. This alarm is automatically cleared when the MonitorServer certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
24015	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The Flume client cannot access the Flume server.

Possible Causes

The MonitorServer certificate file has expired.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing **ALM-24015 MonitorServer Certificate Has Expired**, and view the **Location** information. View the IP address of the instance for which the alarm is generated.

Check whether the certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the node for which the alarm is generated as user **root** and run the **su - omm** command to switch to user **omm**.

Step 3 Run the following command to go to the MonitorServer certificate file directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/conf

Step 4 Run the following command to check the effective time and expiration time of the user certificate to determine whether the certificate file is still in the validity period:

openssl x509 -noout -text -in ms_sChat.crt

- If yes, go to **Step 9**.
- If no, go to **Step 5**.

Step 5 Run the following command to go to the Flume script directory:

cd ${BIGDATA_HOME}/FusionInsight_Porter_*/install/FusionInsight-Flume-*/flume/bin

Step 6 Run the following command to generate a new certificate file. Then, check whether the alarm is automatically cleared one hour later.

sh geneJKS.sh -m *Custom password of the MonitorServer certificate on the server* **-n** *Custom password of the MonitorServer certificate on the client*

- If yes, go to **Step 8**.
- If no, go to **Step 7**.

NOTE

The custom certificate passwords must meet the following complexity requirements:

- Contain at least four types of uppercase letters, lowercase letters, digits, and special characters.
- Contain 8 to 64 characters.
- Be changed periodically (for example, every three months), and certificates and trust lists are generated again to ensure security.

Step 7 Log in to the Flume node for which the alarm is generated as user **omm** and repeat **Step 5** to **Step 6**. Then, check whether the alarm is automatically cleared one hour later.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Select **MonitorServer** in the required cluster for **Service**.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.210 ALM-25000 LdapServer Service Unavailable

Description

The system checks the LdapServer service status every 30 seconds. This alarm is generated when the system detects that both the active and standby LdapServer services are abnormal.

This alarm is cleared when the system detects that one or two LdapServer services are normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
25000	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

When this alarm is generated, no operation can be performed for the KrbServer users and LdapServer users in the cluster. For example, users, user groups, or roles cannot be added, deleted, or modified, and user passwords cannot be changed on the FusionInsight Manager portal. The authentication for existing users in the cluster is not affected.

Possible Causes

- The node where the LdapServer service locates is faulty.
- The LdapServer process is abnormal.

Procedure

Check whether the nodes where the two SlapdServer instances of the LdapServer service are located are faulty.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **LdapServer** > **Instance** to go to the LdapServer instance page to obtain the host name of the node where the two SlapdServer instances locates.

Step 2 Choose **O&M** > **Alarm** > **Alarms**. On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **Node Fault** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Step 3 Check whether the host name in the alarm is consistent with the **Step 1** host name.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 Handle the alarm according to "ALM-12006 Node Fault".

Step 5 Check whether **LdapServer Service Unavailable** is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check whether the LdapServer process is normal.

Step 6 Choose **O&M** > **Alarm** > **Alarms**. On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **Process Fault** exists.

- If yes, go to **Step 7**.
- If no, go to **Step 10**.

Step 7 Check whether the service and host name in the alarm are consistent with the LdapServer service and host name.

- If yes, go to **Step 8**.
- If no, go to **Step 10**.

Step 8 Handle the alarm according to "ALM-12007 Process Fault".

Step 9 Check whether **LdapServer Service Unavailable** is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 11 Select **LdapServer** in the required cluster from the **Service**.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.211 ALM-25004 Abnormal LdapServer Data Synchronization

Description

The system checks the LdapServer data every 30 seconds. This alarm is generated when the data on the active and standby LdapServers of Manager is inconsistent for 12 consecutive times. This alarm is cleared when the data on the active and standby LdapServers is consistent.

The system checks the LdapServer data every 30 seconds. This alarm is generated when the LdapServer data in the cluster is inconsistent with that on Manager for 12 consecutive times. This alarm is cleared when the data is consistent.

Attribute

Alarm ID	Alarm Severity	Auto Clear
25004	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

LdapServer data inconsistency occurs because the LdapServer data in Manager is damaged or the LdapServer data in the cluster is damaged. The LdapServer process with damaged data cannot provide services externally, and the authentication functions of Manager and the cluster are affected.

Possible Causes

- The network of the node where the LdapServer process locates is faulty.
- The LdapServer process is abnormal.
- The OS restart damages data on LdapServer.
- The amount of Oldap data exceeds the threshold (10 MB by default).

Procedure

Check whether the network where the LdapServer nodes reside is faulty.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. Record the IP address of HostName in the alarm locating information as IP1 (if multiple alarms exist, record the IP addresses as IP1, IP2, and IP3 respectively).

Step 2 Contact O&M personnel and log in to the nodes corresponding to IP 1. Run the ping command to check whether the IP address of the management plane of the active OMS node can be pinged.

- If yes, go to **Step 4**.
- If no, go to **Step 3**.

Step 3 Contact the network administrator to recover the network and check whether **Abnormal LdapServer Data Synchronization** is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the LdapServer processes are normal.

Step 4 On the **Alarm** page of FusionInsight Manager, check whether the **OLdap Resource Abnormal** exists.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Clear the alarm by following the steps provided in "ALM-12004 OLdap Resource Abnormal".

Step 6 Check whether **Abnormal LdapServer Data Synchronization** is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 7**.

Step 7 On the **Alarm** page of FusionInsight Manager, check whether **Process Faul**t is generated for the LdapServer service.

- If yes, go to **Step 8**.
- If no, go to **Step 10**.

Step 8 Handle the alarm according to "ALM-12007 Process Fault".

Step 9 Check whether **Abnormal LdapServer Data Synchronization** is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check whether the LdapServer processes are normal.

Step 10 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Record the IP address of HostName in the alarm locating information as "IP1" (if multiple alarms exist, record the IP addresses as "IP1", "IP2", and "IP3" respectively). Choose **Cluster** > *Name of the desired cluster* > **Services** > **LdapServer** > **Configurations**. Record the port number of LdapServer as "PORT". (If the IP address in the alarm locating information is the IP address of the standby management node, choose **System** > **OMS** > **oldap** > **Modify Configuration** and record the listening port number of LdapServer.)

Step 11 Log in to the nodes corresponding to IP1 as user **omm**.

Step 12 Run the following command to check whether errors are displayed in the queried information.

ldapsearch -H ldaps://*IP1*:*PORT* -LLL -x -D cn=root,dc=hadoop,dc=com -W -b ou=Peoples,dc=hadoop,dc=com

After running the command, enter the **LDAP** administrator password. Contact the system administrator to obtain the password.

- If yes, go to **Step 13**.
- If no, go to **Step 15**.

Step 13 Recover the LdapServer and OMS nodes using data backed up before the alarm is generated.

NOTE

Use the OMS data and LdapServer data backed up at the same point in time to recover the data. Otherwise, the service and operation may fail. To recover data when services run properly, you are advised to manually back up the latest management data and then recover the data. Otherwise, Manager data produced between the backup point in time and the recovery point in time will be lost.

Step 14 Check whether alarm **Abnormal LdapServer Data Synchronization** is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Check whether the data volume of the Oldap exceeds the threshold (10 MB by default). (This step applies only to versions earlier than MRS 3.3.0. For MRS 3.3.0 and later versions, go to **Step 18**.)

Step 15 Log in to the active OMS node as user **omm**.

Step 16 Run the following command to check whether the directory contains **.mdb** files.

ll /srv/BigData/ldapData/oldap/data/

- If yes, check and record the size of the **.mdb** file and go to **Step 17**.

- If no, go to **Step 18**.

Step 17 Run the following command to view the Oldap configuration and record the value of **Map size** (the default value is **10485760** bytes, that is, 10 MB)

mdb_stat -e /srv/BigData/ldapData/oldap/data/

Check whether the size of the **.mdb** file with **Step 16** records reaches the value of **Map size.**

- If yes, contact the O&M personnel.
- If no, go to **Step 18**.

Collect fault information.

Step 18 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 19 Select **LdapServer** in the required cluster and **OmsLdapServer** from the **Service**.

Step 20 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 21 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.212 ALM-25005 nscd Service Exception

Alarm Description

The system checks the status of the nscd service every 60 seconds. This alarm is generated when the nscd process fails to be queried for four consecutive times (three minutes) or users in LdapServer cannot be obtained.

This alarm is cleared when the process is restored and users in LdapServer can be obtained.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
25005	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The alarmed node may not be able to synchronize data from LdapServer. The **id** command may fail to obtain the LDAP data, affecting upper-layer services.

Possible Causes

- The nscd service is not started.
- The network is faulty, and cannot access the LDAP server.
- NameService is abnormal.
- Users cannot be queried because the OS executes commands too slowly.

Handling Procedure

Check whether the nscd service is started.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Record the IP address of **HostName** in **Location** of the alarm as **IP1** (if multiple alarms exist, record the IP addresses as **IP1**, **IP2**, and **IP3** respectively).

Step 2 Contact the O&M personnel to access the node using IP1 as user **root**. Run the **ps -ef | grep nscd** command on the node and check whether the **/usr/sbin/nscd** process is started.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Run the **service nscd restart** command as user **root** to restart the nscd service. Then run the **ps -ef | grep nscd** command to check whether the nscd service is started.

- If yes, go to **Step 4**.
- If no, go to **Step 15**.

Step 4 Wait for 5 minutes and run the ps -ef | grep nscd command again as user root. Check whether the service exists.

- If yes, go to **Step 11**.
- If no, go to **Step 15**.

Check whether the network is faulty, and whether the LDAP server can be accessed.

Step 5 Log in to the alarmed node as user **root** and run the **ping** command to check whether the network connectivity between this node and the LdapServer node is normal.

- If yes, go to **Step 6**.
- If no, contact network administrators to troubleshoot the fault.

Check whether the NameService is normal.

Step 6 Log in to the alarmed node as user **root**. Run the **cat /etc/nsswitch.conf** command to check whether the **passwd**, **group**, **services**, **netgroup**, and **aliases** of NameService are correctly configured.

The correct parameter configurations are as follows:

passwd: **compat ldap**; **group**: **compat ldap**; **services**: **files ldap**; **netgroup**: **files ldap**; **aliases**: **files ldap**

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Log in to the alarmed node as user **root**. Run the **cat /etc/nscd.conf** command to check whether the **enable-cache passwd**, **positive-time-to-live passwd**, **enable-cache group**, and **positive-time-to-live group** in the configuration file are correctly configured.

The correct parameter configurations are as follows:

enable-cache passwd: **yes**; **positive-time-to-live passwd**: **600**; **enable-cache group**: **yes**; **positive-time-to-live group**: **3600**

- If yes, go to **Step 8**.
- If no, go to **Step 10**.

Step 8 Run the **/usr/sbin/nscd -i group** and **/usr/sbin/nscd -i passwd** commands as user **root**. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If yes, go to **Step 11**.
- If no, go to **Step 15**.

Step 9 Run the **vi /etc/nsswitch.conf** command as user **root**. Correct the configurations in **Step 6** and save the file. Run the **service nscd restart** command to restart the nscd service. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If yes, go to **Step 11**.
- If no, go to **Step 15**.

Step 10 Run the **vi /etc/nscd.conf** command as user **root**. Correct the configurations in **Step 7** and save the file. Run the **service nscd restart** command to restart the nscd service. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If yes, go to **Step 11**.
- If no, go to **Step 15**.

Step 11 Log in to the FusionInsight Manager portal. Wait for 5 minutes and check whether the **nscd Service Exception** alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check whether frame freezing occurs when running a command in the operating system.

Step 12 Log in to the faulty node as user **root**, run the **id admin** command, and check whether the command execution takes a long time. If the command execution takes more than 3 seconds, the command execution is deemed to be slow.

- If yes, go to **Step 13**.
- If no, go to **Step 15**.

Step 13 Run the **cat /var/log/messages** command to check whether the nscd frequently restarts or the error information "Can't contact LDAP server" exists.

nscd exception example:

```
Feb 11 11:44:42 10-120-205-33 nscd: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780: Can't contact LDAP server
Feb 11 11:44:43 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.55:21780: Can't contact LDAP server
Feb 11 11:44:44 10-120-205-33 ntpq: nss_ldap: failed to bind to LDAP server ldaps://10.120.205.92:21780: Can't contact LDAP server
```

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Run the **vi$BIGDATA_HOME/tmp/random_ldap_ip_order** command to modify the number at the end. If the original number is an odd number, change it to an even number. If the number is an even number, change it to an odd number.

Run the **vi /etc/ldap.conf** command to enter the editing mode, press **Insert** to start editing, and then change the first two IP addresses of the URI configuration item.

After the modification is complete, press **Esc** to exit the editing mode and enter **:wq!** to save the settings and exit.

Run the **service nscd restart** command to restart the nscd service. Wait 5 minutes and run the **id admin** command again. Check whether the command execution is slow.

- If yes, go to **Step 15**.
- If no, log in to other faulty nodes and repeat **Step 12** to **Step 14** to check whether the first LdapServer node in the URI before modifying **/etc/ldap.conf** is faulty. For example, check whether the service IP address is unreachable, the network delay is too long, or other abnormal software is deployed.

Collect the fault information.

Step 15 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 16 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **LdapClient** for the target cluster.

Step 17 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.213 ALM-25006 Sssd Service Exception

Description

The system checks the status of the sssd service every 60 seconds. This alarm is generated when the sssd process fails to be queried for four consecutive times (three minutes) or users in LdapServer cannot be obtained.

This alarm is cleared when the process is restored and users in LdapServer can be obtained.

Attribute

Alarm ID	Alarm Severity	Auto Clear
25006	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

Impact on the System

The alarmed node may not be able to synchronize data from LdapServer. The id command may fail to obtain the LDAP data, affecting upper-layer services.

Possible Causes

- The sssd service is not started or is incorrectly started.
- The network is faulty and cannot access the LDAP server.
- NameService is abnormal.
- Users cannot be queried because the OS executes commands too slowly.

Procedure

Check whether the sssd service is correctly started.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms**. Find the IP address of **HostName** in **Location** of the alarm and record it as IP1 (if multiple alarms exist, record the IP addresses as IP1, IP2, and IP3 respectively).

Step 2 Contact the O&M personnel to access the node using IP1 as user root. Run the **ps -ef | grep sssd** command and check whether the **/usr/sbin/sssd** process is started.

- If the process is started, go to **Step 3**.
- If the process is not started, go to **Step 4**.

Step 3 Check whether the sssd process queried in **Step 2** has three subprocesses.

- If yes, go to **Step 5**.
- If no, go to **Step 4**.

Step 4 Run the **service sssd restart** command as user **roo**t to restart the sssd service. Then run the **ps -ef | grep sssd** command to check whether the sssd process is normal.

In the normal state, the **/usr/sbin/sssd** process has three subprocesses: **/usr/libexec/sssd/sssd_be**, **/usr/libexec/sssd/sssd_nss**, and **/usr/libexec/sssd/sssd_pam**.

- If it exists, go to **Step 9**.
- If it does not exist, go to **Step 13**.

Check whether the LDAP server can be accessed.

Step 5 Log in to the alarmed node as user **root**. Run the **ping** command to check the network connectivity between this node and the LdapServer node.

- If the network is normal, go to **Step 6**.
- If the network is faulty, contact network administrators to troubleshoot the fault.

Check whether NameService is normal.

Step 6 Log in to the alarmed node as user **root**. Run the **cat /etc/nsswitch.conf** command and check the **passwd** and **group** configurations of NameService.

The correct parameter configurations are as follows: **passwd: compat ldap** and **group: compat ldap**.

- If the configurations are correct, go to **Step 7**.
- If the configurations are incorrect, go to **Step 8**.

Step 7 Run the **/usr/sbin/sss_cache -G** and **/usr/sbin/sss_cache -U** commands as user **root**. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If results are queried, go to **Step 9**.
- If no result is queried, go to **Step 13**.

Step 8 Run the **vi /etc/nsswitch.conf** command as user **root**. Correct the configurations in **Step 6** and save the file. Run the **service sssd restart** command to restart the sssd service. Wait for 2 minutes and run the **id admin** and **id backup/manager** commands to check whether results can be queried.

- If results are queried, go to **Step 9**.
- If no result is queried, go to **Step 13**.

Step 9 Log in to the FusionInsight Manager portal. Wait for 5 minutes and check whether the **sssd Service Exception** alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm persists, go to **Step 10**.

Check whether frame freezing occurs when running a command in the operating system.

Step 10 Log in to the faulty node as user **root**, run the **id admin** command, and check whether the command execution takes a long time. If the command execution takes more than 3 seconds, the command execution is deemed to be slow.

- If yes, go to **Step 11**.
- If no, go to **Step 13**.

Step 11 Run the **cat /var/log/messages** command to check whether the sssd frequently restarts or the error information **Can't contact LDAP server** exists.

sssd restart example:

```
Feb  7 11:38:16 10-132-190-105 sssd[pam]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd[nss]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd[be[default]]: Shutting down
Feb  7 11:38:16 10-132-190-105 sssd: Starting up
Feb  7 11:38:16 10-132-190-105 sssd[be[default]]: Starting up
Feb  7 11:38:16 10-132-190-105 sssd[nss]: Starting up
Feb  7 11:38:16 10-132-190-105 sssd[pam]: Starting up
```

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 Run the **vi $BIGDATA_HOME/tmp/random_ldap_ip_order** command to modify the number at the end. If the original number is an odd number, change it to an even number. If the number is an even number, change it to an odd number.

Run the **vi /etc/sssd/sssd.conf** command to reverse the first two IP addresses of the **ldap_uri** configuration item, save the settings, and exit.

Run the **ps -ef | grep sssd** command to query the ID of the sssd process, kill it, and run the **/usr/sbin/sssd -D -f** command to restart the sssd service. Wait 5 minutes and run the **id admin** command again.

Check whether the command execution is slow.

- If yes, go to **Step 13**.
- If no, log in to other faulty nodes and run **Step 10** to **Step 12**. Collect logs and check whether the first ldapserver node in the ldap_uri before modifying **/etc/sssd/sssd.conf** is faulty. For example, check whether the service IP address is unreachable, the network latency is too long, or other abnormal software is deployed.

Collect fault information.

Step 13 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 14 Select **LdapClient** in the required cluster from the **Service**.

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.214 ALM-25007 Number of SlapdServer Connections Exceeds the Threshold

Alarm Description

The system checks the number of process connections on the SlapdServer node every 30 seconds and compares the actual number with the threshold. This alarm is generated when the number of process connections exceeds the threshold (**1000** by default) for multiple times (**5** by default).

Its **Trigger Count** is configurable. If **Trigger Count** is set to **1**, this alarm is cleared when the number of process connections is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the number of process connections is less than or equal to 90% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
25007	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- There are too many SlapdServer connections.
- The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether there are too many SlapdServer process connections.

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **LdapServer**.

Step 2 On the LdapServer dashboard page, observe the SlapdServer process connections and decrease the connections based on service requirements.

Figure 10-73 SlapdServer process connections

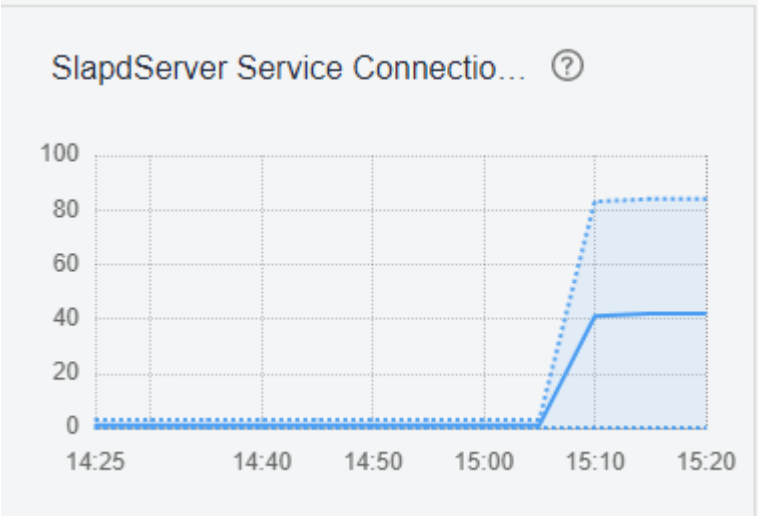

Step 3 Wait about 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 4 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, choose **LdapServer** > **Other** > **SlapdServer Service Connections**, and check whether the alarm trigger count and alarm threshold are set properly.

- If yes, go to **Step 7**.
- If no, go to **Step 5**.

Step 5 Change the trigger count and alarm threshold based on the actual number of process connections, and apply the changes.

Step 6 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **LdapServer** for the target cluster.

Step 9 Specify **Hosts** for collecting logs, which is optional. By default, all hosts are selected.

Step 10 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.215 ALM-25008 SlapdServer CPU Usage Exceeds the Threshold

Alarm Description

The system checks the CPU usage of the SlapdServer node every 30 seconds and compares the actual usage with the threshold. This alarm is generated when the SlapdServer CPU usage exceeds the threshold for multiple times (**5** by default).

Its **Trigger Count** is configurable. If **Trigger Count** is set to **1**, this alarm is cleared when the SlapdServer CPU usage is less than or equal to the threshold. If **Trigger Count** is greater than **1**, this alarm is cleared when the SlapdServer CPU usage is less than or equal to 90% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
25008	Critical (default threshold: 85%) Major (default threshold: 75%)	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- The alarm threshold or alarm trigger count is improperly configured.
- The CPU configuration cannot meet service requirements, and the CPU usage reaches the upper limit.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, choose **LdapServer** > **Other** > **SlapdServer Service Total CPU Percentage**, and check whether the alarm trigger count and alarm threshold are set properly.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Change the trigger count and alarm threshold based on the actual CPU usage, and apply the changes.

Step 3 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the CPU usage reaches the upper limit.

Step 4 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the right pane, click this alarm and obtain the host name in **Location**.

Step 5 Choose **Cluster** > **Services** > **LdapServer**, click the **Instance** tab, and click the SlapdServer instance corresponding to the host name in **Step 4**.

Step 6 On the dashboard of the instance, observe the real-time data of the **CPU Usage of a Single SlapdServer Instance** chart for about 5 minutes and check whether the CPU usage exceeds the threshold (**75%** by default) for multiple times.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Check whether the status of other SlapdServer instances is normal. For details, see **Step 5** to **Step 6**.

- If yes, contact the MRS cluster administrator to evaluate whether to expand the capacity of SlapdServer instances. Then, go to **Step 8**.
- If no, repair the faulty SlapdServer instance and go to **Step 8**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **LdapServer** for the target cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.216 ALM-25500 KrbServer Service Unavailable

Description

The system checks the KrbServer service status every 30 seconds. This alarm is generated when the system detects that the KrbServer service is abnormal.

This alarm is cleared when the system detects that the KrbServer service is normal.

Attribute

Alarm ID	Alarm Severity	Auto Clear
25500	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

When this alarm is generated, no operation can be performed for the KrbServer component in the cluster. The authentication of KrbServer in other components will be affected. The running status of components that depend on KrbServer in the cluster is Bad.

Possible Causes

- The node where the KrbServer service locates is faulty.
- The OLdap service is abnormal.

Procedure

Check whether the node where the KrbServer service locates is faulty.

Step 1 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **KrbServer** > **Instance** to go to the KrbServer instance page to obtain the host name of the node where the KrbServer service locates.

Step 2 On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **Node Fault** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Step 3 Check whether the host name in the alarm is consistent with the **Step 1** host name.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 Handle the alarm according to "ALM-12006 Node Fault".

Step 5 Check whether **KrbServer Service Unavailable** is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the OLdap service is normal.

Step 6 On the **Alarm** page of the FusionInsight Manager system, check whether any alarm of **OLdap Resource Abnormal** exists.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Handle the alarm according to "ALM-12004 OLdap Resource Abnormal".

Step 8 Check whether **KrbServer Service Unavailable** is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On the FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 10 Select **KrbServer** in the required cluster from the **Service**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.217 ALM-26051 Storm Service Unavailable

Description

The system checks the Storm service status every 30 seconds. This alarm is generated when all Nimbus nodes in the cluster are abnormal and the Storm service is unavailable.

This alarm is cleared when the Storm service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26051	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The cluster cannot provide the Storm service, and users cannot perform new Storm tasks.

Possible Causes

- The Kerberos cluster is faulty.
- The ZooKeeper cluster is faulty or suspended.
- The active and standby Nimbus nodes in the Storm cluster are abnormal

Procedure

Check the status of the Kerberos cluster. (Skip this step if the normal mode is used.)

Step 1 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services**.

Step 2 Check whether the running status of the Kerberos service is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 See the related maintenance information of **ALM-25500 KrbServer Service Unavailable**.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the status of the ZooKeeper cluster.

Step 5 Check whether the running status of the ZooKeeper service is **Normal**.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 If ZooKeeper service is stopped, start it, else see the related maintenance information of **ALM-13000 ZooKeeper Service Unavailable**.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check the status of the active and standby Nimbus nodes.

Step 8 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Nimbus** to go to the Nimbus Instances page.

Step 9 Check whether only one Nimbus node that is in the **Active** state in **Roles**.

- If yes, go to **Step 13**.
- If no, go to **Step 10**.

Step 10 Select two Nimbus role instances, choose **More** > **Restart Instance**, and check whether the instances restart successfully.

- If yes, go to **Step 11**.
- If no, go to **Step 13**.

Step 11 Log in to the FusionInsight Manager portal again, choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Nimbus** to check whether the running status is **Normal**.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 Wait for 30 seconds and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collecting Fault Information

Step 13 On the FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 14 Select the following nodes in the required cluster from the **Service** drop-down list:

- KrbServer

  NOTE

  KrbServer logs do not need to be downloaded in normal mode.

- ZooKeeper
- Storm

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.218 ALM-26052 Number of Available Supervisors of the Storm Service Is Less Than the Threshold

Description

The system periodically checks the number of available Supervisors every 60 seconds and compares the number of available Supervisors with the threshold. This alarm is generated when the number of available Supervisors is less than the threshold.

You can change the threshold in **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster*.

This alarm is cleared when the number of available Supervisors is greater than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26052	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Existing tasks in the cluster cannot be performed. The cluster can receive new Storm tasks, but cannot perform these tasks.

Possible Causes

The status of some Supervisors in the cluster is abnormal.

Procedure

Check the Supervisor status.

Step 1 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Supervisor** to go to the Storm service management page.

Step 2 In **Roles**, check whether any instance whose status is **Faulty** or **Restoring** exists.

- If yes, go to **Step 3**.

- If no, go to **Step 5**.

Step 3 Select Supervisor role instances whose status is **Faulty** or **Restoring**, choose **More** > **Restart Instance**, and check whether the instances restart successfully.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Wait for 30 seconds, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

NOTE

Services are interrupted when the Supervisor is being restarted. Then, services are restored after the restarting.

Collect fault information.

Step 5 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 6 Select **Storm** and **ZooKeeper** in the required cluster from the **Service** drop-down list box.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.219 ALM-26053 Storm Slot Usage Exceeds the Threshold

Description

The system checks the slot usage every 60 seconds and compares the actual slot usage with the threshold. This alarm is generated when the slot usage is greater than the threshold.

You can change the threshold in **O&M** > **Alarm** > **Thresholds**.

This alarm is cleared when the slot usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26053	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

New Storm tasks cannot be performed.

Possible Causes

- The status of some Supervisors in the cluster is abnormal.
- The status of all Supervisors is normal, but the processing capability is insufficient.

Procedure

Check the Supervisor status.

Step 1 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Instance** to go to the Storm instance management page.

Step 2 Check whether any instance whose status is **Faulty** or **Restoring** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Select Supervisor role instances whose status is **Faulty** or **Restoring**, choose **More** > **Restart Instance**, and check whether the instances restart successfully.

- If yes, go to **Step 4**.
- If no, go to **Step 10**.

Step 4 Wait several minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Increase the number of slots in each Supervisor.

Step 5 Log in to the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Configurations** > **All Configurations**.

Step 6 Increase the number of ports in the **supervisor.slots.ports** parameter of each Supervisor role and restart the instance.

Step 7 Wait several minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Step 8 Perform capacity expansion for Supervisor.

Step 9 Wait several minutes, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

NOTE

Services are interrupted when the Supervisor is being restarted. Then, services are restored after the restarting.

Collect fault information.

Step 10 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 11 Select **Storm** and **ZooKeeper** in the required cluster from the **Service** drop-down list box.

Step 12 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.220 ALM-26054 Nimbus Heap Memory Usage Exceeds the Threshold

Description

The system checks the heap memory usage of Storm Nimbus every 30 seconds and compares the actual usage with the threshold. The alarm is generated when the heap memory usage of Storm Nimbus exceeds the threshold (80% of the maximum memory by default) for 5 consecutive times.

Users can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Storm** > **Nimbus** to change the threshold.

The alarm is cleared when the heap memory usage is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
26054	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

When the heap memory usage of Storm Nimbus is overhigh, frequent GCs occur. In addition, a memory overflow may occur so that the Yarn service is unavailable.

Possible Causes

The heap memory of the Storm Nimbus instance on the node is overused or the heap memory is inappropriately allocated. As a result, the usage exceeds the threshold.

Procedure

Check the heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Heap Memory Usage of Storm Nimbus Exceeds the Threshold** > **Location**. Check the host name of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the chart area and choose **Customize** > **Nimbus** > **Heap Memory Usage of Nimbus**. Click **OK**.

Step 3 Check whether the used heap memory of Nimbus reaches the threshold (The default value is 80% of the maximum heap memory) specified for Nimbus.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Storm** > **Configurations** > **All Configurations** > **Nimbus** > **System**. Change the value of **-Xmx** in **NIMBUS_GC_OPTS** based on site requirements, and click **Save**. Click **OK**.

NOTE

- You are advised to set **-Xms** and **-Xmx** to the same value to prevent adverse impact on performance when JVM dynamically adjusts the heap memory size.
- The number of Workers grows as the Storm cluster scale increases. You can increase the value of **GC_OPTS** for Nimbus. The recommended value is as follows: If the number of Workers is 20, set **-Xmx** to a value greater than or equal to 1 GB. If the number of Workers exceeds 100, set **-Xmx** to a value greater than or equal to 5 GB.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select the following node in the required cluster from the **Service** drop-down list.

- NodeAgent
- Storm

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.221 ALM-27001 DBService Service Unavailable

Description

The alarm module checks the DBService service status every 30 seconds. This alarm is generated when the system detects that DBService service is unavailable.

This alarm is cleared when DBService service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27001	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The database service is unavailable and cannot provide data import and query functions for upper-layer services, which results in some services exceptions.

Possible Causes

- The floating IP address does not exist.
- There is no active DBServer instance.
- The active and standby DBServer processes are abnormal.

Procedure

Check whether the floating IP address exists in the cluster environment.

Step 1 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService** > **Instance**.

Step 2 Check whether the active instance exists.

- If yes, go to **Step 3**.
- If no, go to **Step 9**.

Step 3 Select the active DBServer instance and record the IP address.

Step 4 Log in to the host that corresponds to the preceding IP address as user **root**, and run the **ifconfig** command to check whether the DBService floating IP address exists on the node.

- If yes, go to **Step 5**.
- If no, go to **Step 9**.

Step 5 Run the **ping** *floatip* command to check whether the DBService floating IP address can be pinged successfully.

- If yes, go to **Step 6**.
- If no, go to **Step 9**.

Step 6 Log in to the host that corresponds to the DBService floating IP address as user **root**, and run the command to delete the floating IP address.

ifconfig *interface* **down**

Step 7 On the FusionInsight Manager home page, choose **Cluster >** *Name of the desired cluster* > **Services** > **DBService** > **More** > **Restart Service** to restart DBService, and check whether DBService is restarted successfully.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Wait for about 2 minutes and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 14**.

Check the status of the active DBServer instance.

Step 9 Select the DBServer instance whose role status is abnormal and record the IP address.

Step 10 On the **Alarm** page, check whether **Process Fault** occurs in the DBServer instance on the host that corresponds to the IP address.

- If yes, go to **Step 11**.

- If no, go to **Step 14**.

Step 11 Handle the alarm according to "ALM-12007 Process Fault".

Step 12 Wait for about 5 minutes and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 19**.

Check the status of the active and standby DBServers.

Step 13 Log in to the host that corresponds to the preceding IP address as user **root**, and run the **su - omm** command to switch to user **omm**.

Step 14 Run the **cd ${DBSERVER_HOME}** command to go to the installation directory of the DBService.

Step 15 Run the **sh sbin/status-dbserver.sh** command to view the status of the active and standby HA processes of DBService. Determine whether the status can be viewed successfully.

```
HAMode
double

NodeName            HostName        HAVersion          StartTime           HAActive
HAAllResOK          HARunPhase
10_5_89_12          host01          V100R001C01        2019-06-13 21:33:09 active
normal              Actived
10_5_89_66          host03          V100R001C01        2019-06-13 21:33:09 standby
normal              Deactived

NodeName            ResName         ResStatus          ResHAStatus         ResType
10_5_89_12          floatip         Normal             Normal              Single_active
10_5_89_12          gaussDB         Active_normal      Normal              Active_standby
10_5_89_66          floatip         Stopped            Normal              Single_active
10_5_89_66          gaussDB         Standby_normal     Normal              Active_standby
```

- If yes, go to **Step 16**.
- If no, go to **Step 19**.

Step 16 Check whether the active and standby HA processes are in the abnormal state.

- If yes, go to **Step 17**.
- If no, go to **Step 19**.

Step 17 On FusionInsight Manager, choose **Cluster >** *Name of the desired cluster* > **Services** > **DBService** > **More** > **Restart Service** to restart DBService, and check whether the system displays a message indicating that the restart is successful.

- If yes, go to **Step 18**.
- If no, go to **Step 19**.

Step 18 Wait for about 2 minutes and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 19**.

Collect fault information.

Step 19 On FusionInsight Manager, choose **O&M** > **Log > Download**.

Step 20 Select **DBService** in the required cluster and **NodeAgent** from the **Service**.

Step 21 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 22 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.222 ALM-27003 DBService Heartbeat Interruption Between the Active and Standby Nodes

Description

This alarm is generated when the active or standby DBService node does not receive heartbeat messages from the peer node for 7 seconds.

This alarm is cleared when the heartbeat recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27003	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local DBService HA Name	Specifies a local DBService HA.
Peer DBService HA Name	Specifies a peer DBService HA.

Impact on the System

During the DBService heartbeat interruption, only one node can provide the service. If this node is faulty, no standby node is available for failover and the service is unavailable.

Possible Causes

The link between the active and standby DBService nodes is abnormal.

Procedure

Check whether the network between the active DBService server and the standby DBService server is normal.

Step 1 In the alarm list on FusionInsight Manager, click ⌄ in the row where the alarm is located in the real-time alarm list and view the standby DBService server address.

Step 2 Log in to the active DBService server as user **root**.

Step 3 Run the **ping** *standby DBService heartbeat IP address* command to check whether the standby DBService server is reachable.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select the following nodes in the required cluster from the **Service**:

- DBService
- Controller
- NodeAgent

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.223 ALM-27004 Data Inconsistency Between Active and Standby DBServices

Description

The system checks the data synchronization status between the active and standby DBService every 10 seconds. This alarm is generated when the synchronization status cannot be queried for six consecutive times or when the synchronization status is abnormal.

This alarm is cleared when the synchronization status becomes normal.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27004	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Local DBService HA Name	Specifies the HA name of the local DBService.
Peer DBService HA Name	Specifies the HA name of the peer DBService.
SYNC_PERCENT	Specifies the synchronization percentage.

Impact on the System

When data is not synchronized between the active and standby DBServices, data may be lost or abnormal if the active instance becomes abnormal.

Possible Causes

- The network between the active and standby nodes is unstable.
- The standby DBService is abnormal.
- The standby node disk space is full.
- The CPU usage of the GaussDB process on the active DBService node is high. You need to locate the failure cause based on logs.

Procedure

Check whether the network between the active and standby nodes is normal.

Step 1 On FusionInsight Manager, choose **Cluster > Services > DBService > Instance**, check the service IP address of the standby DBServer instance.

Step 2 Log in to the active DBService node as user **root**.

Step 3 Run the **ping** *Standby DBService heartbeat IP address* command to check whether the standby DBService node is reachable.

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the standby DBService is normal.

Step 6 Log in to the standby DBService node as user **root**.

Step 7 Run the **su - omm** command to switch to user **omm**.

Step 8 Go to the **${DBSERVER_HOME}/sbin** directory and run the **./status-dbserver.sh** command to check whether the GaussDB resource status of the standby DBService is normal. In the command output, check whether the following information is displayed in the row where **ResName** is **gaussDB**:

For example:

```
10_10_10_231 gaussDB Standby_normal Normal Active_standby
```

- If yes, go to **Step 9**.
- If no, go to **Step 16**.

Check whether the standby node disk space is full. (Skip this check for versions later than MRS 3.1.2.)

Step 9 Log in to the standby DBService node as user **root**.

Step 10 Run the **su - omm** command to switch to user **omm**.

Step 11 Go to the **${DBSERVER_HOME}** directory, and run the following commands to obtain the DBService data directory:

cd ${DBSERVER_HOME}

source .dbservice_profile

echo ${DBSERVICE_DATA_DIR}

Step 12 Run the **df -h** command to view the system disk partition usage information.

Step 13 Check whether the DBService data directory space is full.

- If yes, go to **Step 14**.
- If no, go to **Step 16**.

Step 14 Expand the disk capacity.

Step 15 After the disk capacity is expanded, wait 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Collect fault information.

Step 16 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 17 In the **Service** area, select **DBService** of the target cluster and **OS**, **OS Statistics**, and **OS Performance** under **OMS**, and click **OK**.

Step 18 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 19 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.224 ALM-27005 Database Connections Usage Exceeds the Threshold

Description

The system checks the usage of the number of database connections of the nodes where DBServer instances are located every 30 seconds and compares the usage

with the threshold. If the usage exceeds the threshold for five consecutive times (this number is configurable, and 5 is the default value), the system generates this alarm. The default usage threshold is 90%, and you can configure it based on site requirements.

The trigger count is configurable. This alarm is cleared in the following scenarios:

- The trigger count is 1, and the usage of the number of database connections is less than or equal to the threshold.
- The trigger count is greater than 1, and the usage of the number of database connections is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
27005	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Upper-layer services may fail to connect to the DBService database, affecting services.

Possible Causes

- Too many database connections are used.
- The maximum number of database connections is improperly configured.
- The alarm threshold or alarm trigger count is improperly configured.

Procedure

Checking whether too many data connections are used

Step 1 On FusionInsight Manager, click DBService in the service list on the left navigation pane. The DBService monitoring page is displayed.

Step 2 Observe the number of connections used by the database user, as shown in **Figure 10-74**. Based on the service scenario, reduce the number of database user connections.

Figure 10-74 Number of connections used by database users

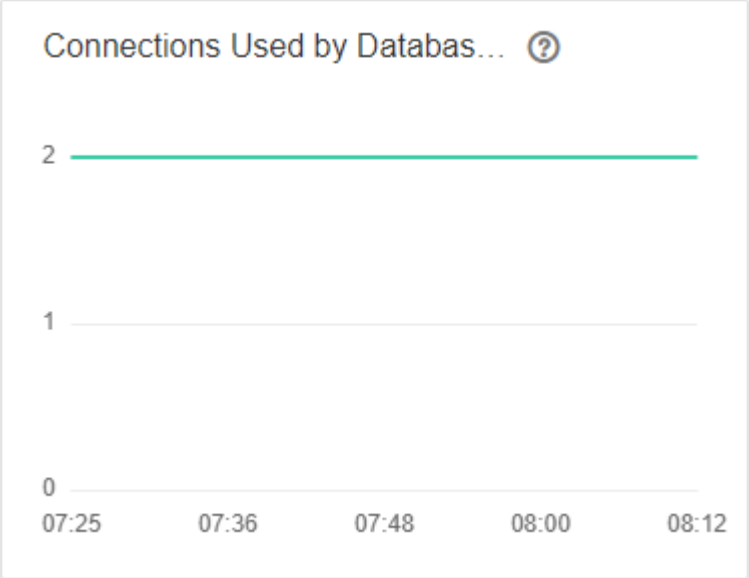

Step 3 Wait for 2 minutes and check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 4**.

Checking whether the maximum number of database connections is properly configured

Step 4 Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService** > **Configurations**. On the displayed page, select the **All Configurations** tab, and increase the maximum number of database connections based on service requirements, as shown in **Figure 10-75**. Click **Save**. In the displayed **Save configuration** dialog box, click **OK**.

Figure 10-75 Setting the maximum number of database connections

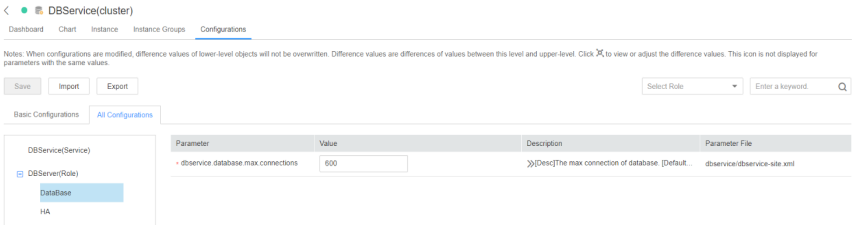

Step 5 After the maximum number of database connections is changed, restart DBService (do not restart the upper-layer services).

Procedure: Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService**. On the displayed page, choose **More** > **Restart Service**. Enter the password of the current login user and click **OK**. Do not select **Restart upper-layer services.**, click **OK**.

Step 6 After the service is restarted, wait for 2 minutes and check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 7**.

Checking whether the alarm threshold or trigger count is properly configured

Step 7 Log in to FusionInsight Manager and change the alarm threshold and alarm trigger count based on the actual database connection usage.

Choose **O&M**> **Alarm** > **Thresholds** > *Name of the desired cluster* > **DBService > Database > Database Connections Usage (DBServer)**. In the **Database Connections Usage (DBServer)** area, click the pencil icon next to **Trigger Count**. In the displayed dialog box, change the trigger count, as shown in **Figure 10-76**.

NOTE

Trigger Count: If the usage of the number of database connections exceeds the threshold consecutively for more than the value of this parameter, an alarm is generated.

Figure 10-76 Setting alarm trigger count

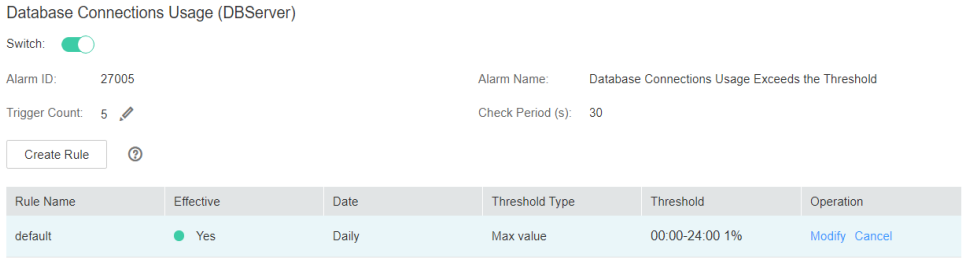

Based on the actual database connection usage, choose **O&M** >**Alarm** > **Thresholds** > *Name of the desired cluster* > **DBService > Database > Database Connections Usage (DBServer)**. In the **Database Connections Usage (DBServer)** area, click **Modify** in the **Operation** column. In the **Modify Rule** dialog box, modify the required parameters and click **OK** as shown in **Figure 10-77**.

Figure 10-77 Set alarm threshold

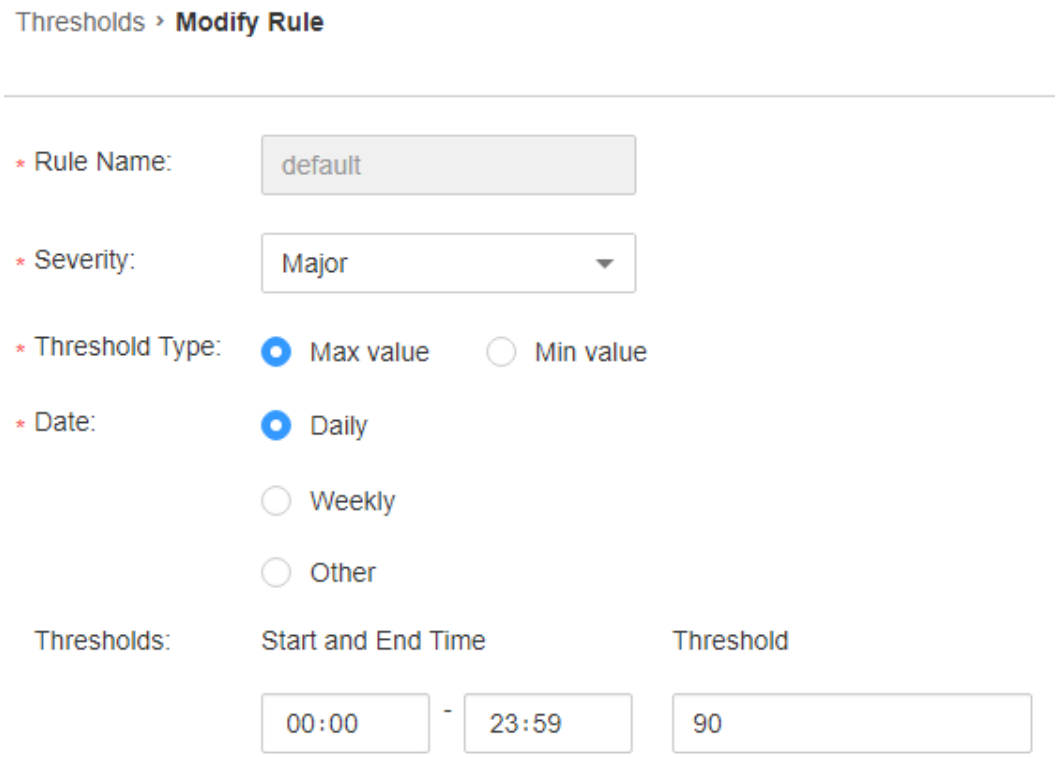

Step 8 Wait for 2 minutes and check whether the alarm is automatically cleared.

- If it is, no further action is required.
- If it is not, go to **Step 9**.

Collect fault information

Step 9 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 10 Select **DBService** in the required cluster from the **Service**.

Step 11 Specify the host for collecting logs by setting the **Host** parameter that is optional. By default, all hosts are selected.

Step 12 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.225 ALM-27006 Disk Space Usage of the Data Directory Exceeds the Threshold

Description

The system checks the disk space usage of the data directory on the active DBServer node every 30 seconds and compares the disk usage with the threshold. The alarm is generated when the disk space usage exceeds the threshold for five consecutive times (the default value). The number of consecutive times is configurable. The disk space usage threshold of the data directory is set to 80% by default, which is configurable as well.

The value of **hit number** is configurable. When the value is set to **1** and the disk space usage is lower than or equal to the threshold, the alarm is cleared. When the value is greater than 1 and the disk space usage is lower than 90% of the threshold, the alarm is cleared.

Attribute

Alarm ID	Alarm Severity	Auto Clear
27006	Major	Yes

Parameters

Name	Meaning
ClusterName	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition where the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the actual indicator value exceeds this threshold, the alarm is generated.

Impact on the System

- Service processes become unavailable.
- When the disk space usage of the data directory exceeds 90%, the database reports the "Database Enters the Read-Only Mode" alarm and enters the read-only mode, which may cause service data loss.

Possible Causes

- The alarm threshold is improperly configured.
- The data volume of the database is too large or the disk configuration cannot meet service requirements, causing excessive disk usage.

Procedure

Check whether the threshold is set properly.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **DBService** > **Database** > **Disk Space Usage of the Data Directory** to check whether the alarm threshold is proper (the default value 80% is a proper value).

- If yes, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 Change the alarm threshold based on the actual service situation.

Step 3 Choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService**. On the **Dashboard** page, view the **Disk Space Usage of the Data Directory** chart and check whether the disk space usage of the data directory is lower than the threshold.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Figure 10-78 Disk Space Usage of the Data Directory

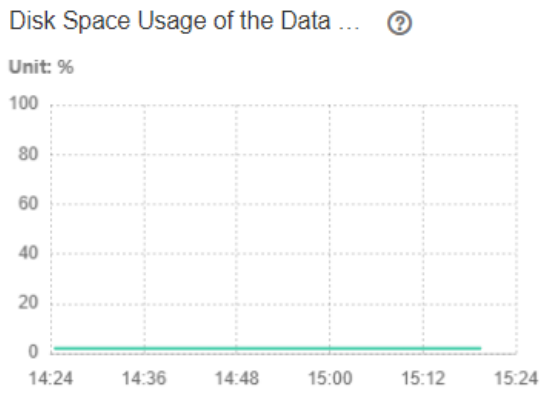

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether large files are incorrectly written into the disk.

Step 5 Log in to the active DBService node as user **omm**.

Step 6 Run the following commands to view the files whose size exceeds 500 MB in the data directory and check whether there are large files incorrectly written into the directory:

source $DBSERVER_HOME/.dbservice_profile

find "$DBSERVICE_DATA_DIR"/../ -type f -size +500M

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 Handle the large files based on the actual scenario and check whether the alarm is cleared 2 minutes later.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **DBService** for the target cluster.

Step 10 Specify the host for collecting logs by setting the **Host** parameter which is optional. By default, all hosts are selected.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.226 ALM-27007 Database Enters the Read-Only Mode

Description

The system checks the disk space usage of the data directory on the active DBServer node every 30 seconds. The alarm is generated when the disk space usage exceeds 90%.

The alarm is cleared when the disk space usage is lower than 80%.

Attribute

Alarm ID	Alarm Severity	Auto Clear
27007	Critical	Yes

Parameters

Name	Meaning
ClusterName	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the actual indicator value exceeds this threshold, the alarm is generated.

Impact on the System

The database enters the read-only mode, causing service data loss.

Possible Causes

The disk configuration cannot meet service requirements. The disk usage reaches the upper limit.

Procedure

Check whether the disk space usage reaches the upper limit.

Step 1 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService**.

Step 2 On the **Dashboard** page, view the **Disk Space Usage of the Data Directory** chart and check whether the disk space usage of the data directory exceeds 90%.

- If yes, go to **Step 3**.
- If no, go to **Step 13**.

Step 3 Log in to the active management node of the DBServer as user **omm** and run the following commands to check whether the database enters the read-only mode:

source $DBSERVER_HOME/.dbservice_profile

gsql -U omm -W *password* **-d postgres -p 20051**

show default_transaction_read_only;

NOTE

In the preceding commands, *password* indicates the password of user **omm** of the DBService database (You can view the initial password of user omm in User **User Account List**.). You can run the **\q** command to exit the database.

Check whether the value of **default_transaction_read_only** is **on**.

```
POSTGRES=# show default_transaction_read_only;
 default_transaction_read_only
-------------------------------
 on
(1 row)
```

- If yes, go to **Step 4**.
- If no, go to **Step 13**.

Step 4 Run the following commands to open the **dbservice.properties** file:

source $DBSERVER_HOME/.dbservice_profile

vi ${DBSERVICE_SOFTWARE_DIR}/tools/dbservice.properties

Step 5 Change the value of **gaussdb_readonly_auto** to **OFF**.

Step 6 Run the following command to open the **postgresql.conf** file:

vi ${DBSERVICE_DATA_DIR}/postgresql.conf

Step 7 Delete **default_transaction_read_only = on**.

Step 8 Run the following command for the configuration to take effect:

gs_ctl reload -D ${DBSERVICE_DATA_DIR}

Step 9 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. On the right of the alarm "Database Enters the Read-Only Mode", click **Clear** in the **Operation** column. In the dialog box that is displayed, click **OK** to manually clear the alarm.

Step 10 Log in to the active management node of the DBServer as user **omm** and run the following commands to view the files whose size exceeds 500 MB in the data directory and check whether there are large files incorrectly written into the directory:

source $DBSERVER_HOME/.dbservice_profile

find "$DBSERVICE_DATA_DIR"/../ -type f -size +500M

- If yes, go to **Step 11**.
- If no, go to **Step 13**.

Step 11 Handle the files that are incorrectly written into the directory based on the actual scenario.

Step 12 Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **DBService**. On the **Dashboard** page, view the **Disk Space Usage of the Data Directory** chart and check whether the disk space usage is lower than 80%.

- If yes, no further action is required.

- If no, go to **Step 13**.

Figure 10-79 Disk Space Usage of the Data Directory

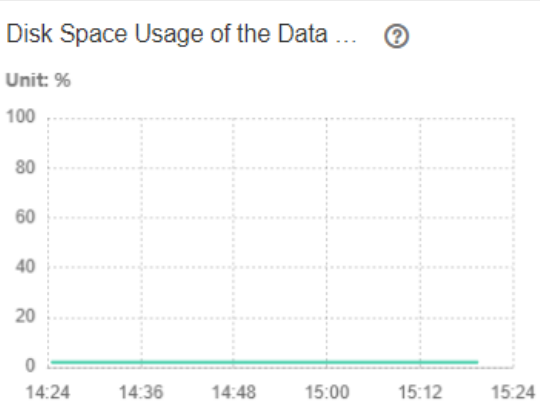

Collect fault information.

Step 13 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 14 Expand the **Service** drop-down list, and select **DBService** for the target cluster.

Step 15 Specify the host for collecting logs by setting the **Host** parameter which is optional. By default, all hosts are selected.

Step 16 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 17 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.227 ALM-29000 Impala Service Unavailable

Alarm Description

The alarm module checks the Impala service status every 30 seconds. This alarm is generated if the Impala service is abnormal.

This alarm is cleared after the Impala service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29000	Critical	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

When the Impala service is abnormal, you cannot perform cluster operations on Impala through FusionInsight Manager. The Impala service functions are unavailable.

Possible Causes

- The Hive service is abnormal.
- The KrbServer service is abnormal.
- The Impala process is abnormal.

Handling Procedure

Check whether the services on which Impala depends are normal.

Step 1 On FusionInsight Manager, choose **Cluster** > **Services** to check whether Hive and KrbServer are stopped.

- If yes, start the stopped services and go to **Step 2**.
- If no, go to **Step 3**.

Step 2 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the alarm list, check whether the **Impala Service Unavailable** alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the alarm list, check whether ALM-16004 Hive Service Unavailable and ALM-25500 KrbServer Service Unavailable exist.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Rectify the fault by following the handling procedure of **ALM-16004 Hive Service Unavailable** or **ALM-25500 KrbServer Service Unavailable**. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the Impala process is normal.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Check whether ALM-12007 Process Fault exists in the alarm list.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Rectify the fault by referring to the handling method of **ALM-12007 Process Fault**, and then check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

Related Information

None

10.228 ALM-29004 Impalad Process Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the Impalad process every 30 seconds. This alarm is generated when the system detects that the memory usage exceeds the default threshold (80%).

This alarm is automatically cleared when the system detects that the memory usage of the process falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29004	Minor	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The memory usage is too high. Some query tasks may fail due to insufficient memory.

Possible Causes

The Impalad process is executing a large number of query tasks.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **CPU and Memory** > **Impalad Process Memory Usage (Impalad)** and check the threshold.

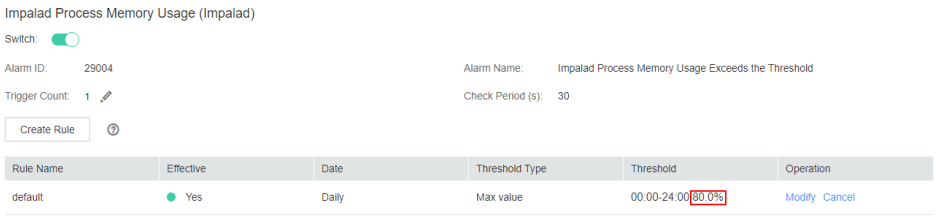

Step 2 If the alarm threshold is smaller than 80%, increase the alarm threshold as required and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 If the threshold is greater than 80%, check whether a large number of concurrent query tasks exist when the alarm is generated. A large number of concurrent query tasks will cause the memory usage to increase sharply. After the tasks are complete, check whether the alarm is automatically cleared. During this period, some tasks may fail to be executed or may be canceled due to insufficient memory. In this case, try again.

NOTE

If the memory usage always exceeds the threshold, the cluster capacity needs to be expanded.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

The alarm is automatically cleared after the burst concurrent tasks are complete.

Related Information

None

10.229 ALM-29005 Number of JDBC Connections to Impalad Exceeds the Threshold

Alarm Description

The system checks the number of client connections to the Impalad node every 30 seconds. This alarm is generated when the number of client connections exceeds the customized threshold (60 by default).

This alarm is automatically cleared when the number of client connections is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29005	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

New client connections may be blocked or even fail.

Possible Causes

The number of client connections maintained by the Impalad service is too large or the threshold is too small.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Connections** > **Number of JDBC Connections to Impalad Process** to check the configured threshold.

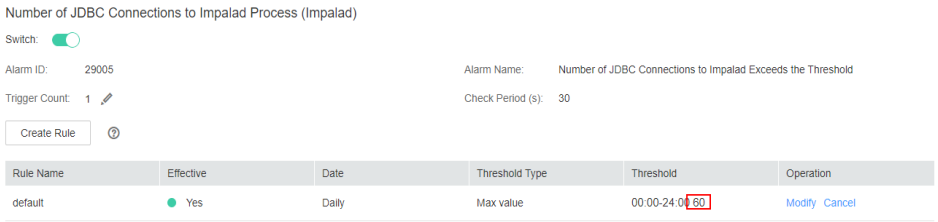

Step 2 Check the number of JDBC applications connected to Impalad and stop idle applications. Check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 3** to change the number of concurrent client connections.

Step 3 On FusionInsight Manager, choose **Cluster** > **Impala** > **Configurations** > **All Configurations** > **Impalad** > **Customization**. Add the customized parameter **--fe_service_threads**. The default value of this parameter is **64**. Change the value as required and click **Save**.

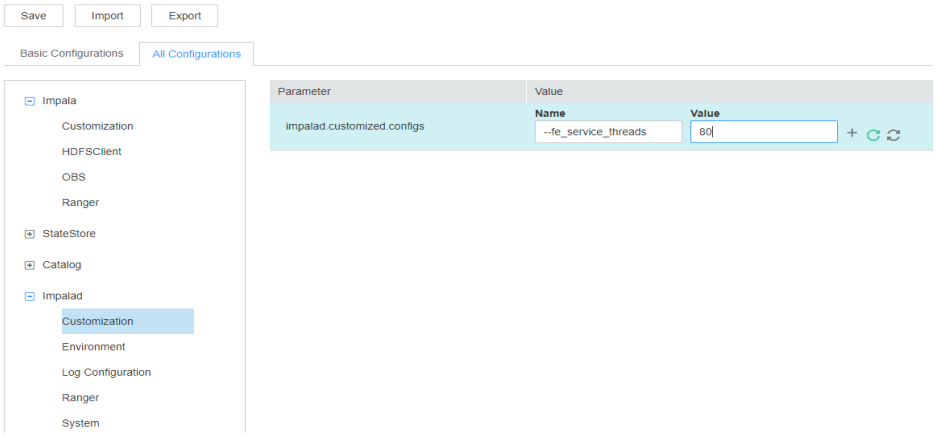

Step 4 After the query tasks on all clients are complete, click the **Instance** tab. Select all Impalad instances, and restart them.

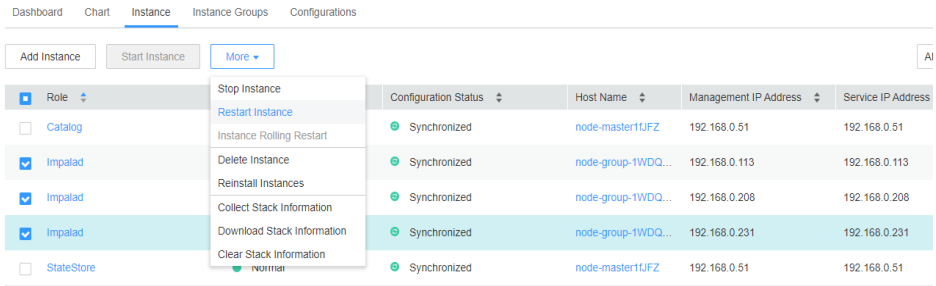

Step 5 After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

Related Information

None

10.230 ALM-29006 Number of ODBC Connections to Impalad Exceeds the Threshold

Alarm Description

The system checks the number of client connections to the Impalad node every 30 seconds. This alarm is generated when the number of client connections exceeds the customized threshold (60 by default).

This alarm is automatically cleared when the number of client connections is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29006	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.

Type	Parameter	Description
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

New client connections may be blocked or even fail.

Possible Causes

The number of client connections maintained by the Impalad service is too large or the threshold is too small.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Connections** > **Number of ODBC Connections to Impalad Process (Impalad)** to check the threshold.

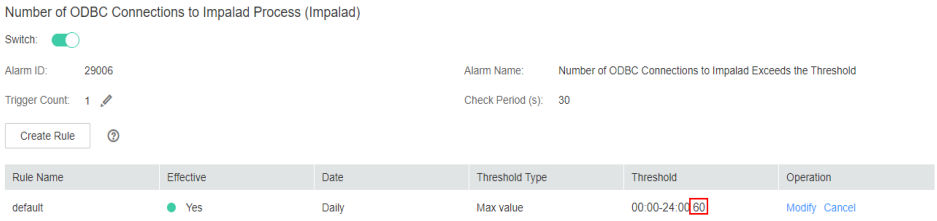

Step 2 Check the number of ODBC applications connected to Impalad and stop idle applications. Check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 3** to change the number of concurrent connections supported by Impalad.

Step 3 On FusionInsight Manager, choose **Cluster** > **Impala** > **Configurations** > **All Configurations** > **Impalad** > **Customization**. Add the customized parameter **--fe_service_threads**. The default value of this parameter is **64**. Change the value as required and click **Save**.

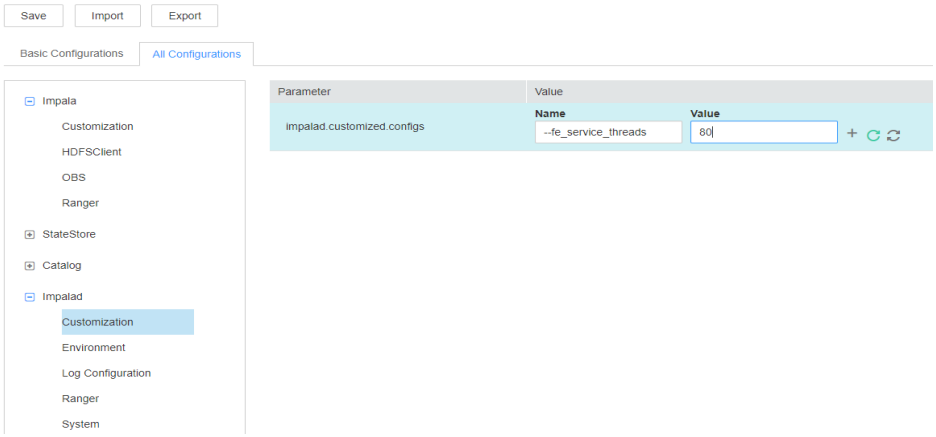

Step 4 After the query tasks on all clients are complete, click the **Instance** tab. Select all Impalad instances, and restart them.

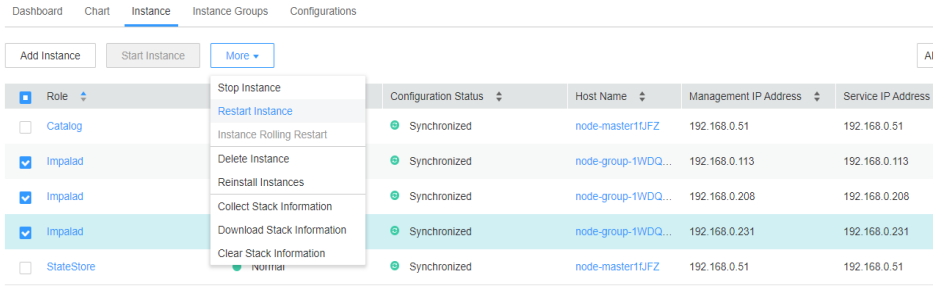

Step 5 After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

Related Information

None

10.231 ALM-29007 Impalad Process Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the Impalad process every 30 seconds. This alarm is generated when the system detects that the memory usage exceeds the default threshold (80%).

This alarm is automatically cleared when the system detects that the memory usage of the process falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29007	Minor	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The memory usage is too high. Some query tasks may fail due to insufficient memory.

Possible Causes

The Impalad process is executing a large number of query tasks.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **CPU and Memory** > **Impalad Process Memory Usage (Impalad)** and check the threshold.

Step 2 If the alarm threshold is smaller than 80%, increase the alarm threshold as required and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 If the threshold is greater than 80%, check whether a large number of concurrent query tasks exist when the alarm is generated. A large number of concurrent query tasks will cause the memory usage to increase sharply. After the tasks are complete, check whether the alarm is automatically cleared. During this period, some tasks may fail to be executed or may be canceled due to insufficient memory. In this case, try again.

NOTE

If the memory usage always exceeds the threshold, the cluster capacity needs to be expanded.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

The alarm is automatically cleared after the burst concurrent tasks are complete.

Related Information

None

10.232 ALM-29008 Number of ODBC Connections to Impalad Exceeds the Threshold

Alarm Description

The system checks the number of client connections to the Impalad node every 30 seconds. This alarm is generated when the number of client connections exceeds the customized threshold (60 by default).

This alarm is automatically cleared when the number of client connections is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29008	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

New client connections may be blocked or even fail.

Possible Causes

The number of client connections maintained by the Impalad service is too large or the threshold is too small.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Connections** > **Number of ODBC Connections to Impalad Process (Impalad)** to check the threshold.

Step 2 Check the number of ODBC applications connected to Impalad and stop idle applications. Check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 3** to change the number of concurrent connections supported by Impalad.

Step 3 On FusionInsight Manager, choose **Cluster** > **Impala** > **Configurations** > **All Configurations** > **Impalad** > **Customization**. Add the custom parameter **--fe_service_threads**. The default value of this parameter is **64**. Change the value as required and click **Save**.

Step 4 After the query tasks on all clients are complete, click the **Instances** tab. Select all Impalad instances, and restart them.

Step 5 After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If yes, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.233 ALM-29010 Number of Queries Being Submitted by Impalad Exceeds the Threshold

Alarm Description

The system checks the total number of queries being submitted by the Impalad node every 60 seconds. This alarm is generated when the number of queries exceeds the customized threshold (150 by default).

This alarm is automatically cleared when the number of queries is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29010	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The queries may be blocked or even fail.

Possible Causes

The Impalad service has maintained a large number of queries, or the threshold is too small.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Query Task Sum Statistics** > **Total number of Queries Being Submitted (Impalad)** and check the threshold.

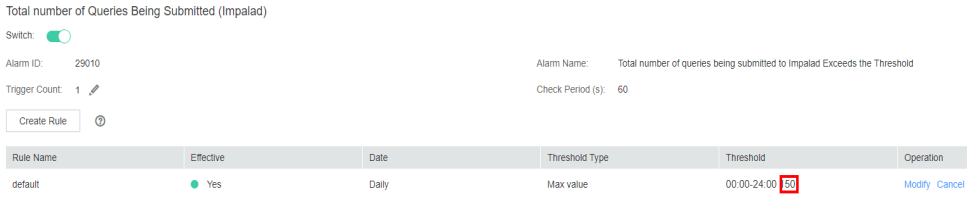

Step 2 Change the threshold.

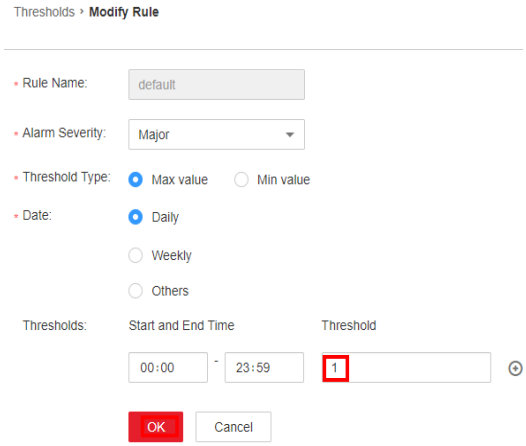

Step 3 Click the **Instances** tab, select all Impalad instances, and restart them.

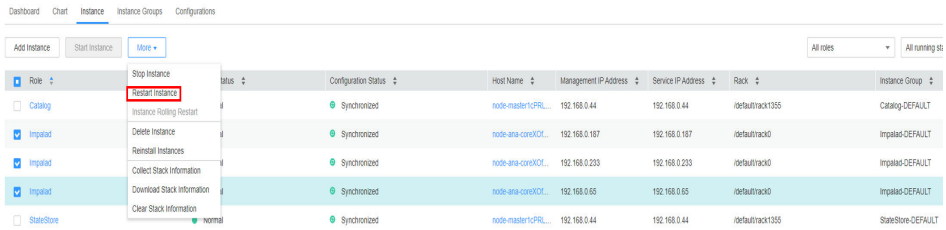

Step 4 After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.234 ALM-29011 Number of Queries Being Executed by Impalad Exceeds the Threshold

Alarm Description

The system checks the total number of queries being executed by the Impalad node every 60 seconds. This alarm is generated when the number of queries exceeds the customized threshold (150 by default).

This alarm is automatically cleared when the number of queries is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29011	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The queries may be blocked or even fail.

Possible Causes

The Impalad service has maintained a large number of queries, or the threshold is too small.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Query Task Sum Statistics** > **Total number of Queries Being Executed (Impalad)** and check the threshold.

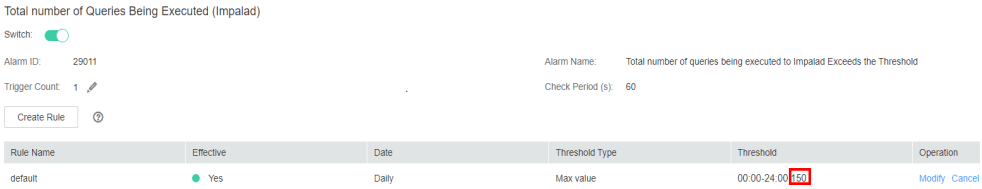

Step 2 Change the threshold.

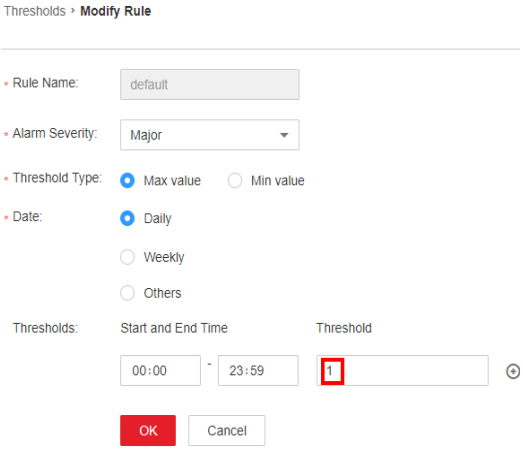

Step 3 Click the **Instances** tab, select all Impalad instances, and restart them.

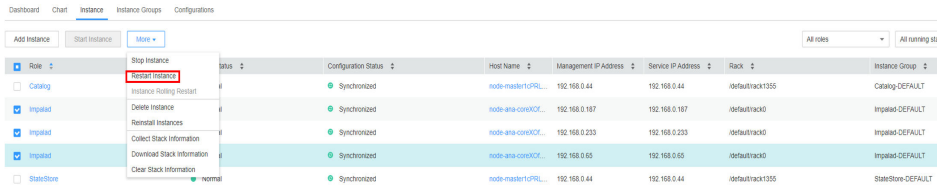

Step 4 After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.235 ALM-29012 Number of Queries Being Waited by Impalad Exceeds the Threshold

Alarm Description

The system checks the total number of queries being waited by the Impalad node every 60 seconds. This alarm is generated when the number of queries exceeds the customized threshold (150 by default).

This alarm is automatically cleared when the number of queries is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29012	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The queries may be blocked or even fail.

Possible Causes

The Impalad service has maintained a large number of queries, or the threshold is too small.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Query Task Sum Statistics** > **Total number of Waiting Queries (Impalad)** and check the threshold.

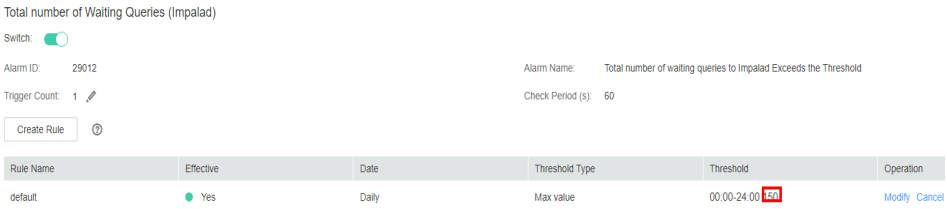

Step 2 Change the threshold.

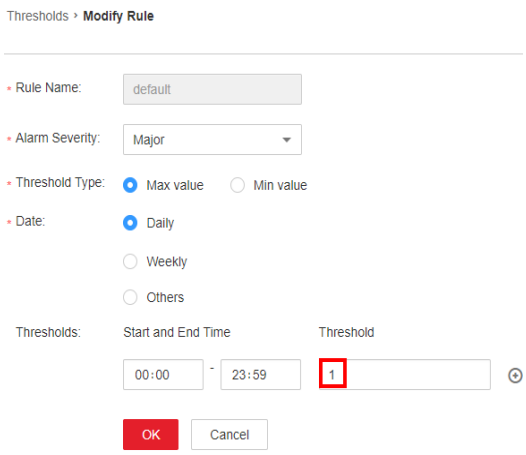

Step 3 Click the **Instances** tab, select all Impalad instances, and restart them.

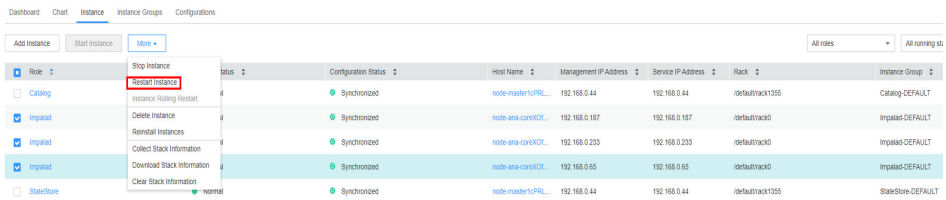

Step 4 After the restart is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.236 ALM-29013 Impalad FGC Time Exceeds the Threshold

Alarm Description

The system checks the FGC time of the Impalad service every 60 seconds. This alarm is generated when the FGC time exceeds the threshold (12 seconds) for five consecutive times. This alarm is cleared when the FGC time is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29013	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Type	Parameter	Description
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Data read and write are affected.

Possible Causes

The memory of the node instance is overused or the heap memory is inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC time.

Step 1 Choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Process FGCT** > **Process FGCT of Impalad (Impalad)**, and check the threshold (12s by default).

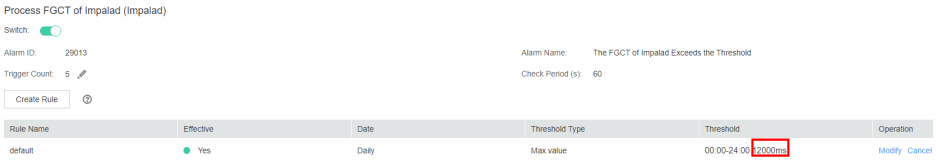

Step 2 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and check whether the alarm whose **Alarm ID** is **29013** exists in the alarm list.

- If yes, go to **3**.
- If no, no further action is required.

Step 3 On FusionInsight Manager, choose **Cluster** > **Impala**, click the **Instances** tab, select the Impalad instance for which the alarm is generated, then click the **Chart** tab, locate the **Process FGCT** chart, and check whether the FGC time is greater than the threshold in **1**.

- If yes, go to **4**.
- If no, go to **Step 5**.

Step 4 Choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **Process FGCT** > **Process FGCT of Impalad (Impalad)**, and change the threshold to a value less than the time obtained in **3**. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.237 ALM-29014 Catalog FGC Time Exceeds the Threshold

Alarm Description

The system checks the FGC time of the Catalog service every 60 seconds. This alarm is generated when the FGC time exceeds the threshold (12 seconds) for five consecutive times. This alarm is cleared when the FGC time is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29014	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Data read and write are affected.

Possible Causes

The memory of the node instance is overused or the heap memory is inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC time.

Step 1 Choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, choose **Impala** > **Process FGCT** > **Process FGCT of Catalog (Catalog)**, and check the threshold (12s by default).

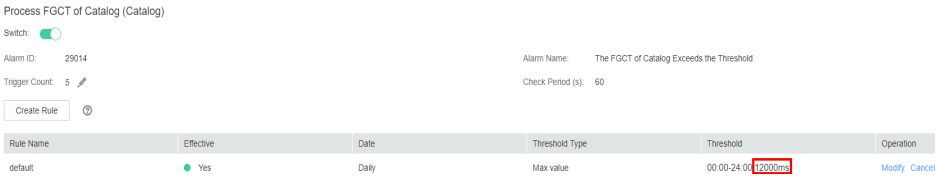

Step 2 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and check whether the alarm whose **Alarm ID** is **29014** exists in the alarm list.

- If yes, go to **3**.
- If no, no further action is required.

Step 3 On FusionInsight Manager, choose **Cluster** > **Impala**, click the **Instance** tab, select the Catalog instance for which the alarm is generated, then click the **Chart** tab, locate the **Process FGCT** chart, and check whether the FGC time is greater than the threshold in **1**.

- If yes, go to **4**.
- If no, go to **5**.

Step 4 Choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, choose **Impala** > **Process FGCT** > **Process FGCT of Catalog (Catalog)**, and change the threshold to a value less than the time obtained in **3**. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **5**.

Collect fault information.

Step 5 On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.238 ALM-29015 Catalog Process Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the Catalog process every 30 seconds. This alarm is generated when the system detects that the memory usage exceeds the default threshold (80%).

This alarm is automatically cleared when the system detects that the memory usage of the process falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29015	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The memory usage is too high. Some query tasks may fail due to insufficient memory.

Possible Causes

The memory of the node instance is overused or the memory is inappropriately configured.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds** > **Impala** > **CPU and Memory** > **Catalog Process Memory Usage (Impalad)** and check the threshold.

Step 2 If the alarm threshold is smaller than 80%, increase the alarm threshold as required and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 If the threshold is greater than 80%, check whether a large number of concurrent query tasks exist when the alarm is generated. A large number of concurrent query tasks will cause the memory usage to increase sharply. After the tasks are complete, check whether the alarm is automatically cleared. During this period, some tasks may fail to be executed or may be canceled due to insufficient memory. In this case, try again.

NOTE

If the memory usage always exceeds the threshold, the cluster capacity needs to be expanded.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.239 ALM-29016 Impalad Instance in the Sub-healthy State

Alarm Description

The system checks every 60 seconds whether Impalad can execute **select 1**. This alarm is generated when the returned result has been incorrect for 20 seconds two consecutive times. This alarm is cleared when the SQL statement is correctly executed within 20 seconds.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29016	Minor	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

Impalad cannot execute SQL statements or SQL statement execution times out, which affects data read and write.

Possible Causes

The Impalad service maintains too many queries.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Impala** > **Impalad Web UI**. On the displayed page, click any node to go to the web UI.

Step 2 On the web UI, click **/backends** to view the Impala instance list. Locate the instance for which the alarm is generated and click **Web UI**. After the web UI of the subhealthy node is displayed, click **/queries** to check the task execution status and check whether any task is executed slowly.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 After the task is complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Impala** > **Instances**, select the Impala instance for which the alarm is generated, click **More**, and select **Restart Instance**. Then, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager of the active or standby cluster, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Impala** for the target cluster.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.240 ALM-29100 Kudu Service Unavailable

Alarm Description

The system checks the Kudu service status every 60 seconds. This alarm is generated when the system detects that all Kudu instances are abnormal and considers that the Kudu service is unavailable.

This alarm is cleared when at least one Kudu instance becomes normal and the system considers that the Kudu instance service is restored.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29100	Critical	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

Users cannot use the Kudu service.

Possible Causes

Some Kudu instances are abnormal.

Handling Procedure

Handle the Kudu instance exceptions.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed page, locate the alarm ALM-29100 Kudu Service Unavailable.

Step 2 In the **Location Information** column, record the host name and role name.

Step 3 Choose **Cluster** > **Services** > **Kudu** > **Instances**. Click the role name for the host name obtained in **Step 2**, view the instance logs, and restore the instance. Then, check whether the alarm is cleared.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 6 In the **Service** area, select the following nodes of the desired cluster.

- Kudu

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

Related Information

None

10.241 ALM-29104 Tserver Process Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the Kudu Tserver process every 60 seconds. This alarm is generated when the system detects that the memory usage of the Kudu Tserver process exceeds the threshold.

This alarm is cleared when the memory usage of the Tserver process becomes normal and the system considers that the Kudu instance service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29104	Critical	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.

Type	Parameter	Description
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Users cannot use the Kudu service.

Possible Causes

The memory usage of a KuduTserver instance is too high.

Handling Procedure

Handle the Kudu instance exceptions.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed page, locate alarm **ALM-29104 Tserver Process Memory Usage Exceeds the Threshold** and view the alarm source.

Step 2 Choose **O&M** > **Alarm** > **Thresholds** > **Kudu**. Locate the threshold of this alarm and check whether the Kudu instance memory usage exceeds the threshold. If yes, rectify the fault or change the threshold.

Step 3 Choose **O&M** > **Alarm** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select the following service for the target cluster:

- Kudu

Step 6 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.242 ALM-29106 Tserver Process CPU Usage Exceeds the Threshold

Alarm Description

The system checks the Kudu service status every 60 seconds. This alarm is generated when the system detects that the CPU usage of the Kudu Tserver process is too high.

This alarm is cleared when the CPU usage of the Tserver process becomes normal and the system considers that the Kudu instance service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29106	Critical	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Users cannot use the Kudu service.

Possible Causes

The CPU usage of a KuduTserver instance is too high.

Handling Procedure

Handle the Kudu instance exceptions.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed page, locate alarm **ALM-29106 Tserver Process CPU Usage Exceeds the Threshold** and view the alarm source.

Step 2 Choose **O&M** > **Alarm** > **Thresholds** > **Kudu**. Locate the alarm threshold and check whether the CPU usage of the cluster Kudu instance exceeds the threshold. If yes, rectify the fault or change the threshold.

Step 3 Choose **O&M** > **Alarm** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select the following service for the target cluster:

- Kudu

Step 6 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.243 ALM-29107 Tserver Process Memory Usage Exceeds the Threshold

Alarm Description

The system checks the Kudu service status every 60 seconds. This alarm is generated when the memory usage of the Kudu Tserver process exceeds the threshold.

This alarm is cleared when the memory usage of the Tserver process becomes normal and the system considers that the Kudu instance service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
29107	Critical	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Users cannot use the Kudu service.

Possible Causes

The memory usage of the KuduTserver instance is too high.

Handling Procedure

Handle the Kudu instance exceptions.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed page, locate alarm **ALM-29107 Tserver Process Memory Usage Exceeds the Threshold** and view the alarm source.

Step 2 Choose **O&M** > **Alarm** > **Thresholds** > **Kudu**. Locate the alarm threshold, compare the memory usage of the KuduTserver instance in the cluster with the threshold, and find the node whose memory usage exceeds the threshold.

Add nodes or reschedule jobs to reduce the memory usage of the Tserver node or change the threshold.

Step 3 Choose **O&M** > **Alarm** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select the following service for the target cluster:

- Kudu

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.244 ALM-38000 Kafka Service Unavailable

Description

The system checks the Kafka service status every 30 seconds. This alarm is generated when the Kafka service is unavailable.

This alarm is cleared when the Kafka service recovers.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38000	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The cluster cannot provide the Kafka service, and users cannot perform new Kafka tasks.

Possible Causes

- The KrbServer service is abnormal.(Skip this step if the normal mode is used.)
- The ZooKeeper service is abnormal or does not respond.
- The Broker instance in the Kafka cluster are abnormal.

Procedure

Check the status of the KrbServer service. (Skip this step if the normal mode is used.)

Step 1 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services > KrbServer**.

Step 2 Check whether the running status of the KrbServer service is **Normal**.

- If yes, go to **Step 5**.
- If no, go to **Step 3**.

Step 3 Rectify the fault by following the steps provided in **ALM-25500 KrbServer Service Unavailable**.

Step 4 Perform **Step 2** again.

Check the status of the ZooKeeper cluster.

Step 5 Check whether the running status of the ZooKeeper service is **Normal**.

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 If ZooKeeper service is stopped, start it, else rectify the fault by following the steps provided in **ALM-13000 ZooKeeper Service Unavailable**.

Step 7 Perform **Step 5** again.

Check the Broker status.

Step 8 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance** to go to the Kafka instances page.

Step 9 Check whether all instances in **Roles** are running properly.

- If yes, go to **Step 11**.
- If no, go to **Step 10**.

Step 10 Select all Broker instances, choose **More** > **Restart Instance**, and check whether the instances restart successfully.

- If yes, go to **Step 11**.
- If no, go to **Step 13**.

Step 11 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** to check whether the running status is **Normal**.

- If yes, go to **Step 12**.
- If no, go to **Step 13**.

Step 12 Wait for 30 seconds and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Collecting Fault Information

Step 13 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 14 Select **Kafka** in the required cluster from the **Service** drop-down list.

Step 15 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.245 ALM-38001 Insufficient Kafka Disk Capacity

Description

The system checks the Kafka disk usage every 60 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold. This alarm is generated when the disk usage is greater than the threshold.

You can change the threshold in **O&M** > **Alarm** > **Thresholds**. Under the service list, choose **Kafka > Disk > Broker Disk Usage (Broker)** and change the threshold.

When the **Trigger Count** is 1, this alarm is cleared when the Kafka disk usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this

alarm is cleared when the Kafka disk usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38001	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition where the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Kafka data write operations are affected.

Possible Causes

- The configuration (such as number and size) of the disks for storing Kafka data cannot meet the requirement of the current service traffic, due to which the disk usage reaches the upper limit.
- Data retention time is too long, due to which the data disk usage reaches the upper limit.
- The service plan does not distribute data evenly, due to which the usage of some disks reaches the upper limit.

Procedure

Check the disk configuration of Kafka data.

Step 1 On the FusionInsight Manager portal and click **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, locate the alarm and obtain **HostName** from **Location**.

Step 3 Click **Cluster** > *Name of the desired cluster* > **Hosts**.

Step 4 In the host list, click the host name obtained in **Step 2**.

Step 5 Check whether the **Disk** area contains the partition name in the alarm.

- If yes, go to **Step 6**.
- If no, manually clear the alarm and no further operation is required.

Step 6 Check whether the disk partition usage contained in the alarm reaches 100% in the **Disk** area.

- If yes, handle the alarm by following the instructions in **Related Information**.
- If no, go to **Step 7**.

Check the Kafka data storage duration.

Step 7 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations**.

Step 8 Check whether the value of parameter **disk.adapter.enable** is set to **true**.

- If yes, go to **Step 10**.
- If no, go to **Step 9**.

Step 9 Set the value of **disk.adapter.enable** to **true**. Check whether the value of **adapter.topic.min.retention.hours** is properly set.

- If yes, go to **Step 10**.
- If no, adjust the data retention period based on service requirements.

NOTICE

If the disk auto-adaptation function is enabled, some historical data of specified topics is deleted. If the retention period of some topics cannot be adjusted, click **All Configurations** and add the topics to the value of the **disk.adapter.topic.blacklist** parameter.

Step 10 Wait 10 minutes and check whether the usage of faulty disks reduces.

- If yes, wait until the alarm is cleared.
- If no, go to **Step 11**.

Check the Kafka data plan.

Step 11 In the **Instance** area, click **Broker**. In the **Real Time** area of Broker, Click the drop-down menu in the Chart area and choose **Customize** to customize monitoring items.

Step 12 In the dialog box, select **Disk** > **Broker Disk Usage** and click **OK**.

The Kafka disk usage information is displayed.

Figure 10-80 Broker Disk Usage

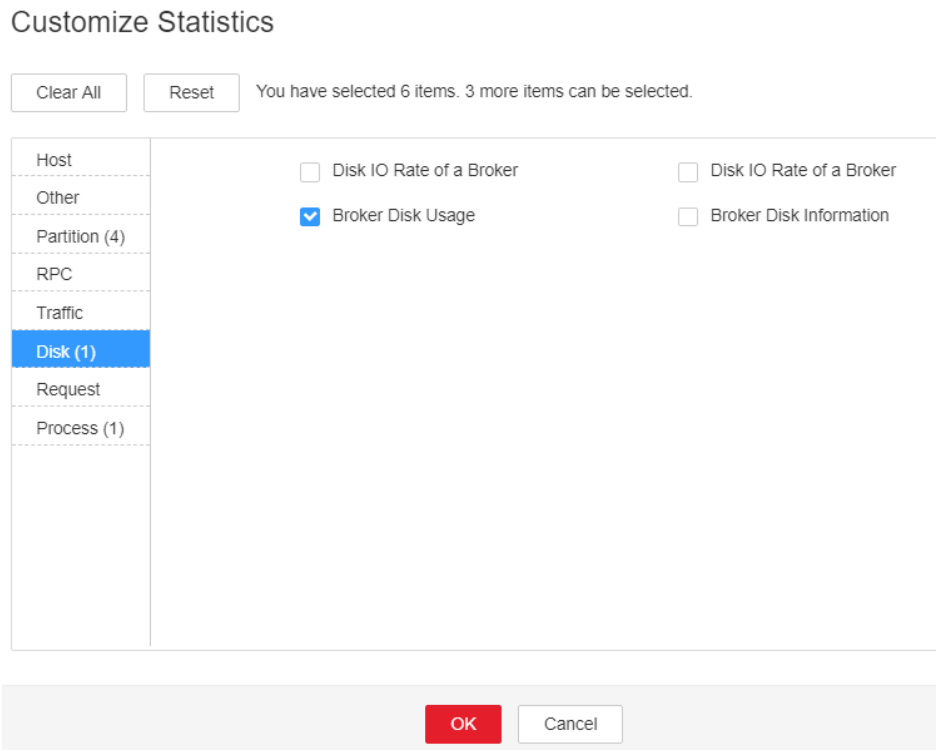

Step 13 View the information in **Step 12** to check whether there is only the disk parathion for which the alarm is generated in **Step 2**.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 Perform disk planning and mount a new disk again. Go to the **Instance Configurations** page of the node for which the alarm is generated, modify **log.dirs**, add other disk directories, and restart the Kafka instance.

Step 15 Determine whether to shorten the data retention time configured on Kafka based on service requirements and service traffic.

- If yes, go to **Step 16**.
- If no, go to **Step 17**.

Step 16 Log in to FusionInsight Manager, select **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations**, and click **All Configurations**. In the search box on the right, enter **log.retention.hours**. The value of the parameter indicates the default data retention time of the topic. You can change the value to a smaller one.

NOTE

- For a topic whose data retention time is configured alone, the modification of the data retention time on the Kafka Service Configuration page does not take effect.
- To modify the data retention time for a topic, use the Kafka client command-line interface (CLI) to configure the topic.

  Example: **kafka-topics.sh --zookeeper "***ZooKeeper IP address***:2181/kafka" --alter --topic "***Topic bane***" --config retention.ms= "***retention time***"**

Step 17 Check whether the usage of some disks reaches the upper limit due to unreasonable configuration of the partitions of some topics. For example, the number of partitions configured for a topic with large data volume is smaller than the number of disks. In this case, the data is not evenly allocated to disks.

NOTE

If you do not know which topics have a large amount of service data, perform the following steps:

1. Log in to an instance node based on the host node information obtained in **Step 2**.
2. Go to the data directory (directory specified by **log.dirs** before the modification in **Step 14**).
3. Run the following command to check whether there is topic with partition that use large disk space.

   du -h --max-depth=1 ./

- If yes, go to **Step 18**.
- If no, go to **Step 19**.

Step 18 In the Kafka client CLI, run the following command to perform partition capacity expansion for the topic:

kafka-topics.sh --zookeeper "*ZooKeeper IP address***:2181/kafka" --alter --topic "***Topic name***" --partitions="***New number of partitions***"**

NOTE

- You are advised to set the new number of partitions to a multiple of the number of Kafka data disks.
- The step may not quickly clear the alarm, and you need to modify the data retention time in **Step 11** to gradually balance data allocation.

Step 19 Determine whether to perform capacity expansion.

NOTE

You are advised to perform capacity expansion for Kafka when the current disk usage exceeds 80%.

- If yes, go to **Step 20**.
- If no, go to **Step 21**.

Step 20 Expand the disk capacity and check whether the alarm is cleared after capacity expansion.

- If yes, no further action is required.
- If no, go to **Step 22**.

Step 21 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 22**.

Collect fault information.

Step 22 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 23 Select **Kafka** in the required cluster from the **Service** drop-down list.

Step 24 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 25 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

Step 1 Log in to FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**, stop the Broker instance whose status is **Restoring**, record the management IP address of the node where the Broker instance is located, and record **broker.id**. The value can be obtained by using the following method: Click the role name. On the **Configurations** page, select **All Configurations**, and search for the **broker.id** parameter.

Step 2 Log in to the recorded management IP address as user **root**, and run the **df -lh** command to view the mounted directory whose disk usage is 100%, for example, **${BIGDATA_DATA_HOME}/kafka/data1**.

Step 3 Go to the directory, run the **du -sh *** command to view the size of each file in the directory,check whether files other than **kafka-logs** exist, and determine whether these files can be deleted or migrated.

- If yes, go to **Step 8**.
- If no, go to **Step 4**.

Step 4 Go to the **kafka-logs** directory, run the **du -sh *** command, select a partition folder to be moved. The naming rule is **Topic name-Partition ID**. Record the topic and partition.

Step 5 Modify the **recovery-point-offset-checkpoint** and **replication-offset-checkpoint** files in the **kafka-logs** directory in the same way.

1. Decrease the number in the second line in the file. (To remove multiple directories, the number deducted is equal to the number of files to be removed.)
2. Delete the line of the to-be-removed partition. (The line structure is "Topic name Partition ID Offset". Save the data before deletion. Subsequently, the content must be added to the file of the same name in the destination directory.)

Step 6 Modify the **recovery-point-offset-checkpoint** and **replication-offset-checkpoint** files in the destination data directory. For example, **${BIGDATA_DATA_HOME}/kafka/data2/kafka-logs** in the same way.

- Increase the number in the second line in the file. (To move multiple directories, the number added is equal to the number of files to be moved.)
- Add the to-be moved partition to the end of the file. (The line structure is "Topic name Partition ID Offset". You can copy the line data saved in **Step 5**.)

Step 7 Move the partition to the destination directory. After the partition is moved, run the **chown omm:wheel -R** *Partition directory* command to modify the directory owner group for the partition.

Step 8 Log in to FusionInsight Manager and choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance** to start the Broker instance.

Step 9 Wait for 5 to 10 minutes and check whether the health status of the Broker instance is **Normal**.

- If yes, resolve the disk capacity insufficiency problem according to the handling method of "ALM-38001 Insufficient Kafka Disk Space" after the alarm is cleared.
- If no, contact the O&M personnel.

----End

10.246 ALM-38002 Kafka Heap Memory Usage Exceeds the Threshold

Description

The system checks the Kafka service status every 30 seconds. The alarm is generated when the heap memory usage of a Kafka instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the heap memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the heap memory usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38002	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.

Name	Meaning
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available Kafka heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the Kafka instance is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Kafka Heap Memory Usage Exceeds the Threshold** > **Location**. Check the host name of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Process** > **Heap Memory Usage of Kafka**, and click **OK**.

Figure 10-81 Heap Memory Usage of Kafka

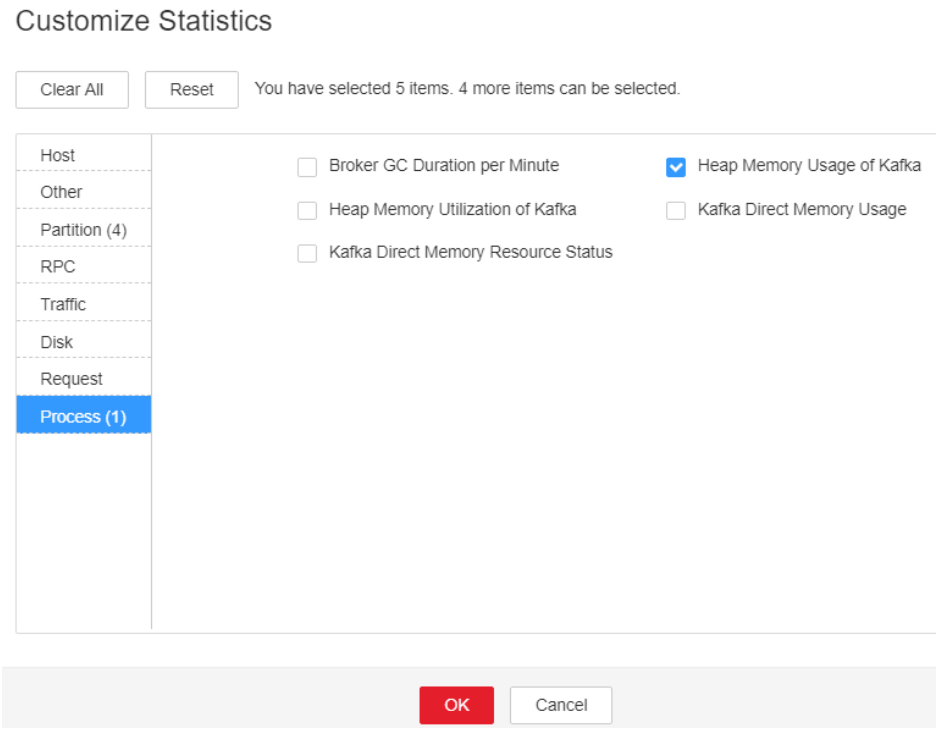

Step 3 Check whether the used heap memory of Kafka reaches 95% of the maximum heap memory specified for Kafka.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Check the heap memory size configured for Kafka.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations**> **Broker(Role)** > **Environment**. Increase the value of **KAFKA_HEAP_OPTS** by referring to the Note.

Figure 10-82 KAFKA_HEAP_OPTS

NOTE

- It is recommended that **-Xmx** and **-Xms** be set to the same value.
- You are advised to view **Heap Memory Usage of Kafka** by referring to **Step 2**, and set the value of **KAFKA_HEAP_OPTS** to twice the value of **Heap Memory Used by Kafka.**

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 7 Select **Kafka** in the required cluster from the **Service** drop-down list.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.247 ALM-38004 Kafka Direct Memory Usage Exceeds the Threshold

Description

The system checks the direct memory usage of the Kafka service every 30 seconds. This alarm is generated when the direct memory usage of a Kafka instance exceeds the threshold (80% of the maximum memory) for 10 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the direct memory usage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the direct memory usage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38004	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available direct memory of the Kafka service is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the Kafka instance is overused or the direct memory is inappropriately allocated.

Procedure

Check the direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Kafka Direct Memory Usage Exceeds the Threshold** > **Location** to check the host name of the instance for which the alarm is generated.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down menu in the Chart area and choose **Customize** > **Process** > **Kafka Direct Memory Usage**, and click **OK**.

Figure 10-83 Kafka Direct Memory Usage

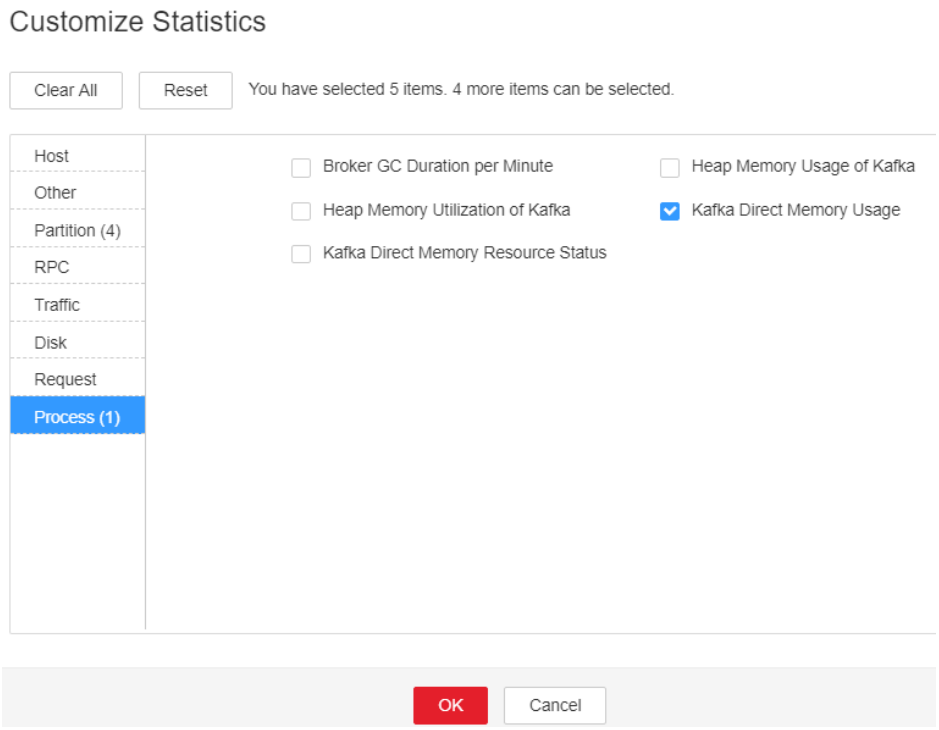

Step 3 Check whether the used direct memory of Kafka reaches 80% of the maximum direct memory specified for Kafka.

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Check the direct memory size configured for the Kafka.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations** > **Broker(Role)** > **Environment** to increase the value of **-Xmx** configured in the **KAFKA_HEAP_OPTS** parameter by referring to the Note.

NOTE

- It is recommended that **-Xmx** and **-Xms** be set to the same value.
- You are advised to view **Kafka Direct Memory Usage** by referring to **Step 2**, and set the value of **KAFKA_HEAP_OPTS** to twice the value of **Direct Memory Used by Kafka.**

Step 5 Save the configuration and restart the Kafka service.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **Kafka** in the required cluster from the **Service** drop-down list.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.248 ALM-38005 GC Duration of the Broker Process Exceeds the Threshold

Description

The system checks the garbage collection (GC) duration of the Broker process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default) for 3 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the GC duration is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the GC duration is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38005	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

A long GC duration of the Broker process may interrupt the services.

Possible Causes

The Kafka GC duration of the node is too long or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC duration.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **GC Duration of the Broker Process Exceeds the Threshold** > **Location**. Check the host name of the instance involved in this alarm.

Step 2 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Process** > **Broker GC Duration per Minute**, and click **OK**.

Figure 10-84 Broker GC Duration per Minute

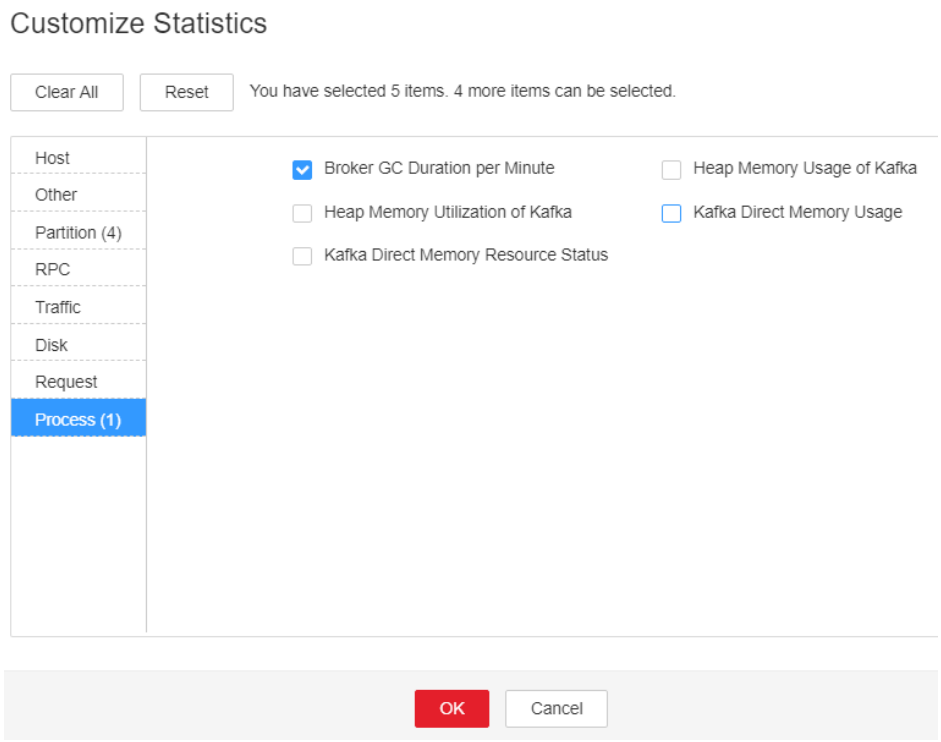

Step 3 Check whether the GC duration of the Broker process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 7**.

Check the direct memory size configured for the Kafka.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations** > **Broker(Role)** > **Environment** to increase the value of **-Xmx** configured in the **KAFKA_HEAP_OPTS** parameter by referring to the Note.

NOTE

- It is recommended that **-Xmx** and **-Xms** be set to the same value.
- You are advised to set the value of **KAFKA_HEAP_OPTS** to twice the value of **Direct Memory Used by Kafka.**

  On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **Process** > **Kafka Direct Memory Resource Status** to check the value of **Direct Memory Used by Kafka**.

Step 5 Save the configuration and restart the Kafka service.

Step 6 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **Kafka** in the required cluster from the **Service** drop-down list.

Step 9 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.249 ALM-38006 Percentage of Kafka Partitions That Are Not Completely Synchronized Exceeds the Threshold

Description

The system checks the percentage of Kafka partitions that are not completely synchronized to the total number of partitions every 60 seconds. This alarm is generated when the percentage exceeds the threshold (50% by default) for 3 consecutive times.

When the **Trigger Count** is 1, this alarm is cleared when the percentage is less than or equal to the threshold. When the **Trigger Count** is greater than 1, this alarm is cleared when the percentage is less than or equal to 90% of the threshold.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38006	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.

Name	Meaning
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

Too many Kafka partitions that are not completely synchronized affect service reliability. In addition, data may be lost when leaders are switched.

Possible Causes

Some nodes where the Broker instance resides are abnormal or stop running. As a result, replicas of some partitions in Kafka are out of the in-sync replicas (ISR) set.

Procedure

Check Broker instances.

Step 1 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. The Kafka instances page is displayed.

Step 2 Check whether faulty nodes exist among all Broker nodes.

- If yes, record the host name of the node and go to **Step 3**.
- If no, go to **Step 5**.

Step 3 On the FusionInsight Manager portal, click **O&M** > **Alarm** > **Alarms** to check whether the fault described in **Step 2** exists in the alarm information and handle the alarm based on corresponding methods.

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. The Kafka instances page is displayed.

Step 5 Check whether stopped nodes exist among all Broker instance.

- If yes, go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Select all stopped Broker instances and click **Start Instance**.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 9 Select **Kafka** in the required cluster from the **Service** drop-down list.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.250 ALM-38007 Status of Kafka Default User Is Abnormal

Description

The system checks the default user of Kafka every 60 seconds. This alarm is generated when the system detects that the user status is abnormal.

Trigger Count is set to **1**. This alarm is cleared when the user status becomes normal.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38007	Critical	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Meaning
HostName	Specifies the host name for which the alarm is generated.
Trigger Condition	Specifies the condition that the Kafka default user status is abnormal.

Impact on the System

If the Kafka default user status is abnormal, metadata synchronization between Brokers and interaction between Kafka and ZooKeeper will be affected, affecting service production, consumption, and topic creation and deletion.

Possible Causes

- The Sssd service is abnormal.
- Some Broker instances stop running.

Procedure

Check whether the Sssd service is abnormal.

Step 1 On the FusionInsight Manager portal, choose **O&M** > **Alarm** > **Alarms** > **Status of Kafka Default User Is Abnormal** > **Location** to check the host name of the instance for which the alarm is generated.

Step 2 Find the host information in the alarm information and log in to the host.

Step 3 Run the **id -Gn kafka** command and check whether "No such user" is displayed in the command output.

- If yes, record the host name of the node and go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Alarms**. Check whether there is **Sssd Service Exception** in the alarm information. If there is, handle the alarm based on alarm information.

Check the running status of the Broker instance.

Step 5 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. The Kafka instance page is displayed.

Step 6 Check whether there are stopped nodes on all Broker instances.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 Select all stopped Broker instances and click **Start Instance**.

Step 8 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 10 In the **Service** area, select **Kafka** in the required cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.251 ALM-38008 Abnormal Kafka Data Directory Status

Description

The system checks the Kafka data directory status every 60 seconds. This alarm is generated when the system detects that the status of a data directory is abnormal.

Trigger Count is set to **1**. This alarm is cleared when the data directory status becomes normal.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38008	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host name for which the alarm is generated.
DirName	Specifies the directory name for which the alarm is generated.
Trigger Condition	Specifies the condition that the Kafka data directory status is abnormal.

Impact on the System

If the Kafka data directory status is abnormal, the current replicas of all partitions in the data directory are brought offline, and the data directory status of multiple nodes is abnormal at the same time. As a result, some partitions may become unavailable.

Possible Causes

- The data directory permission is tampered with.
- The disk where the data directory is located is faulty.

Procedure

Check the permission on the faulty data directory.

Step 1 Find the host information in the alarm information and log in to the host.

Step 2 In the alarm information, check whether the data directory and its subdirectories belong to the omm:wheel group.

- If yes, record the host name of the node and go to **Step 4**.
- If no, go to **Step 3**.

Step 3 Restore the owner group of the data directory and its subdirectories to omm:wheel.

- If yes, go to **Step 6**.
- If no, go to **Step 5**.

Check whether the disk where the data directory is located is faulty.

Step 4 In the upper-level directory of the data directory, create and delete files as user **omm**. Check whether data read/write on the disk is normal.

Step 5 Replace or repair the disk where the data directory is located to ensure that data read/write on the disk is normal.

Step 6 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. On the Kafka instance page that is displayed, restart the Broker instance on the host recorded in **Step 2**.

Step 7 After Broker is started, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 9 In the **Service** area, select **Kafka** in the required cluster.

Step 10 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.252 ALM-38009 Busy Broker Disk I/Os (Applicable to Versions Later Than MRS 3.1.0)

NOTE

- This section applies to versions later than MRS 3.1.0.
- If the alarm name is **ALM-38009 Kafka Topic Overload**, handle the alarm by following the instructions provided in **ALM-38009 Kafka Topic Overload (Applicable to MRS 3.1.0 and Earlier Versions)**.

Alarm Description

The system checks the I/O status of each Kafka disk every 60 seconds. This alarm is generated when the disk I/O of a Kafka data directory on a broker exceeds the threshold (80% by default).

Its **Trigger Count** is **3**. This alarm is cleared when the disk I/O is lower than the threshold (80% by default).

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
38009	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
DataDirectoryName	Specifies the name of the Kafka data directory with frequent disk I/Os.

Impact on the System

The disk partition has frequent I/Os. Data may fail to be written to the Kafka topic for which the alarm is generated.

Possible Causes

- There are many replicas configured for the topic.
- The parameter for batch writing producer's messages is inappropriately configured. The service traffic of this topic is too heavy, and the current partition configuration is inappropriate.

Handling Procedure

Check the number of topic replicas.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm, click ⌄, and view the host name in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services** > **Kafka** > **KafkaTopic Monitor**, search for the topic for which the alarm is generated, and check the number of replicas.

Step 3 Reduce the replication factors of the topic (for example, reduce to **3**) if the number of replicas is greater than 3.

Run the following command on the FusionInsight client to replan the replicas of Kafka topics:

kafka-reassign-partitions.sh --zookeeper *{zk_host}:{port}***/kafka --reassignment-json-file** *{manual assignment json file path}* **--execute**

For example:

/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper 10.149.0.90:2181,10.149.0.91:2181,10.149.0.92:2181/kafka --reassignment-json-file expand-cluster-reassignment.json --execute

NOTE

In the **expand-cluster-reassignment.json** file, describe the brokers to which the partitions of the topic are migrated in the following format: {"partitions":[{"topic": "*topicName*","partition": 1,"replicas": [1,2,3] }],"version":1}

Step 4 Observe for a period of time and check whether the alarm is cleared. If the alarm persists, go to **Step 5**.

Check the partition planning of the topic.

Step 5 On the **KafkaTopic Monitor** page, view **Topic Input Traffic** in the **Topic Traffic** area of each topic, obtain the topic with the largest value, and check the partitions of this topic as well as information about the host of these partitions.

Step 6 Log in to the host queried in **Step 5** and run the **iostat -d -x** command to check the **%util** value of each disk.

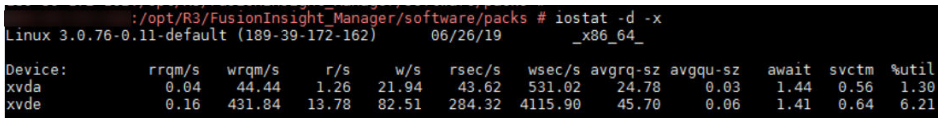

- If the **%util** value of each disk exceeds the threshold (**80%** by default), expand the Kafka disk capacity. After the capacity expansion, replan the topic partitions by referring to **Step 3**.
- If the **%util** values of the disks vary greatly, check the disk partition configuration of Kafka. For example, check the value of **log.dirs** in the **${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/1_14_Broker/etc/server.properties** file.

  Run the following command to view the **Filesystem** information:

  df -h *log.dirs value*

  The command output is as follows.

  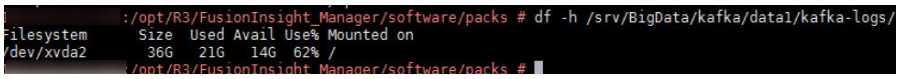

- If the partition where Filesystem is located matches the partition with a high **%util** value, plan Kafka partitions on idle disks, configure **log.dirs** as an idle disk directory, and replan topic partitions by referring to **Step 3**. Ensure that the partitions of the topic are evenly distributed to each disk.

Step 7 Observe for a period of time and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, repeat **Step 5** to **Step 6** three times. Then, go to **Step 8**.

Step 8 Observe for a period of time and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Kafka** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.253 ALM-38009 Kafka Topic Overload (Applicable to MRS 3.1.0 and Earlier Versions)

NOTE

- This section applies to MRS 3.1.0 or earlier.
- If the alarm name is **ALM-38009 Busy Broker Disk I/Os**, handle the alarm by following the instructions provided in **ALM-38009 Busy Broker Disk I/Os (Applicable to Versions Later Than MRS 3.1.0)**.

Alarm Description

The system checks the overload status of each Kafka topic every 60 seconds. This alarm is generated when the percentage of partitions of a topic on the overloaded disk exceeds the threshold (40% by default).

Its **Trigger Count** is **1**. This alarm is cleared when the percentage of partitions of a topic on the overloaded disk is lower than the threshold (40% by default).

An overloaded disk refers to the disk whose I/O usage of a disk partition is greater than 80%.

For example:

The partitions of Topic A are distributed on three brokers. The I/O usages of the disk partitions on two brokers are greater than 80%.

The percentage of partitions on the overloaded disk is 2/3, greater than 40%, this alarm is generated.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
38009	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
TopicName	Specifies the Kafka topic for which the alarm was generated.

Impact on the System

The disk partition has frequent I/Os. Data may fail to be written to the Kafka topic for which the alarm is generated.

Possible Causes

- There are many replicas configured for the topic.
- The parameter for batch writing producer's messages is inappropriately configured. The service traffic of this topic is too heavy, and the current partition configuration is inappropriate.

Handling Procedure

Check the number of topic replicas.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm, click ⌄, and view the host name in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, choose **Services** > **Kafka** > **KafkaTopic Monitor**, search for the topic for which the alarm is generated, and check the number of replicas.

Step 3 Reduce the replication factors of the topic (for example, reduce to **3**) if the number of replicas is greater than 3.

Run the following command on the FusionInsight client to replan the replicas of Kafka topics:

kafka-reassign-partitions.sh --zookeeper *{zk_host}:{port}***/kafka --reassignment-json-file** *{manual assignment json file path}* **--execute**

For example:

/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper 10.149.0.90:2181,10.149.0.91:2181,10.149.0.92:2181/kafka --reassignment-json-file expand-cluster-reassignment.json --execute

NOTE

In the **expand-cluster-reassignment.json** file, describe the brokers to which the partitions of the topic are migrated in the following format: {"partitions":[{"topic": "*topicName*","partition": 1,"replicas": [1,2,3] }],"version":1}

Step 4 Observe for a period of time and check whether the alarm is cleared. If the alarm persists, go to **Step 5**.

Check the partition planning of the topic.

Step 5 On the **KafkaTopic Monitor** page, view **Topic Input Traffic** in the **Topic Traffic** area of each topic, obtain the topic with the largest value, and check the partitions of this topic as well as information about the host of these partitions.

Step 6 Log in to the host queried in **Step 5** and run the **iostat -d -x** command to check the **%util** value of each disk.

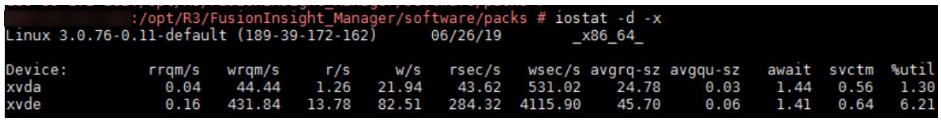

- If the **%util** value of each disk exceeds the threshold (**80%** by default), expand the Kafka disk capacity. After the capacity expansion, replan the topic partitions by referring to **Step 3**.
- If the **%util** values of the disks vary greatly, check the disk partition configuration of Kafka. For example, check the value of **log.dirs** in the **${BIGDATA_HOME}/FusionInsight_HD_8.1.0.1/1_14_Broker/etc/server.properties** file.

  Run the following command to view the **Filesystem** information:

  df -h *log.dirs value*

  The command output is as follows.

  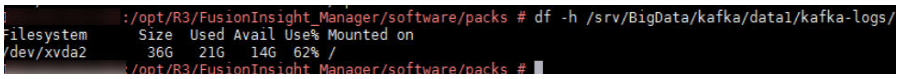

- If the partition where Filesystem is located matches the partition with a high **%util** value, plan Kafka partitions on idle disks, configure **log.dirs** as an idle disk directory, and replan topic partitions by referring to **Step 3**. Ensure that the partitions of the topic are evenly distributed to each disk.

Step 7 Observe for a period of time and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, repeat **Step 5** to **Step 6** three times. Then, go to **Step 8**.

Step 8 Observe for a period of time and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Kafka** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.254 ALM-38010 Topics with Single Replica

Description

The system checks the number of replicas of each topic every 60 seconds on the node where the Kafka Controller resides. This alarm is generated when there is one replica for a topic.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38010	Major	No

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
TopicName	Specifies the list of topics for which the alarm is generated.

Impact on the System

There is the single point of failure (SPOF) risk for topics with only one replica. When the node where the replica resides becomes abnormal, the partition does not have a leader, and services on the topic are affected.

Possible Causes

- The number of replicas for the topic is incorrectly configured.

Procedure

Check the number of replicas for the topic.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, click ⌄ of this alarm, and view the **TopicName** list in **Location**.

Step 2 Check whether replicas need to be added for the topic for which the alarm is generated.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 On the FusionInsight client, re-plan topic replicas and describe the partition distribution of the topic in the **add-replicas-reassignment.json** file in the following format: {"partitions":[{"topic": "*topic name*","partition": 1,"replicas": [1,2] }],"version":1}. Then, run the following command to add replicas:

kafka-reassign-partitions.sh --zookeeper *{zk_host}:{port}*/**kafka --reassignment-json-file** *{manual assignment json file path}* **--execute**

For example:

/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper 192.168.0.90:2181,192.168.0.91:2181,192.168.0.92:2181/kafka --reassignment-json-file add-replicas-reassignment.json --execute

Step 4 Run the following command to check the task execution progress:

kafka-reassign-partitions.sh --zookeeper *{zk_host}:{port}*/**kafka --reassignment-json-file** *{manual assignment json file path}* **--verify**

For example:

/opt/client/Kafka/kafka/bin/kafka-reassign-partitions.sh --zookeeper 192.168.0.90:2181,192.168.0.91:2181,192.168.0.92:2181/kafka --reassignment-json-file add-replicas-reassignment.json --verify

Step 5 After completing the handling operations or confirming that the alarm has no impact, manually clear the alarm on FusionInsight Manager.

Step 6 After a period of time, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 In the **Service** area, select **Kafka** in the required cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

If the alarm has no impact, manually clear the alarm.

Related Information

None

10.255 ALM-38011 User Connection Usage on Broker Exceeds the Threshold

Description

The system checks the number of connections of each user on Broker every 30 seconds. This alarm is generated when the connection usage of a user on the Broker exceeds the threshold (80% by default) for 5 consecutive times.

The number of times that smoothing is performed is **5**. This alarm is cleared when the connection usage of a user on the Broker is less than the threshold.

The alarm can be automatically cleared. However, if the number of connections of a user suddenly becomes **0** and no connection is created, the alarm cannot be automatically cleared. You need to manually clear it.

Attribute

Alarm ID	Alarm Severity	Automatically Cleared
38011	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
UserName	Specifies the username for which the alarm is generated.

Impact on the System

If the number of connections of a user is excessive, the user cannot create new connections to the Broker.

Possible Causes

- The number of connections (created by a user) used by the client exceeds the preset threshold.
- The threshold for the connection usage does not meet service requirements.

Procedure

Check the number of connections established by the same user on the client.

Step 1 On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Alarms** > **User Connection Usage on Broker Exceeds the Threshold**. Check the host name and username of the Broker instance for which the alarm is generated in **Location**.

Step 2 On FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Instance**. Click the instance for which the alarm is generated to go to the page for the instance. Click the drop-down list in the upper right corner of the chart area, choose **Customize > Other**, and select **User Connection Usage on Broker**, **Maximum Number of User Connections on Broker**, and **Number of User Connections on Broker** to view the number of the current user connections on the Broker.

Step 3 Observe the number of real-time connections of the current alarm user and check whether the real-time monitoring data of the current user exists.

- If yes, go to **Step 4**.
- If no, the current user has disconnected all connections. You need to clear the alarm manually, and no further action is required.

  NOTE

  After the alarm user disconnects all connections, the monitoring data of the user disappears. In this case, the alarm will not be automatically cleared. You need to manually clear it.

Step 4 Check whether the user is authorized by the service side.

If yes, go to **Step 7**.

If no, go to **Step 5**.

Step 5 Run the following command on the client to limit the number of connections of the user. There are two configuration rules based on the following commands:

1. For the specific Broker and user, run the following command:

   kafka-configs.sh --bootstrap-server *<broker ip:port>* **--alter --add-config 'max.connections.per.user.overrides=[***<username>***:***<connection.number>***]' --entity-type brokers --entity-name** *<broker.id>* **--command-config Kafka/kafka/config/producer.properties**

   NOTE

   For unauthorized users, confirm with the service side to reduce the maximum number of connections of an unauthorized user or set the maximum number of connections to **0**.

   In the command, you need to specify the IP address and port number of Broker, set values of configuration items, and specify the **brokerId** and **username**. Here, the user refers to the authorized Kerberos user.

   The configuration updated using the command line tool can take effect dynamically. The configuration becomes invalid after the service is restarted. To make the configuration take effect after the restart, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations**> **Broker** > **Server** on the FusionInsight Manager home page and update the configuration to **max.connections.per.user.overrides**.

2. For the specific use and default Broker (that is, all Broker instances in the cluster), run the following command:

   kafka-configs.sh --bootstrap-server *<broker ip:port>* **--alter --add-config 'max.connections.per.user.overrides=[***<username>***:***<connection.number>***]' --entity-type brokers ---entity-default --command-config Kafka/kafka/config/client.properties**

   Example:

   kafka-configs.sh --bootstrap-server 10.153.3.26:21007 --alter --add-config 'max.connections.per.user.overrides=[showcase:4]' --entity-type brokers --entity-name 1 --command-config Kafka/kafka/config/client.properties

Step 6 Check whether the maximum number of connections is **0** and whether the number of connections of the current user decreases or remains unchanged according to **Step 2**.

- If yes, manually clear the alarm and no further action is required.
- If no, go to **Step 7**.

Step 7 Check whether the number of real-time connections and connection usage of the current user are sharply increased when they are compared with historical data, and whether have exceeded the specified maximum number of connections.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

NOTE

If there is an obvious increase after the comparison and the maximum number of connections has reached the preset value, the connections of the user may be abnormal. You need to confirm with the service party.

Check whether the number of user connections meets service requirements.

Step 8 Check whether the number of connections of the user meets service requirements.

- If yes, go to **Step 9**.
- If no, contact the service party to rectify the fault.

NOTE

If the number of user connections is abnormal, contact the service party to rectify the fault from the following aspects:

- Check whether new services are added so that the number of user connections increases sharply.
- Check whether handle leakage occurs on the code at the service side.

Step 9 Consider whether to increase the maximum number of connections of the user.

- If yes, go to **Step 10**.
- If no, go to **Step 12**.

Step 10 Increase the maximum number of connections based on the service requirements. Set the number of connections of the user on the Kafka client. For details, see **Step 5**.

Step 11 Wait for several minutes and then check whether the alarm is automatically cleared.

- If yes, go to **Step 12**.
- If no, go to **Step 2**.

Step 12 Determine whether to add the user to the whitelist based on service requirements on the service side.

- If yes, go to **Step 13**.
- If no, go to **Step 15**.

NOTE

To add a user to the whitelist, you need to restart the Kafka service. However, this operation will cause service interruption and affect service running. Therefore, you must confirm with the service side before performing this operation.

Step 13 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Kafka** > **Configurations** > **All Configurations** > **Broker(Role)** > **Server** to add the user to the **max.connections.per.user.whitelist** configuration item.

Step 14 Restart the service for the modification to take effect. In addition, you need to manually clear the alarm, and no further action is required.

Collect the fault information.

Step 15 On the FusionInsight Manager homepage, choose **O&M** > **Log** > **Download**.

Step 16 Expand the **Service** drop-down list, and select **Kafka** for the target cluster.

Step 17 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.256 ALM-43001 Spark2x Service Unavailable

Alarm Description

The system checks the Spark2x service status every 300 seconds. This alarm is generated when the Spark2x service is unavailable.

This alarm is cleared when the Spark2x service recovers.

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
43001	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The Spark tasks submitted by users fail to be executed.

Possible Causes

- The KrbServer service is abnormal.
- The LdapServer service is abnormal.
- ZooKeeper is abnormal.
- HDFS is abnormal.
- Yarn is abnormal.
- The corresponding Hive service is abnormal.
- The Spark2x assembly package is abnormal.
- The NameNode memory is insufficient.
- The memory of the Spark process is insufficient.

Handling Procedure

If the alarm is abnormal Spark2x assembly packet, the Spark packet is abnormal. Wait for about 10 minutes. The alarm is automatically cleared.

Check whether service unavailability alarms exist in services that Spark2x depends on.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**.

Step 2 Check whether the following alarms exist in the alarm list:

- ALM-25500 KrbServer Service Unavailable
- ALM-25000 LdapServer Service Unavailable
- ALM-13000 ZooKeeper Service Unavailable
- ALM-14000 HDFS Service Unavailable
- ALM-18000 Yarn Service Unavailable
- ALM-16004 Hive Service Unavailable
- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Handle the alarms based on the troubleshooting methods provided in the alarm help.

After the alarm is cleared, wait a few minutes and check whether the alarm GuardianService Unavailable is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check whether the NameNode memory is insufficient.

Step 4 Check whether the NameNode memory is insufficient.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Restart the NameNode to release the memory. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the memory of the Spark process is insufficient.

Step 6 Check whether the memory of the Spark process is insufficient due to memory-related modifications.

- If yes, go to **Step 7**.
- If no, go to **Step 8**.

Step 7 Ensure that the memory of the Spark process is sufficient or expand the cluster capacity. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 9 In the **Service** area, select the following nodes of the desired cluster. (Hive is the specific Hive service determined based on **ServiceName** in the alarm location information).

- KrbServer
- LdapServer
- ZooKeeper
- HDFS
- Yarn
- Hive

Step 10 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.257 ALM-43006 Heap Memory Usage of the JobHistory2x Process Exceeds the Threshold

Description

The system checks the JobHistory2x Process status every 30 seconds. The alarm is generated when the heap memory usage of a JobHistory2x Process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43006	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available JobHistory2x Process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the JobHistory2x Process is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43006**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **Memory > JobHistory2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used heap memory of the JobHistory2x Process reaches the threshold(default value is 95%) of the maximum heap memory specified for JobHistory2x.

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Figure 10-85 JobHistory2x Memory Usage Statistics

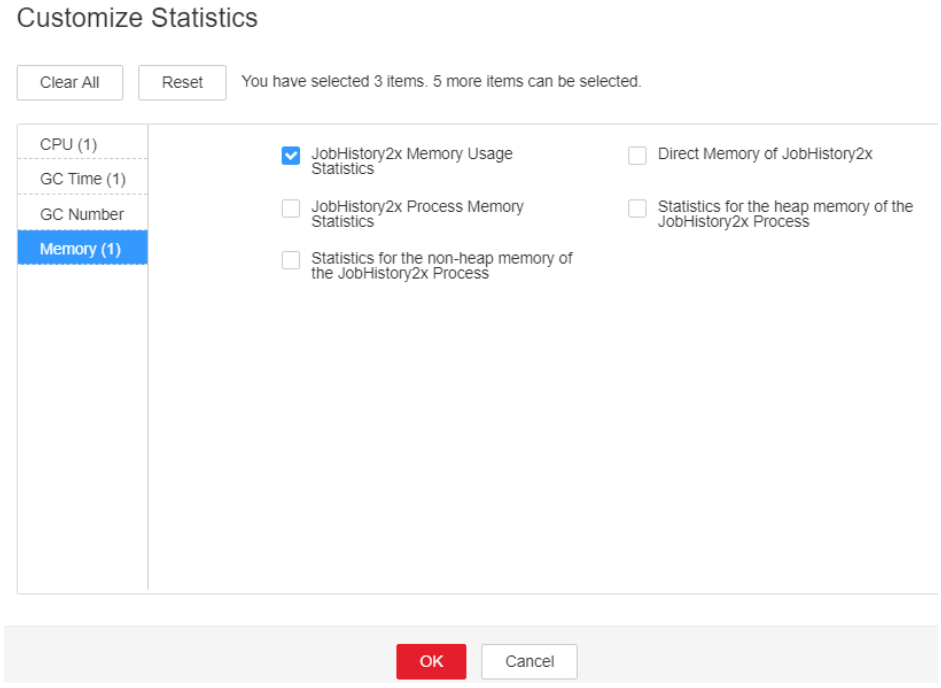

Step 3 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click **JobHistory2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Memory > Statistics for the heap memory of the JobHistory2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used heap memory of the JobHistory2x process in the corresponding period and obtain the maximum value.

Figure 10-86 Statistics for the heap memory of the JobHistory2x Process

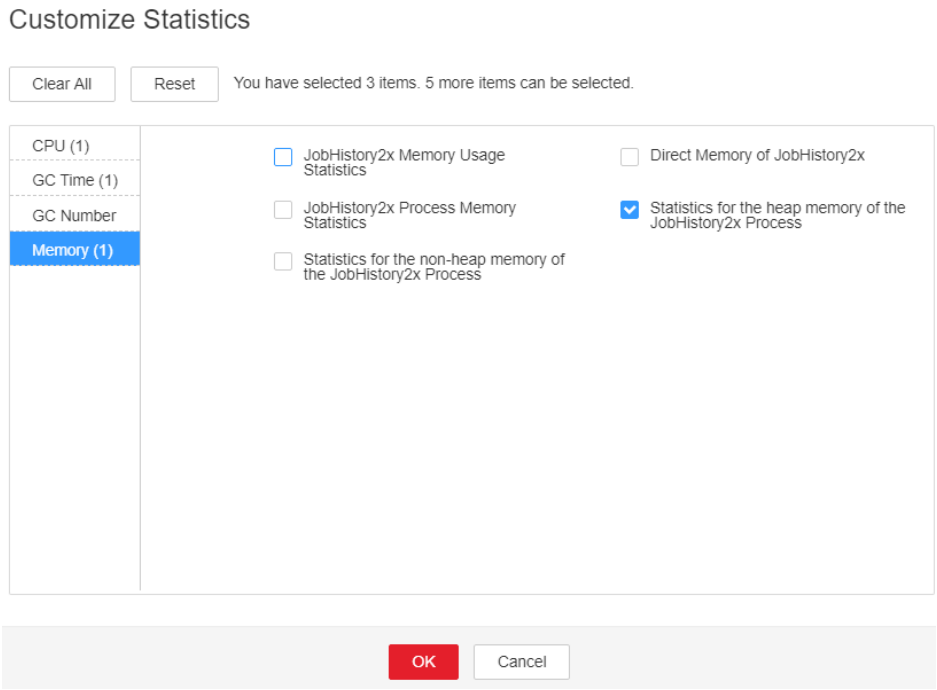

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JobHistory2x** > **Default**. The default value of **SPARK_DAEMON_MEMORY** is 4GB. You can change the value according to the following rules: Ratio of the maximum heap memory usage of the JobHistory2x to the **Threshold** of the **JobHistory2x Heap Memory Usage Statistics (JobHistory2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time.

NOTE

On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **Memory** >**JobHistory2x Heap Memory Usage Statistics (JobHistory2x)** to view **Threshold**.

Step 5 Restart all JobHistory2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **Spark2x** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.258 ALM-43007 Non-Heap Memory Usage of the JobHistory2x Process Exceeds the Threshold

Description

The system checks the JobHistory2x Process status every 30 seconds. The alarm is generated when the non-heap memory usage of a JobHistory2x Process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43007	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.

Name	Meaning
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available JobHistory2x Process non-heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The non-heap memory of the JobHistory2x Process is overused or the non-heap memory is inappropriately allocated.

Procedure

Check non-heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43007**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **Memory** > **JobHistory2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**, Check whether the used non-heap memory of the JobHistory2x Process reaches the threshold(default value is 95%) of the maximum non-heap memory specified for JobHistory2x.

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Figure 10-87 JobHistory2x Memory Usage Statistics

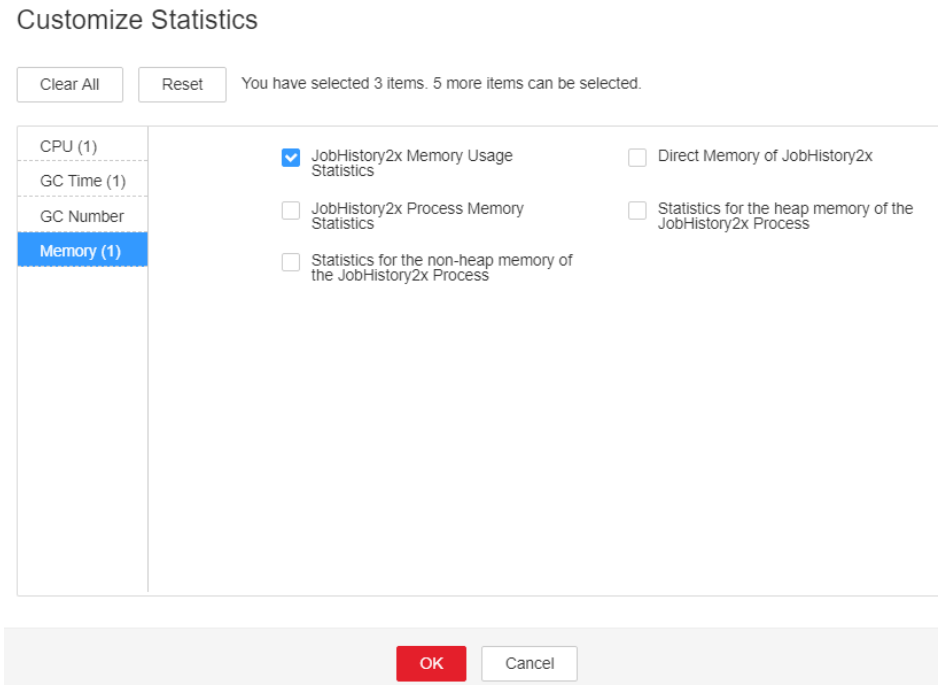

Step 3 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click **JobHistory2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Memory** > **Statistics for the non-heap memory of the JobHistory2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used non-heap memory of the JobHistory2x process in the corresponding period and obtain the maximum value.

Figure 10-88 Statistics for the non-heap memory of the JobHistory2x Process

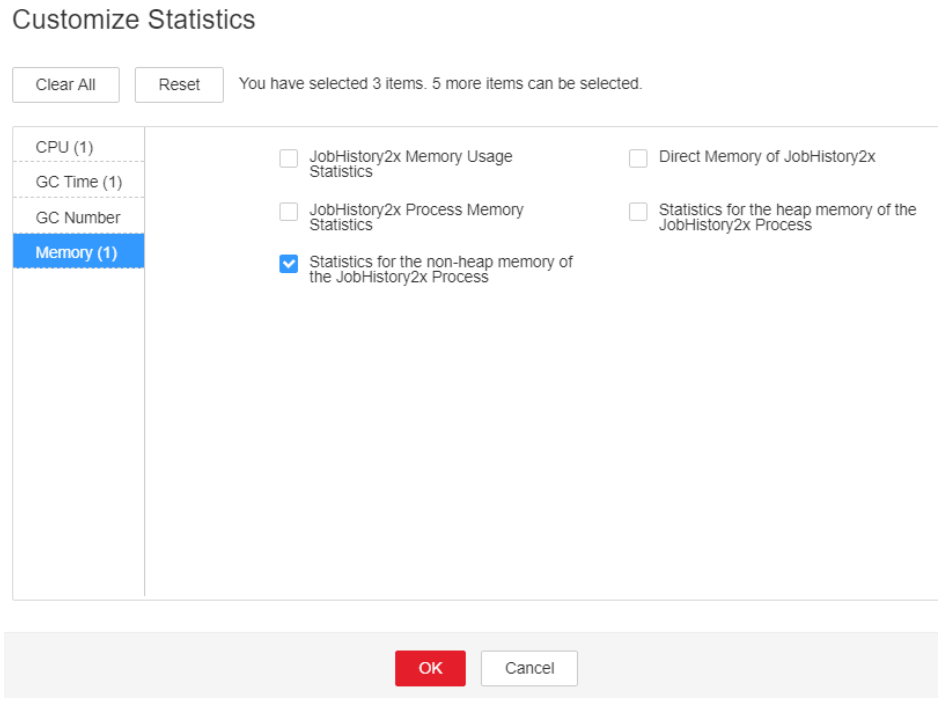

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JobHistory2x** > **Default**. You can change the value of **-XX:MaxMetaspaceSize** in **SPARK_DAEMON_JAVA_OPTS** according to the following rules: Ratio of the JobHistory2x non-heap memory usage to the **Threshold** of **JobHistory2x Non-Heap Memory Usage Statistics (JobHistory2x)** in the alarm period.

NOTE

On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* **> Spark2x** > **Memory** >**JobHistory2x Non-Heap Memory Usage Statistics (JobHistory2x)** to view **Threshold**.

Step 5 Restart all JobHistory2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **Spark2x** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.259 ALM-43008 The Direct Memory Usage of the JobHistory2x Process Exceeds the Threshold

Description

The system checks the JobHistory2x Process status every 30 seconds. The alarm is generated when the direct memory usage of a JobHistory2x Process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43008	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available JobHistory2x Process direct memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct memory of the JobHistory2x Process is overused or the direct memory is inappropriately allocated.

Procedure

Check direct memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43008**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **Memory** > **JobHistory2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used direct memory of the JobHistory2x Process reaches the threshold(default value is 95%) of the maximum direct memory specified for JobHistory2x.

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Figure 10-89 JobHistory2x Memory Usage Statistics

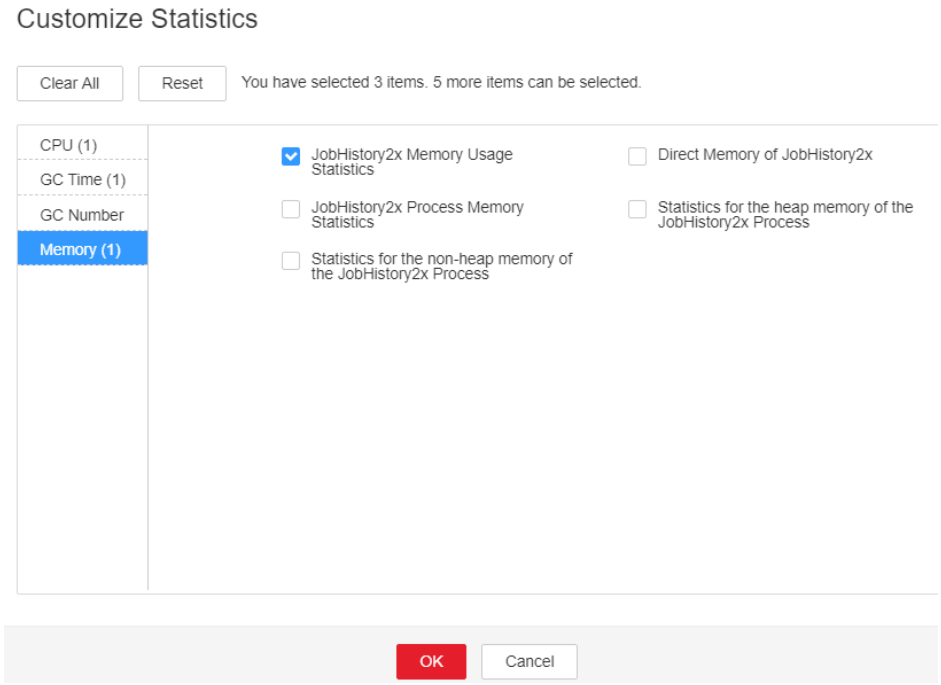

Step 3 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click **JobHistory2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Memory > Direct Memory of JobHistory2x**, and click **OK**. Based on the alarm generation time, check the values of the used direc memory of the JobHistory2x process in the corresponding period and obtain the maximum value.

Figure 10-90 Direct Memory of JobHistory2x

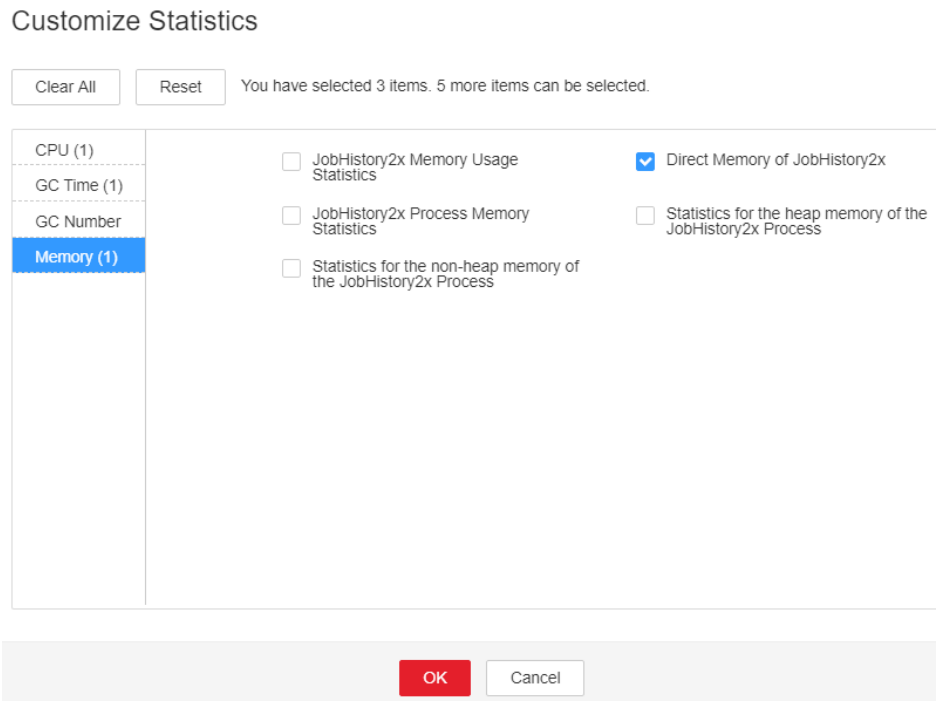

Step 4 On the FusionInsight Manager portal, choose **Cluster** > ***Name of the desired cluster*** **> Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JobHistory2x** > **Default**. The default value of **-XX:MaxDirectMemorySize** in **SPARK_DAEMON_JAVA_OPTS** is 512 MB. You can change the value according to the following rules: Ratio of the maximum direct memory usage of the JobHistory2x to the **Threshold** of the **JobHistory2x Direct Memory Usage Statistics (JobHistory2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time. It is recommended that the value be less than or equal to the value of **SPARK_DAEMON_MEMORY**.

NOTE

On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* **> Spark2x** > **Memory** >**JobHistory2x Direct Memory Usage Statistics (JobHistory2x)** to view **Threshold**.

Step 5 Restart all JobHistory2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **Spark2x** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.260 ALM-43009 JobHistory2x Process GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) time of the JobHistory2x Process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (exceeds 5 seconds for three consecutive checks.) To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **GC Time** > **Total GC time in milliseconds (JobHistory2x)**. This alarm is cleared when the JobHistory2x GC time is shorter than or equal to the threshold.

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43009	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.

Name	Meaning
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the GC time exceeds the threshold, JobHistory2x maybe run in low performance.

Possible Causes

The memory of JobHistory2x is overused, the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC time.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43009**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JobHistory2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **GC Time** > **Garbage Collection (GC) Time of JobHistory2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC time is longer than the threshold(default value: 12 seconds).

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Figure 10-91 Garbage Collection (GC) Time of JobHistory2x

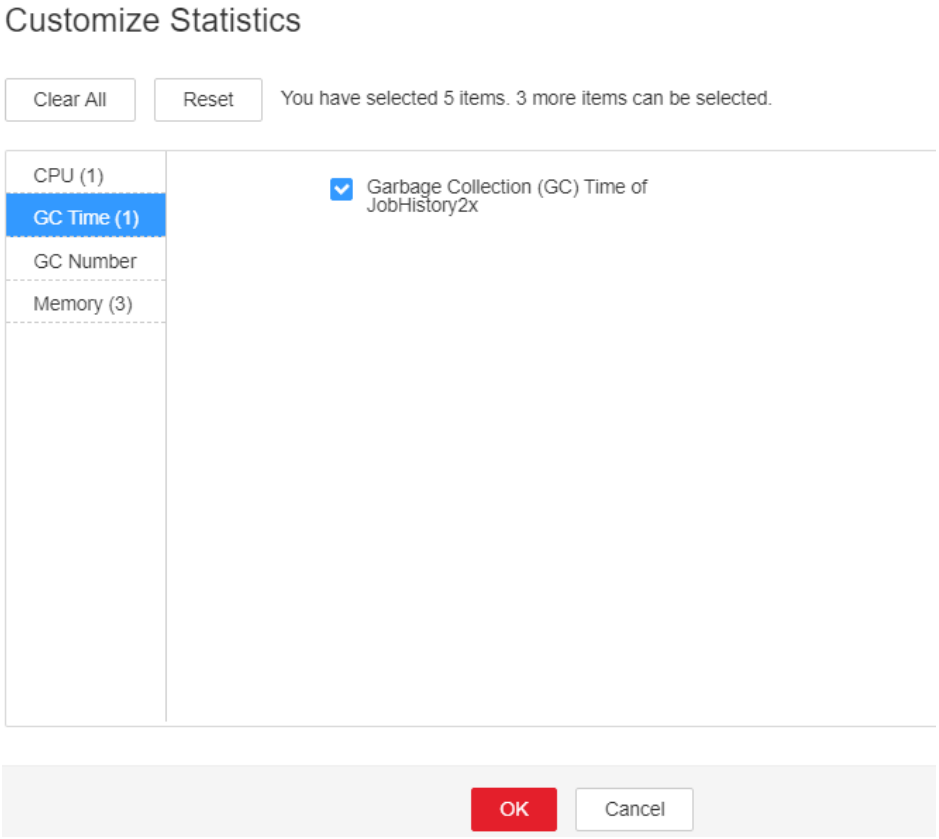

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JobHistory2x** > **Default**. The default value of **SPARK_DAEMON_MEMORY** is 4GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported, increase the value by 1 time.

Step 4 Restart all JobHistory2x instances.

Step 5 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager interface of active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 7 Select **Spark2x** in the required cluster from the **Service**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.261 ALM-43010 Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold

Description

The system checks the JDBCServer2x Process status every 30 seconds. The alarm is generated when the heap memory usage of a JDBCServer2x Process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43010	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available JDBCServer2x Process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the JDBCServer2x Process is overused or the heap memory is inappropriately allocated.

Procedure

Check heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43010**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **Memory** > **JDBCServer2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used heap memory of the JDBCServer2x Process reaches the threshold(default value is 95%) of the maximum heap memory specified for JDBCServer2x.

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Figure 10-92 JDBCServer2x Memory Usage Statistics

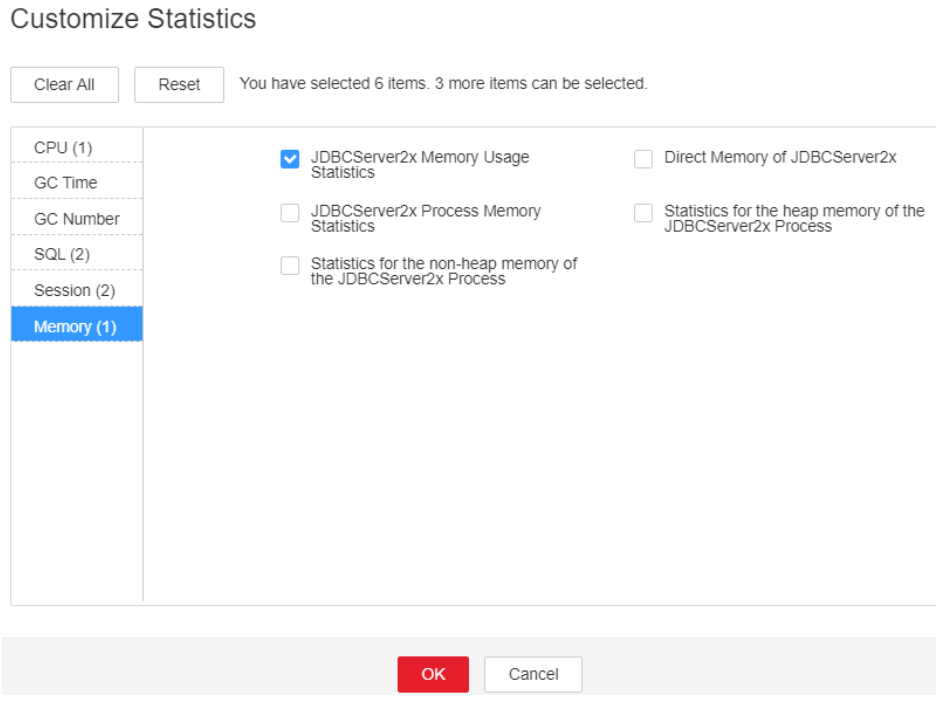

Step 3 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click **JDBCServer2x** by which the alarm

is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Memory > Statistics for the heap memory of the JDBCServer2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used heap memory of the JDBCServer2x process in the corresponding period and obtain the maximum value.

Figure 10-93 Statistics for the heap memory of the JDBCServer2x Process

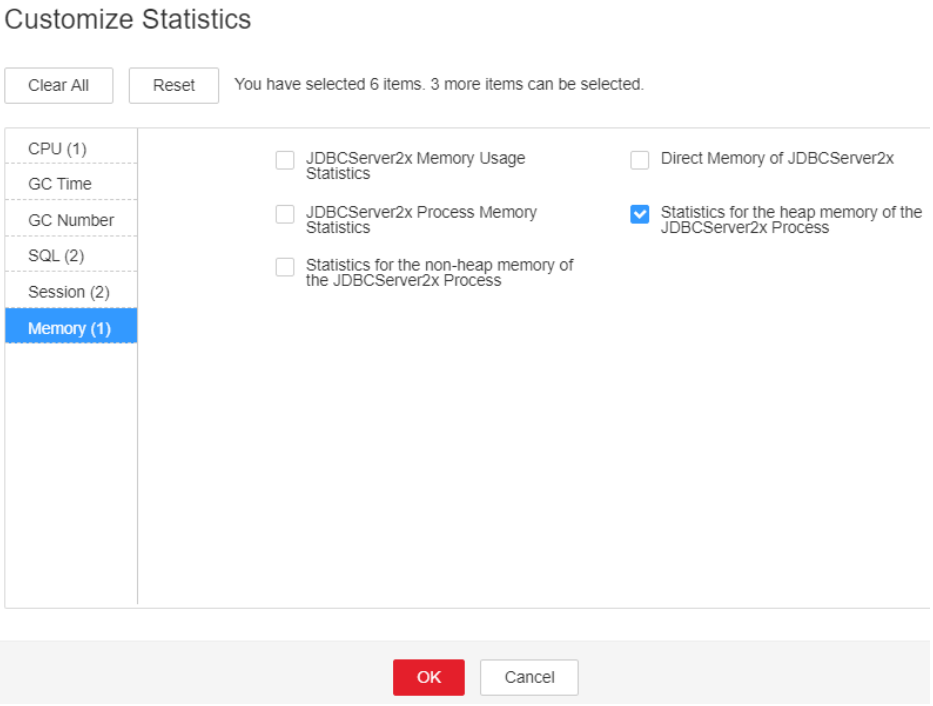

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JDBCServer2x** > **Tuning**. The default value of **SPARK_DRIVER_MEMORY** is 4 GB. You can change the value according to the following rules: Ratio of the maximum heap memory usage of the JobHistory2x to the **Threshold** of the **JDBCServer2x Heap Memory Usage Statistics (JDBCServer2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time. In the case of large service volume and high concurrency, add instances.

NOTE

On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* **> Spark2x** > **Memory** >**JDBCServer2x Heap Memory Usage Statistics (JDBCServer2x)** to view **Threshold**.

Step 5 Restart all JDBCServer2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log > Download**.

Step 8 Select **Spark2x** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.262 ALM-43011 Non-Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold

Description

The system checks the JDBCServer2x Process status every 30 seconds. The alarm is generated when the non-heap memory usage of an JDBCServer2x Process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43011	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.

Name	Meaning
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available JDBCServer2x Process non-heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The non-heap memory of the JDBCServer2x Process is overused or the non-heap memory is inappropriately allocated.

Procedure

Check non-heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43011**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **Memory** > **JDBCServer2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**, Check whether the used non-heap memory of the JDBCServer2x Process reaches the threshold(default value is 95%) of the maximum non-heap memory specified for JDBCServer2x.

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Figure 10-94 JDBCServer2x Memory Usage Statistics

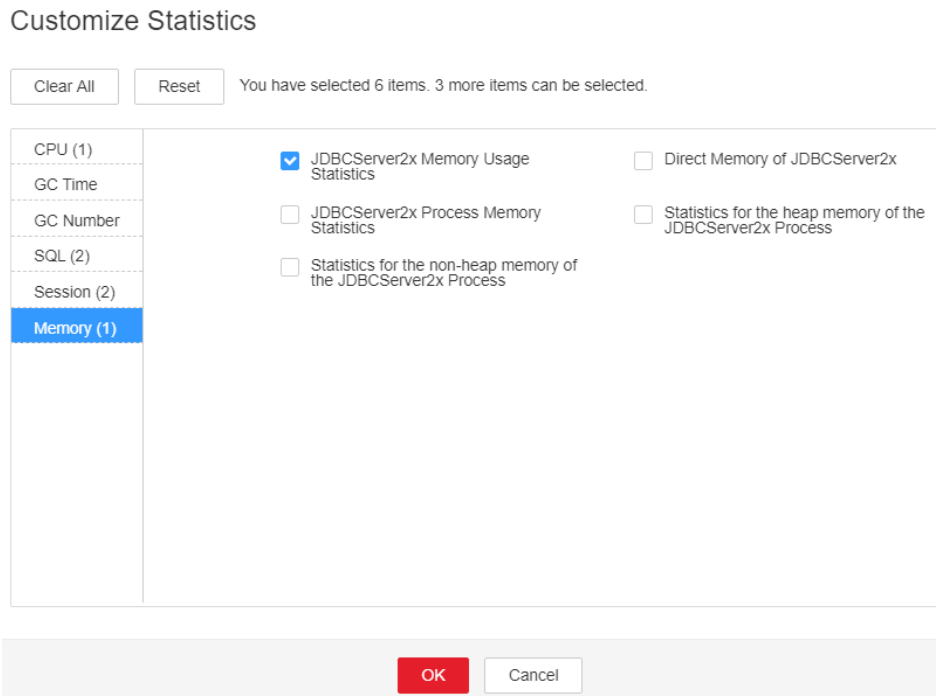

Step 3 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click **JDBCServer2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Memory** > **Statistics for the non-heap memory of the JDBCServer2x Process**, and click **OK**. Based on the alarm generation time, check the values of the used non-heap memory of the JDBCServer2x process in the corresponding period and obtain the maximum value.

Figure 10-95 Statistics for the non-heap memory of the JDBCServer2x Process

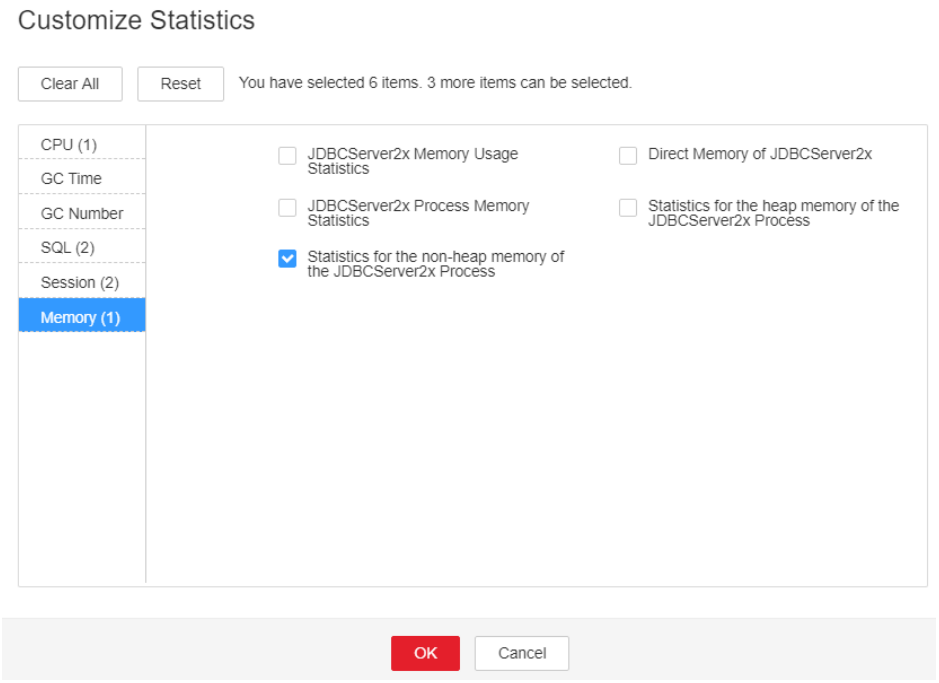

Step 4 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JDBCServer2x** > **Tuning**. You can change the value of **-XX: MaxMetaspaceSize** in **spark.driver.extraJavaOptions** according to the following rules: Ratio of the JDBCServer2x non-heap memory usage to the **Threshold** of **JDBCServer2x Non-Heap Memory Usage Statistics (JDBCServer2x)** in the alarm period.

NOTE

On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* **> Spark2x** > **Memory** > **JDBCServer2x Non-Heap Memory Usage Statistics (JDBCServer2x)** to view **Threshold**.

Step 5 Restart all JDBCServer2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log >Download**.

Step 8 Select **Spark2x** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.263 ALM-43012 Direct Heap Memory Usage of the JDBCServer2x Process Exceeds the Threshold

Description

The system checks the JDBCServer2x Process status every 30 seconds. The alarm is generated when the direct heap memory usage of a JDBCServer2x Process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43012	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Specifies the threshold triggering the alarm. If the current indicator value exceeds this threshold, the alarm is generated.

Impact on the System

If the available JDBCServer2x Process direct heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct heap memory of the JDBCServer2x Process is overused or the direct heap memory is inappropriately allocated.

Procedure

Check direct heap memory usage.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43012**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > Memory > **JDBCServer2x Memory Usage Statistics** from the drop-down list box in the upper right corner and click **OK**. Check whether the used direct heap memory of the JDBCServer2x Process reaches the threshold(default value is 95%) of the maximum direct heap memory specified for JDBCServer2x.

- If yes, go to **Step 3**.
- If no, go to **Step 7**.

Figure 10-96 JDBCServer2x Memory Usage Statistics

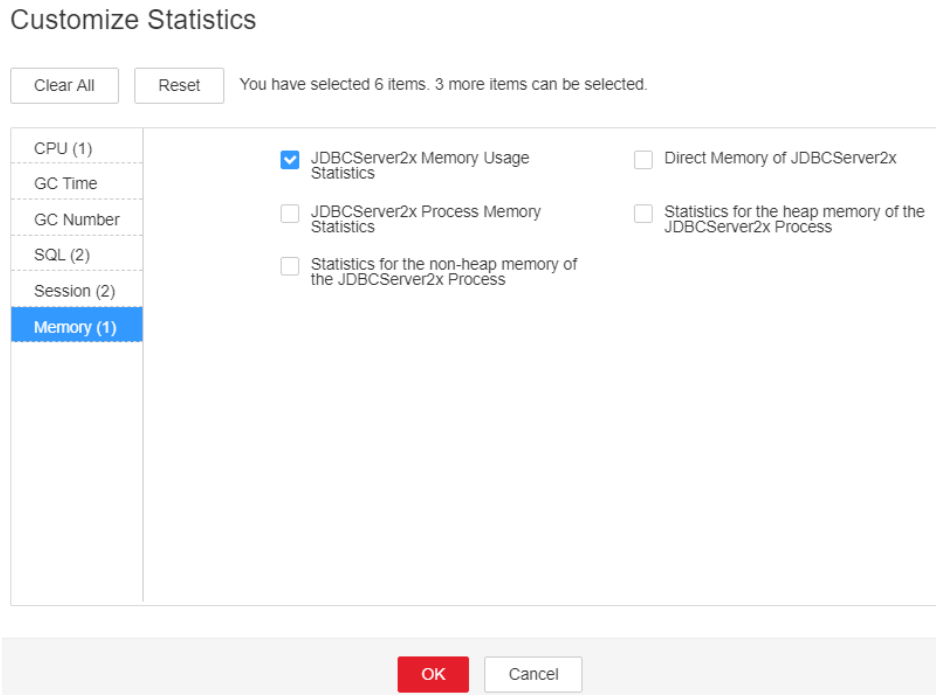

Step 3 On the FusionInsight Manager home page, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click **JDBCServer2x** by which the alarm is reported to go to the **Dashboard** page, click the drop-down list in the upper right corner of the chart area, choose **Customize** > **Memory > Direct Memory of JDBCServer2x**, and click **OK**. Based on the alarm generation time, check the values of the used direct memory of the JDBCServer2x process in the corresponding period and obtain the maximum value.

Figure 10-97 Direct Memory of JDBCServer2x

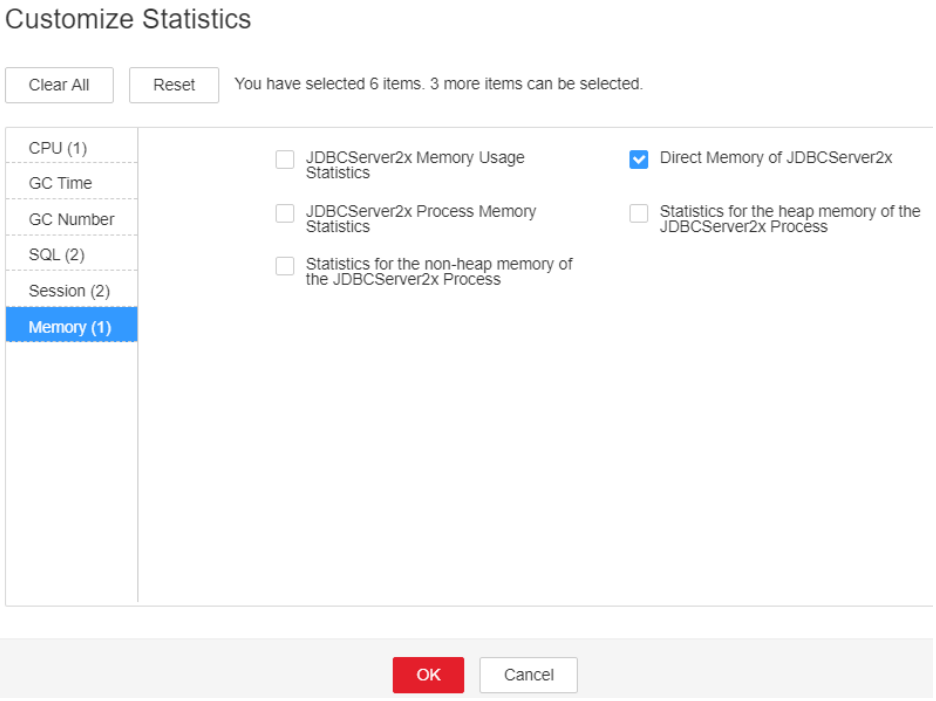

Step 4 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JDBCServer2x** > **Tuning**. The default value of **-XX:MaxDirectMemorySize** in **spark.driver.extraJavaOptions** is 512 MB. You can change the value according to the following rules: Ratio of the maximum direct memory usage of the JDBCServer2x to the **Threshold** of the **JDBCServer2x Direct Memory Usage Statistics (JDBCServer2x)** in the alarm period. If this alarm is generated occasionally after the parameter value is adjusted, increase the value by 0.5 times. If the alarm is frequently reported after the parameter value is adjusted, increase the value by 1 time. In the case of large service volume and high service concurrency, you are advised to add instances.

NOTE

On the FusionInsight Manager home page, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **Memory** >**JDBCServer2x Direct Memory Usage Statistics (JDBCServer2x)** to view **Threshold**.

Step 5 Restart all JDBCServer2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On the FusionInsight Manager portal, choose **O&M** > **Log** > **Download**.

Step 8 Select **Spark2x** in the required cluster from the **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.264 ALM-43013 JDBCServer2x Process GC Time Exceeds the Threshold

Description

The system checks the garbage collection (GC) time of the JDBCServer2x Process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (exceeds 5 seconds for three consecutive checks.) To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **GC Time** > **Total GC time in milliseconds (JDBCServer2x)**. This alarm is cleared when the JDBCServer2x GC time is shorter than or equal to the threshold.

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43013	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service name for which the alarm is generated.
RoleName	Specifies the role name for which the alarm is generated.

Name	Meaning
HostName	Specifies the object (host ID) for which the alarm is generated.
Trigger Condition	Generates an alarm when the actual indicator value exceeds the specified threshold.

Impact on the System

If the GC time exceeds the threshold, JDBCServer2x maybe run in low performance.

Possible Causes

The memory of JDBCServer2x is overused, the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Procedure

Check the GC time.

Step 1 On the FusionInsight Manager portal, choose **O&M > Alarm > Alarms** and select the alarm whose **ID** is **43013**. Check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On the FusionInsight Manager portal, choose **Cluster >** *Name of the desired cluster* **> Services** > **Spark2x** > **Instance** and click the JDBCServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **GC Time** > **Garbage Collection (GC) Time of JDBCServer2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC time is longer than the threshold(default value: 12 seconds).

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Figure 10-98 Garbage Collection (GC) Time of JDBCServer2x

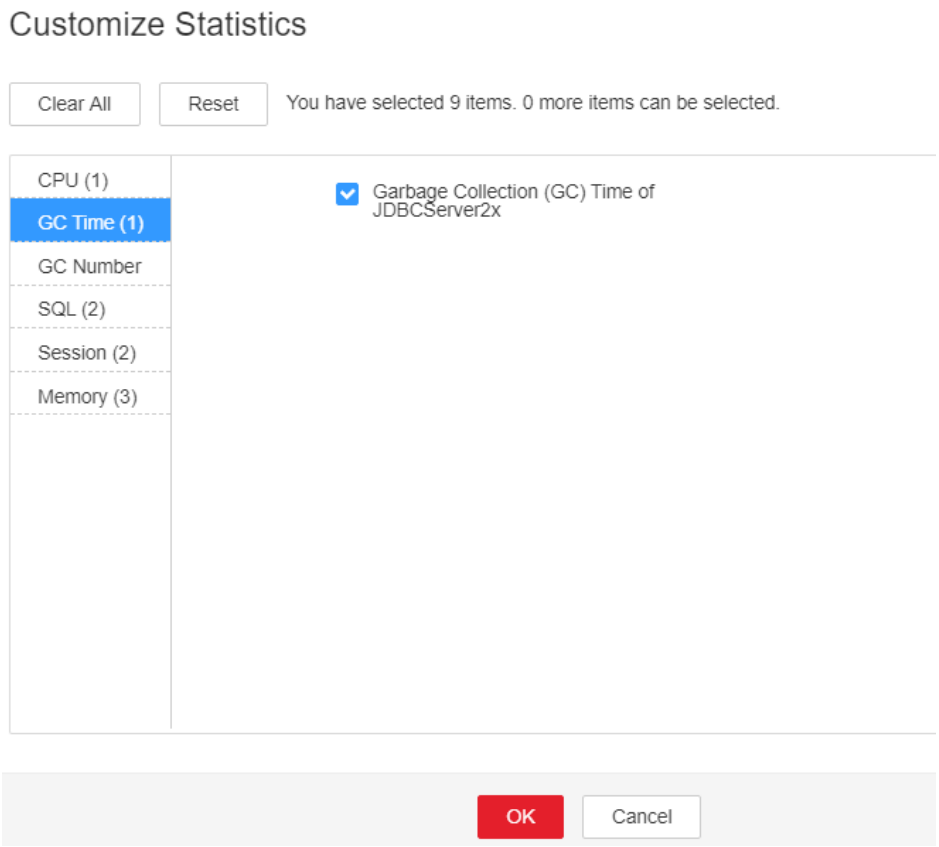

Step 3 On the FusionInsight Manager portal, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations**, and click **All Configurations**. Choose **JDBCServer2x** > **Default**. The default value of **SPARK_DRIVER_MEMORY** is 4 GB. If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported, increase the value by 1 time. In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Restart all JDBCServer2x instances.

Step 5 After 10 minutes, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On the FusionInsight Manager interface of active and standby clusters, choose **O&M** > **Log** > **Download**.

Step 7 Select **Spark2x** in the required cluster from the **Service**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Related Information

None

10.265 ALM-43017 JDBCServer2x Process Full GC Number Exceeds the Threshold

Description

The system checks the number of Full garbage collection (GC) times of the JDBCServer2x process every 60 seconds. This alarm is generated when the detected Full GC number exceeds the threshold (exceeds 12 for three consecutive checks.) You can change the threshold by choosing **O&M > Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **GC number** > **Full GC Number of JDBCServer2x**. This alarm is cleared when the Full GC number of the JDBCServer2x process is less than or equal to the threshold.

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JDBCServer2x is changed to JDBCServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43017	Major	Yes

Parameters

Name	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Name	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The performance of the JDBCServer2x process is affected, or even the JDBCServer2x process is unavailable.

Possible Causes

The heap memory usage of the JDBCServer2x process is excessively large, or the heap memory is inappropriately allocated. As a result, Full GC occurs frequently.

Procedure

Check the number of Full GCs.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm > Alarms**, select this alarm, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. On the displayed page, click the JDBCServer2x for which the alarm is reported. On the **Dashboard** page that is displayed, click the drop-down menu in the Chart area and choose **Customize** > **GC Number** > **Full GC Number of JDBCServer2x** in the upper right corner and click **OK**. Check whether the number of Full GCs of the JDBCServer2x process is greater than the threshold(default value: 12).

- If it is, go to **Step 3**.
- If it is not, go to **Step 6**.

Figure 10-99 Full GC Number of JDBCServer2x

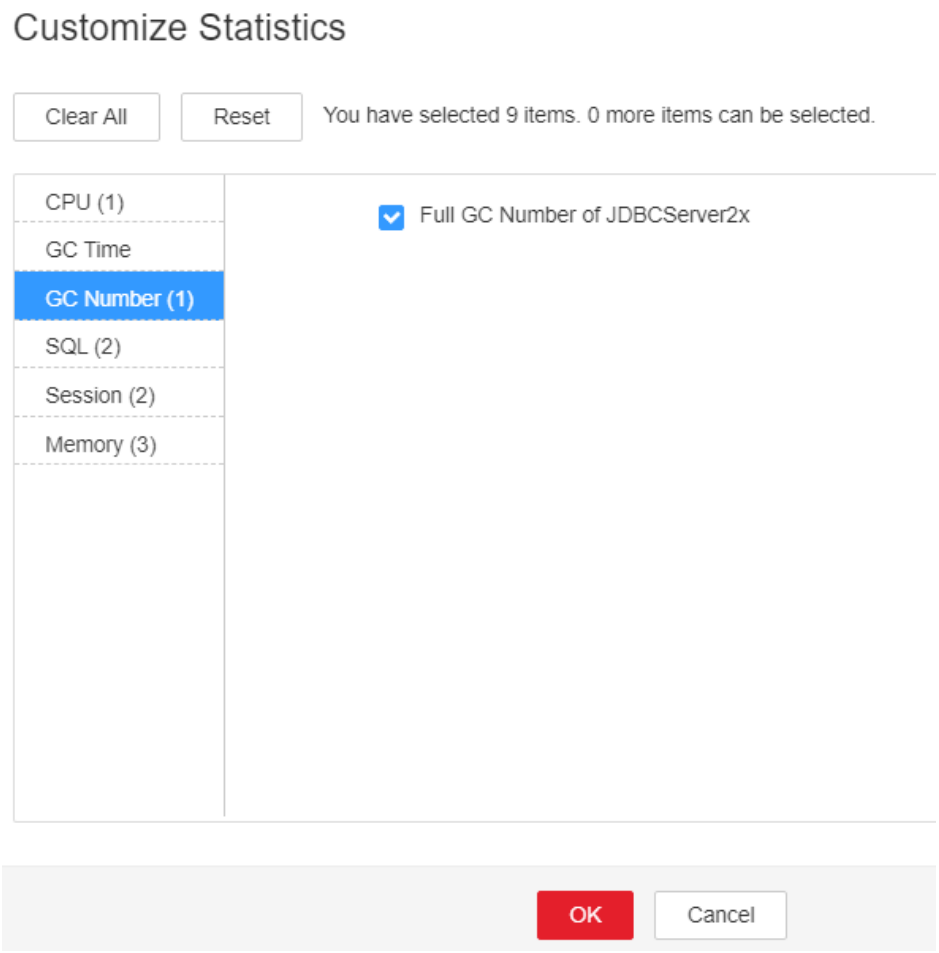

Step 3 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations** > **All Configurations**. On the displayed page, choose **JDBCServer2x** > **Tuning**. The default value of **SPARK_DRIVER_MEMORY** is 4GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported, increase the value by 1 time. In the case of large service volume and high concurrency, add instances.

Step 4 Restart all JDBCServer2x instances.

Step 5 After 10 minutes, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 Log in to FusionInsight Manager, and choose **O&M** > **Log** > **Download**.

Step 7 Select **Spark2x** in the required cluster from the **Service** drop-down list.

Step 8 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.266 ALM-43018 JobHistory2x Process Full GC Number Exceeds the Threshold

Description

The system checks the number of Full garbage collection (GC) times of the JobHistory2x process every 60 seconds. This alarm is generated when the detected Full GC number exceeds the threshold (exceeds 12 for three consecutive checks.) You can change the threshold by choosing **O&M** > **Alarm > Thresholds** > *Name of the desired cluster* > **Spark2x** > **GC number** > **Full GC Number of JobHistory2x**. This alarm is cleared when the Full GC number of the JobHistory2x process is less than or equal to the threshold.

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, JobHistory2x is changed to JobHistory. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Alarm Severity	Auto Clear
43018	Major	Yes

Parameters

Name	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Name	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The performance of the JobHistory2x process is affected, or even the JobHistory2x process is unavailable.

Possible Causes

The heap memory usage of the JobHistory2x process is excessively large, or the heap memory is inappropriately allocated. As a result, Full GC occurs frequently.

Procedure

Check the number of Full GCs.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm > Alarms**, select this alarm, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. On the displayed page, click the JobHistory2x for which the alarm is reported. On the **Dashboard** page that is displayed, click the drop-down menu in the Chart area and choose **Customize** > **GC Number** > **Full GC Number of JobHistory2x** in the upper right corner and click **OK**. Check whether the number of Full GCs of the JobHistory2x process is greater than the threshold(default value: 12).

- If it is, go to **Step 3**.
- If it is not, go to **Step 6**.

Figure 10-100 Full GC Number of JobHistory2x

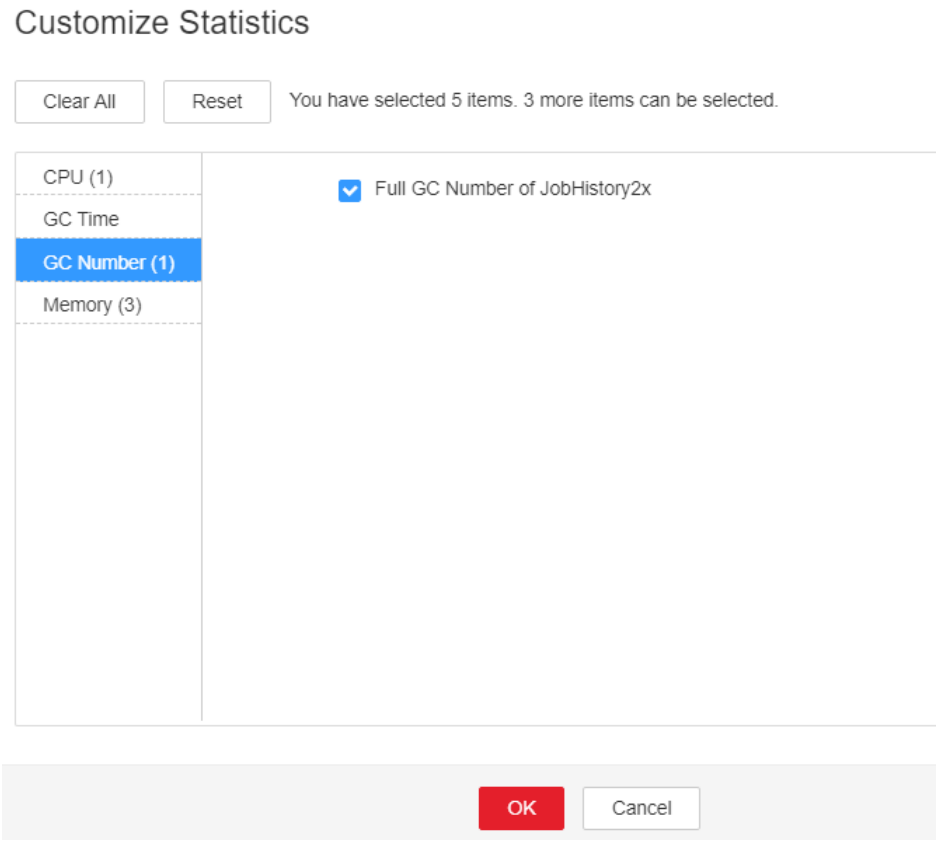

Step 3 Choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations** > **All Configurations**. On the displayed page, choose **JobHistory2x** > **Default.** The default value of **SPARK_DAEMON_MEMORY** is 4GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. If the alarm is frequently reported, increase the value by 1 time.

Step 4 Restart all JobHistory2x instances.

Step 5 After 10 minutes, check whether the alarm is cleared.

- If it is, no further action is required.
- If it is not, go to **Step 6**.

Collect fault information.

Step 6 Log in to FusionInsight Manager, and choose **O&M** > **Log** > **Download**.

Step 7 Select **Spark2x** in the required cluster from the **Service**.

Step 8 Click in the upper right corner. In the displayed dialog box, set **Start Date** and **End Date** to 10 minutes before and after the alarm generation time respectively and click **OK**. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected logs.

----End

Alarm Clearing

This alarm will be automatically cleared after the fault is rectified.

Related Information

None

10.267 ALM-43019 Heap Memory Usage of the IndexServer2x Process Exceeds the Threshold

Description

The system checks the IndexServer2x process status every 30 seconds. The alarm is generated when the heap memory usage of a IndexServer2x process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Severity	Auto Clear
43019	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the available IndexServer2x process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the IndexServer2x process is overused or the heap memory is inappropriately allocated.

Procedure

Check the heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm > Alarms**. In the displayed alarm list, choose the alarm for which the ID is **43019**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize** > **Memory** > **IndexServer2x Memory Usage Statistics** > **OK**. Check whether the heap memory used by the IndexServer2x process reaches the maximum heap memory threshold (95% by default).

- If the threshold is reached, go to **Step 3**.
- If the threshold is not reached, go to **Step 7**.

Figure 10-101 IndexServer2x Memory Usage Statistics

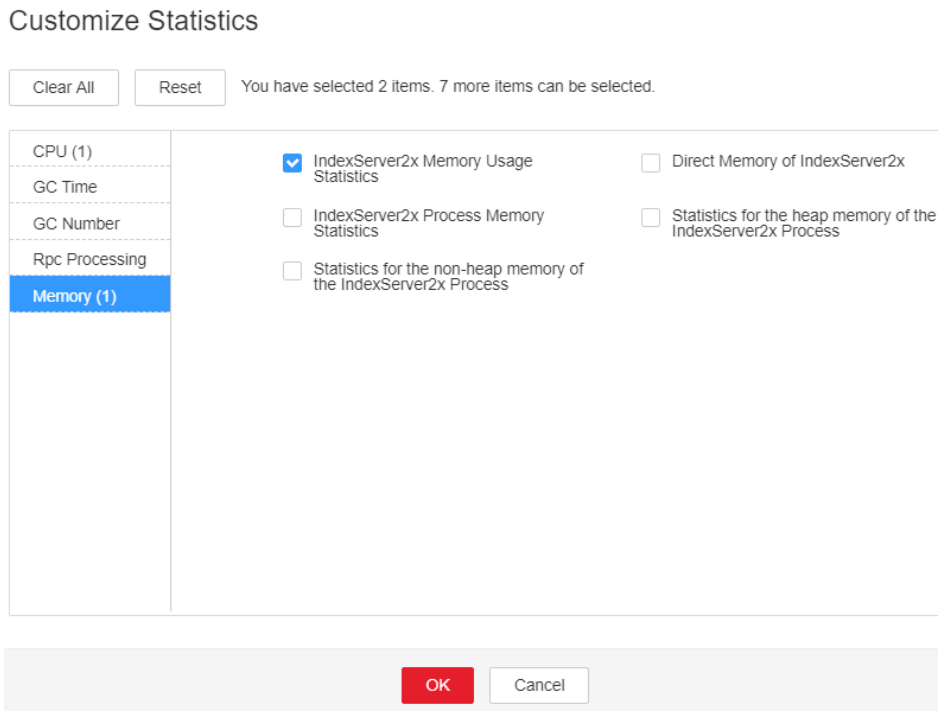

Step 3 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click the IndexServer2x that reported the alarm to

go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize** > **Memory** > **Statistics for the heap memory of the IndexServer2x Process** > **OK**. Based on the alarm generation time, check the values of the used heap memory of the IndexServer2x process in the corresponding period and obtain the maximum value.

Figure 10-102 Statistics for the heap memory of the IndexServer2x Process

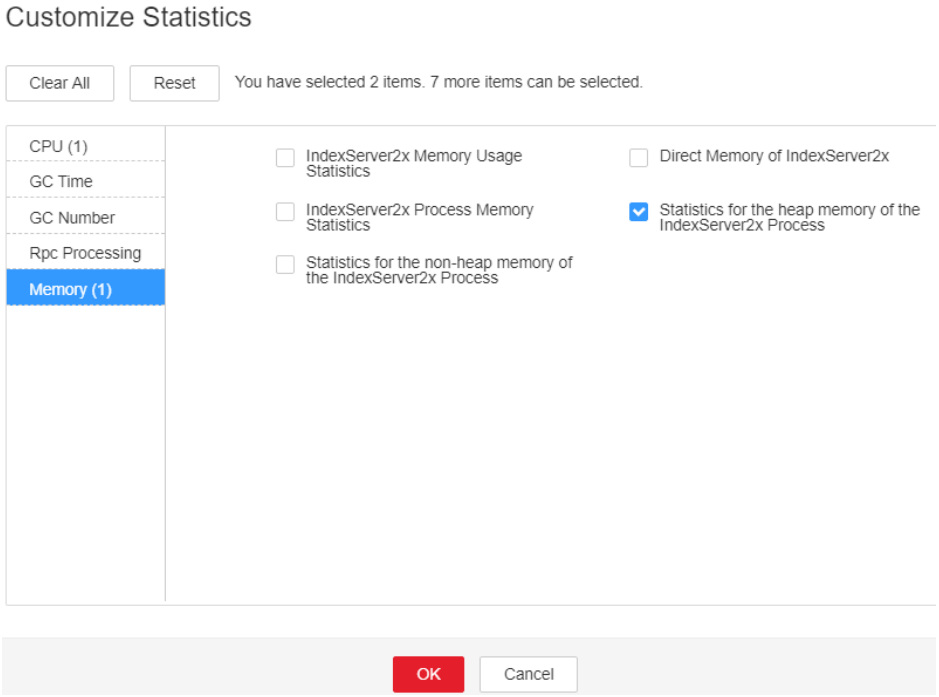

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations** > **All Configuration** > **IndexServer2x**> **Tuning**. The default value of the **SPARK_DRIVER_MEMORY** parameter is 4 GB. You can change the value based on the ratio of the maximum heap memory used by the IndexServer2x process to the threshold specified by **IndexServer2x Heap Memory Usage Statistics (IndexServer2x)** in the alarm period. If the alarm persists after the parameter value is changed, increase the value by 0.5 times. If the alarm is generated frequently, double the rate.

NOTE

On FusionInsight Manager, you can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **Memory** > **IndexServer2x Heap Memory Usage Statistics (IndexServer2x)** to view the threshold.

Step 5 Restart all IndexServer2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and provide the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Reference

None

10.268 ALM-43020 Non-Heap Memory Usage of the IndexServer2x Process Exceeds the Threshold

Description

The system checks the IndexServer2x process status every 30 seconds. The alarm is generated when the non-heap memory usage of the IndexServer2x process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Severity	Auto Clear
43020	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the available IndexServer2x process non-heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The non-heap memory of the IndexServer2x process is overused or the non-heap memory is inappropriately allocated.

Procedure

Check non-heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm > Alarms**. In the displayed alarm list, choose the alarm for which the ID is **43020**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize** > **Memory > IndexServer2x Memory Usage Statistics** > **OK**. Check whether the non-heap memory used by the IndexServer2x process reaches the maximum non-heap memory threshold (95% by default).

- If the threshold is reached, go to **Step 3**.
- If the threshold is not reached, go to **Step 7**.

Figure 10-103 IndexServer2x Memory Usage Statistics

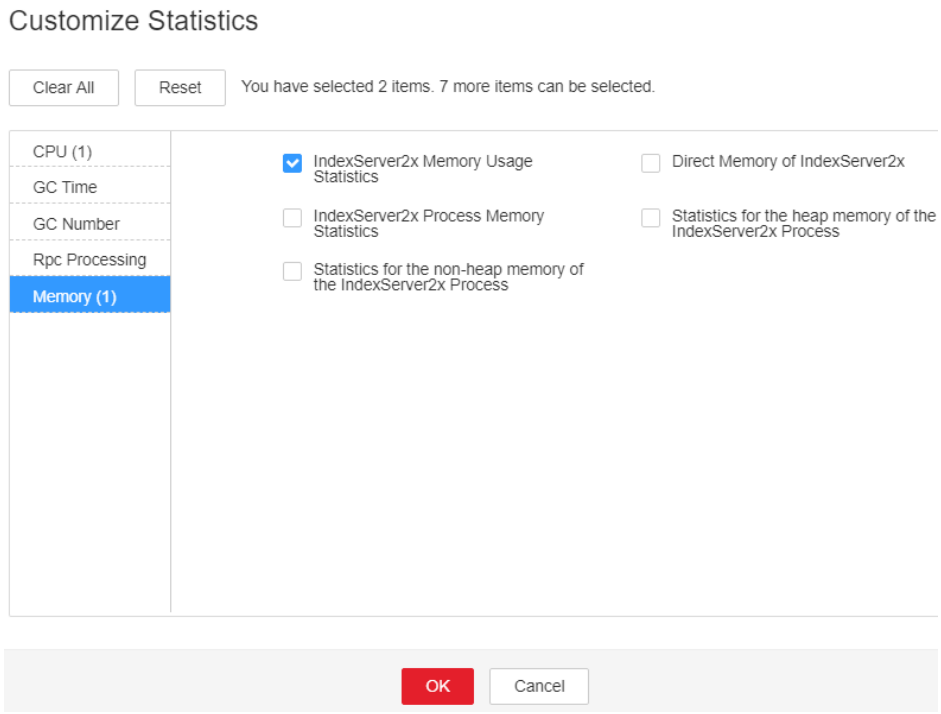

Step 3 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize** > **Memory** > **Statistics for the non-heap memory of the IndexServer2x Process** > **OK**. Based on the alarm generation time, check the values of the used non-heap memory of the IndexServer2x process in the corresponding period and obtain the maximum value.

Figure 10-104 Statistics for the non-heap memory of the IndexServer2x Process

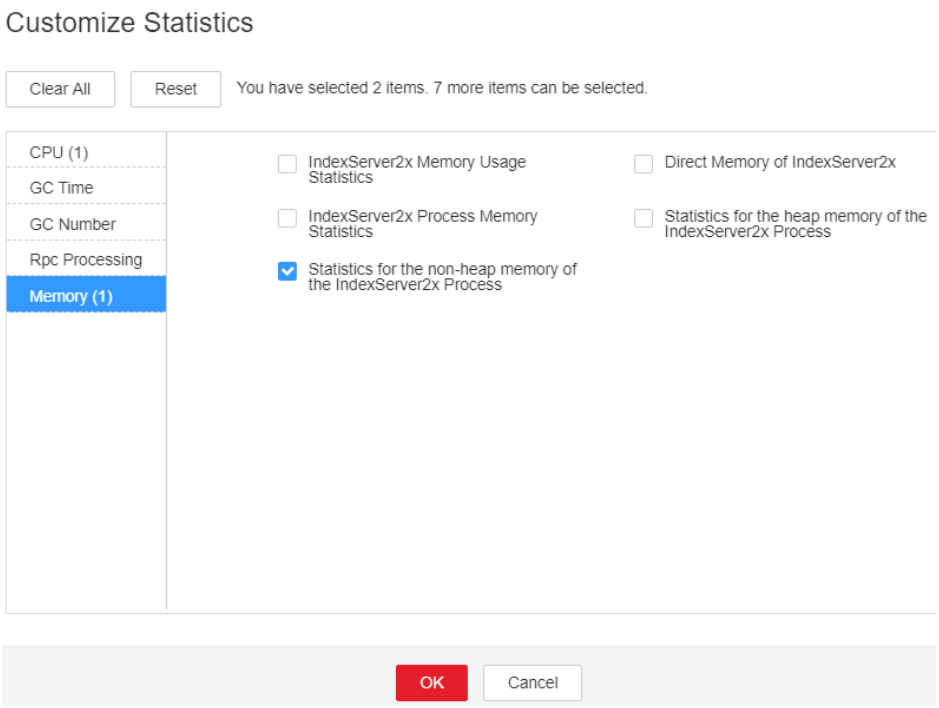

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations** > **All Configurations** > **IndexServer2x**> **Tuning**. You can change the value of **XX:MaxMetaspaceSize** in the **spark.driver.extraJavaOptions** parameter based on the ratio of the maximum non-heap memory used by the IndexServer2x process to the threshold specified by **IndexServer2x Non-Heap Memory Usage Statistics (IndexServer2x)** in the alarm period.

NOTE

On FusionInsight Manager, you can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **Memory** > **IndexServer2x Non-Heap Memory Usage Statistics (IndexServer2x)** to view the threshold.

Step 5 Restart all IndexServer2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and provide the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Reference

None

10.269 ALM-43021 Direct Memory Usage of the IndexServer2x Process Exceeds the Threshold

Description

The system checks the IndexServer2x process status every 30 seconds. The alarm is generated when the direct heap memory usage of a IndexServer2x process exceeds the threshold (95% of the maximum memory).

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Severity	Auto Clear
43021	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the available IndexServer2x process direct memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The direct heap memory of the IndexServer2x process is overused or the direct heap memory is inappropriately allocated.

Procedure

Check direct heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the displayed alarm list, choose the alarm for which the ID is **43021**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of the chart area, and choose **Customize** > **Memory** > **IndexServer2x Memory Usage Statistics** > **OK**. Check whether the direct memory used by the IndexServer2x process reaches the maximum direct memory threshold.

- If the threshold is reached, go to **Step 3**.
- If the threshold is not reached, go to **Step 7**.

Figure 10-105 IndexServer2x Memory Usage Statistics

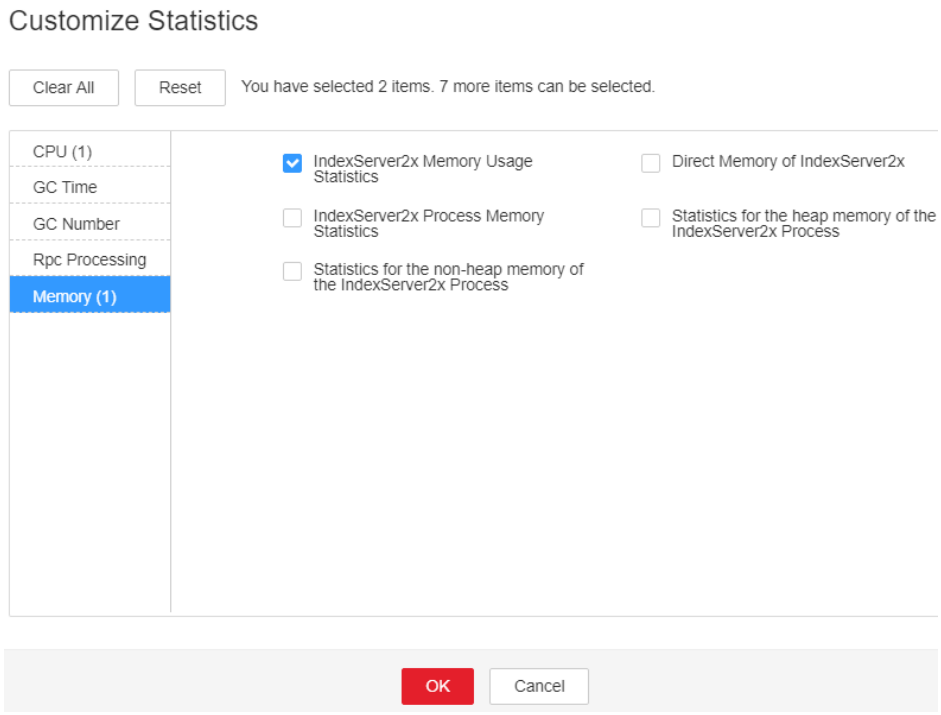

Step 3 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance**. Click the IndexServer2x that reported the alarm to go to the **Dashboard** page. Click the drop-down list in the upper right corner of

the chart area, and choose **Customize** > **Memory** > **Direct Memory of IndexServer2x** > **OK**. Based on the alarm generation time, check the values of the used direct memory of the IndexServer2x process in the corresponding period and obtain the maximum value.

Figure 10-106 Direct Memory of IndexServer2x

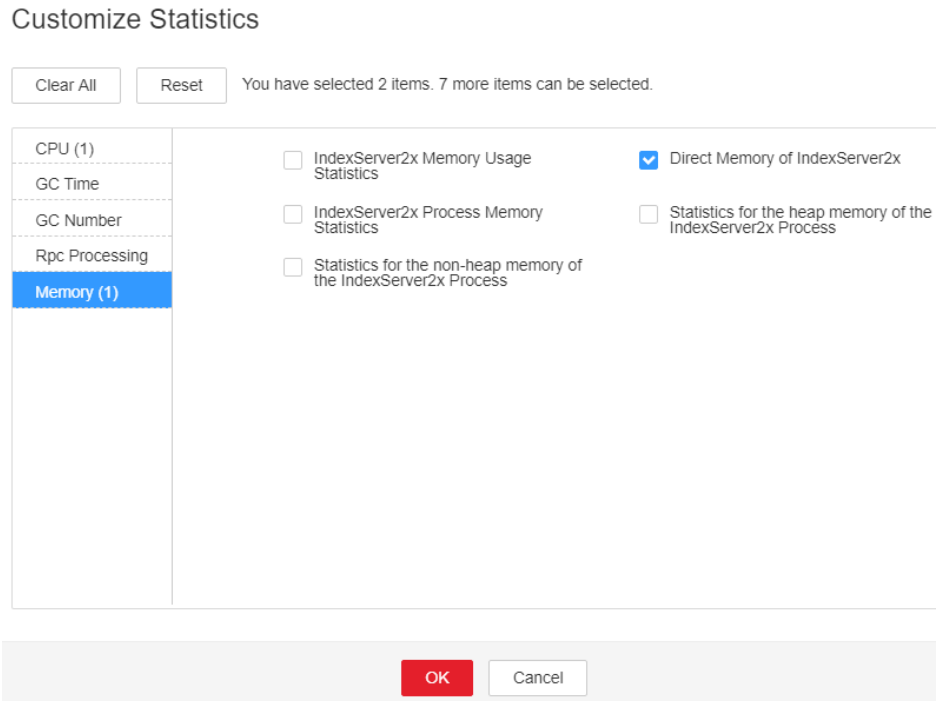

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations** > **All Configurations** > **IndexServer2x**> **Tuning**. You can change the value of **XX:MaxDirectMemorySize** (the default value is 512 MB) in the **spark.driver.extraJavaOptions** parameter based on the ratio of the maximum direct memory used by the IndexServer2x process to the threshold specified by **IndexServer2x Direct Memory Usage Statistics (IndexServer2x)** in the alarm period. If the alarm persists after the parameter value is changed, increase the value by 0.5 times. If the alarm is generated frequently, double the rate.

NOTE

On FusionInsight Manager, you can choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **Memory** > **IndexServer2x Direct Memory Usage Statistics (IndexServer2x)** to view the threshold.

Step 5 Restart all IndexServer2x instances.

Step 6 After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact the O&M personnel and provide the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Reference

None

10.270 ALM-43022 IndexServer2x Process GC Time Exceeds the Threshold

Description

The system checks the GC time of the IndexServer2x process every 60 seconds. This alarm is generated when the detected GC time exceeds the threshold (12 seconds) for three consecutive times. To change the threshold, choose **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **GC Time** > **Total GC time in milliseconds (IndexServer2x)**. This alarm is cleared when the IndexServer2x GC time is shorter than or equal to the threshold.

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Severity	Auto Clear
43022	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC time exceeds the threshold, IndexServer2x may run in low performance or even unavailable.

Possible Causes

The heap memory of the IndexServer2x process is overused or the heap memory is inappropriately allocated. As a result, GC occurs frequently.

Procedure

Check the GC time.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm > Alarms**. In the displayed alarm list, choose the alarm with ID **43022**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance** and click the IndexServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the Chart area and choose **Customize** > **GC Time** > **Garbage Collection (GC) Time of IndexServer2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC time is longer than the threshold (default value: 12 seconds).

- If the threshold is reached, go to **Step 3**.
- If the threshold is not reached, go to **Step 6**.

Figure 10-107 Garbage Collection (GC) Time of IndexServer2x

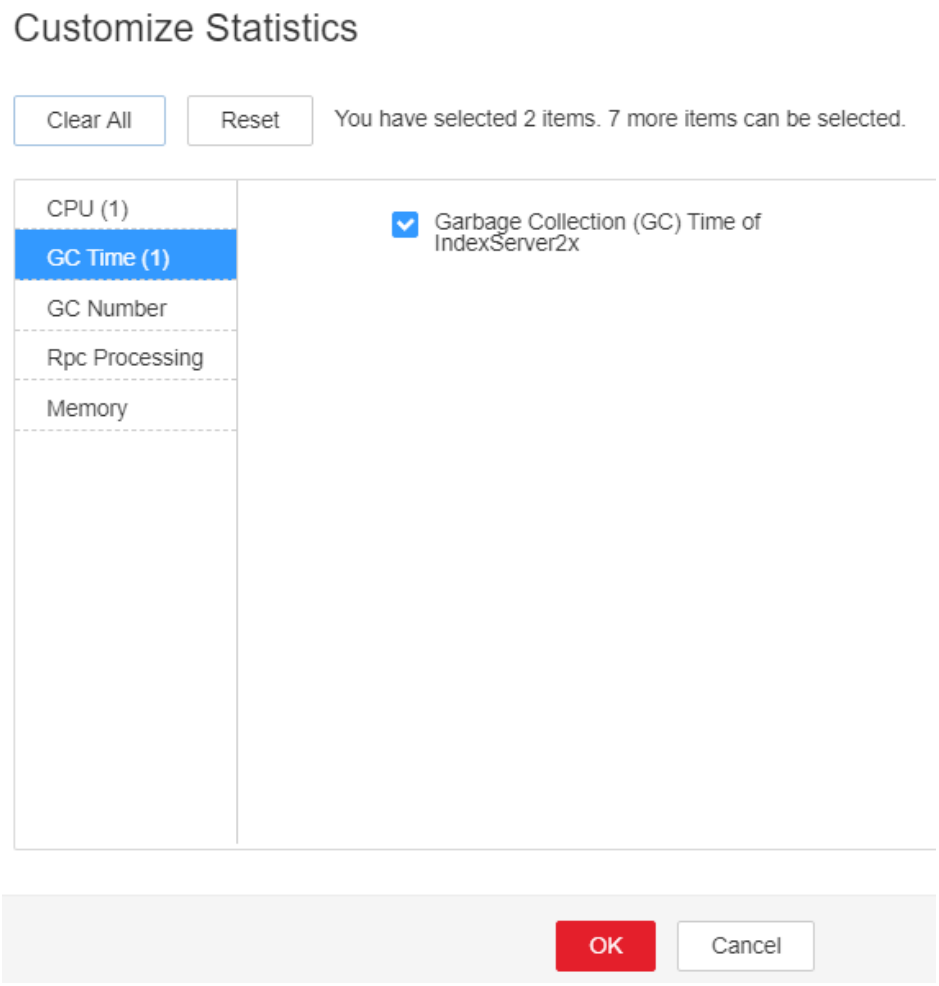

Step 3 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations** > **All Configurations** > **IndexServer2x** > **Default**. The default value of the **SPARK_DRIVER_MEMORY** is 4 GB. You can change the value according to the following rules: Increase the value of the **SPARK_DRIVER_MEMORY** parameter 1.5 times to its default value. If this alarm is still generated occasionally after the adjustment, increase the value by 0.5 times. Double the value if the alarm is reported frequently.

Step 4 Restart all IndexServer2x instances.

Step 5 After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and provide the collected logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Reference

None

10.271 ALM-43023 IndexServer2x Process Full GC Number Exceeds the Threshold

Description

The system checks the Full GC number of the IndexServer2x process every 60 seconds. This alarm is generated when the detected Full GC number exceeds the threshold (12) for three consecutive times. You can change the threshold by choosing **O&M** > **Alarm** > **Thresholds** > *Name of the desired cluster* > **Spark2x** > **GC Number** > **Full GC Number of IndexServer2x**. This alarm is cleared when the Full GC number of the IndexServer2x process is less than or equal to the threshold. This alarm is cleared when the Full GC number of the IndexServer2x process is less than or equal to the threshold.

NOTE

In MRS 3.3.0-LTS and later versions, the Spark2x component is renamed Spark, and the role names in the component are also changed. For example, IndexServer2x is changed to IndexServer. Refer to the descriptions and operations related to the component name and role names in the document based on your MRS version.

Attribute

Alarm ID	Severity	Auto Clear
43023	Major	Yes

Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC number exceeds the threshold, IndexServer2x maybe run in low performance or even unavailable.

Possible Causes

The heap memory of the IndexServer2x process is overused or the heap memory is inappropriately allocated. As a result, Full GC occurs frequently.

Procedure

Check the number of Full GCs.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm > Alarms**. In the displayed alarm list, choose the alarm with the ID **43023**, and check the **RoleName** in **Location** and confirm the IP address of **HostName**.

Step 2 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Instance** and click the IndexServer2x for which the alarm is generated to go to the **Dashboard** page. Click the drop-down menu in the chart area and choose **Customize** > **GC Number > Full GC Number of IndexServer2x** from the drop-down list box in the upper right corner and click **OK** to check whether the GC number is larger than the threshold (default value: 12).

- If the threshold is reached, go to **Step 3**.
- If the threshold is not reached, go to **Step 6**.

Figure 10-108 Full GC Number of IndexServer2x

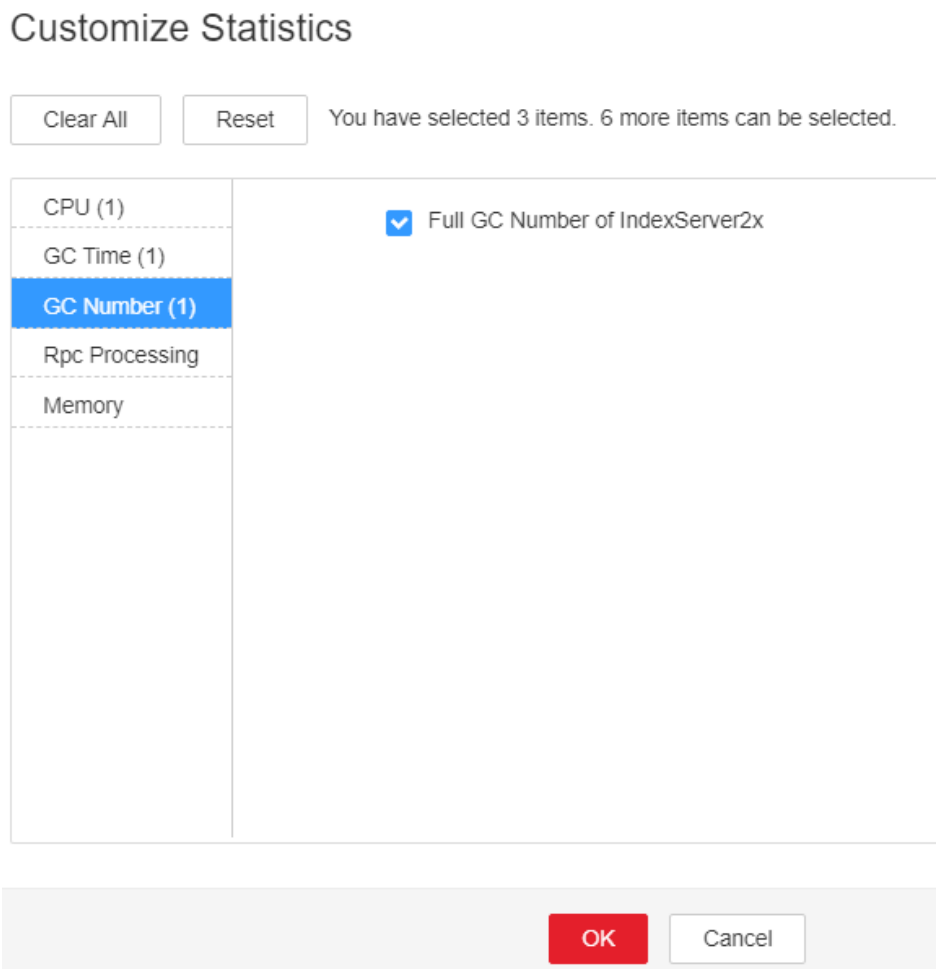

Step 3 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **Spark2x** > **Configurations** > **All Configurations** > **IndexServer2x** > **Tuning**. The default value of the **SPARK_DRIVER_MEMORY** is 4 GB. You can change the value according to the following rules: If this alarm is generated occasionally, increase the value by 0.5 times. Double the value if the alarm is reported frequently. In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Restart all IndexServer2x instances.

Step 5 After 10 minutes, check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Spark2x** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 9 Contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

After the fault is rectified, the system automatically clears this alarm.

Reference

None

10.272 ALM-44000 Presto Service Unavailable

Alarm Description

The system checks the Presto service status every 60 seconds. This alarm is generated when the system detects that Presto is unavailable.

This alarm is cleared when the Presto service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
44000	Critical	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Presto cannot run SQL queries.

Possible Causes

- The Presto coordinator or worker process is faulty.
- The network communication between Presto coordinator and worker instances is interrupted.

Handling Procedure

Step 1 Check the status of the coordinator and worker processes.

1. Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Presto**. On the page that is displayed, click the **Instance** tab. In the Presto instance list, check whether the status of all coordinator or worker instances is **Unknown**.
   - If yes, go to **2**.
   - If no, go to **1**.
2. In the upper part of the Presto instance list, choose **More** > **Restart Service** to restart the coordinator and worker processes.
3. In the alarm list, check whether ALM-44000 Presto Service Unavailable is cleared.
   - If yes, no further action is required.
   - If no, go to **1** in **Step 2**.

Step 2 Collect fault information.

1. On FusionInsight Manager, choose **System** > **Export Log**.
2. Select **Presto** for **Service**.
3. Click  in the upper right corner.

   Set **Start Time** and **End Time** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **OK**.
4. Contact the O&M engineers and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

10.273 ALM-44004 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold

Alarm Description

This alarm is generated when the system detects that the number of queuing tasks in a resource group exceeds the threshold. The system queries the number of queuing tasks in a resource group through the JMX interface. You can choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Presto** > **resource-groups** to configure a resource group. You can choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Coordinator** > **Customize** > **resourceGroupAlarm** to configure the threshold of each resource group.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
44004	Major	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

If the number of queuing tasks in a resource group exceeds the threshold, a large number of tasks may be in the queuing state. The Presto task time exceeds the expected value. When the number of queuing tasks in a resource group exceeds the maximum number (**maxQueued**) of queuing tasks in the resource group, new tasks cannot be executed.

Possible Causes

The resource group configuration is improper or too many tasks in the resource group are submitted.

Handling Procedure

Step 1 Choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Presto** > **resource-groups** to adjust the resource group configuration.

Step 2 You can choose **Components** > **Presto** > **Service Configuration** (switch **Basic** to **All**) > **Coordinator** > **Customize** > **resourceGroupAlarm** to modify the threshold of each resource group.

Step 3 Collect the fault information.

1. Log in to the cluster node based on the host name in the fault information and query the number of queuing tasks based on **Resource Group** in the additional information on the Presto client.
2. Log in to the cluster node based on the host name in the fault information, view the **/var/log/Bigdata/nodeagent/monitorlog/monitor.log** file, and search for resource group information to view the monitoring collection information of the resource group.

3. Contact O&M personnel and send the collected logs.

----End

Related Information

None

10.274 ALM-44005 Presto Coordinator Process GC Time Exceeds the Threshold

Description

The system collects GC time of the Presto Coordinator process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System** > **Configure Alarm Threshold** > **Service** > **Presto** > **Coordinator** > **Presto Process Garbage Collection Time** > **Garbage Collection Time of the Coordinator Process** on MRS Manager. This alarm is cleared when the Coordinator process GC time is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44005	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the GC time of the Coordinator process is too long, the Coordinator process running performance will be affected and the Coordinator process will even be unavailable.

Possible Causes

The heap memory of the Coordinator process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.

   NOTE

   For MRS 1.8.10 or earlier, log in to MRS Manager and choose **Alarms**.

2. Select the alarm whose **Alarm ID** is **44005** and then check the role name in **Location** and confirm the IP adress of the instance.
3. Choose **Components** > **Presto** > **Instances** > **Coordinator** (business IP address of the instance for which the alarm is generated) > **Customize** > **Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Coordinator process is longer than 5 seconds.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Presto** > **Service Configuration**, and switch **Basic** to **All**. Choose **Presto** > **Coordinator**. Increase the value of **-Xmx** (maximum heap memory) in the **JAVA_OPTS** parameter based on the site requirements.
6. Check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M personnel and send the collected logs.

----End

Reference

None

10.275 ALM-44006 Presto Worker Process GC Time Exceeds the Threshold

Description

The system collects GC time of the Presto Worker process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). You can change the threshold by choosing **System** > **Configure Alarm Threshold** > **Service** > **Presto** > **Worker** > **Presto Garbage Collection Time** > **Garbage Collection Time of the Worker Process** on MRS Manager. This alarm is cleared when the Worker process GC time is shorter than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
44006	Major	Yes

Parameter

Parameter	Description
ServiceName	Service for which the alarm is generated.
RoleName	Role for which the alarm is generated.
HostName	Host for which the alarm is generated.

Impact on the System

If the GC time of the Worker process is too long, the Worker process running performance will be affected and the Worker process will even be unavailable.

Possible Causes

The heap memory of the Worker process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Procedure

Step 1 Check the GC time.

1. Go to the cluster details page and choose **Alarms**.

   NOTE

   For MRS 1.8.10 or earlier, log in to MRS Manager and choose **Alarms**.

2. Select the alarm whose **Alarm ID** is **44006**. Then check the role name in **Location** and confirm the IP adress of the instance.
3. Choose **Components** > **Presto** > **Instances** > **Worker** (business IP address of the instance for which the alarm is generated) > **Customize** > **Presto Garbage Collection Time**. Click **OK** to view the GC time.
4. Check whether the GC time of the Worker process is longer than 5 seconds.
   - If yes, go to **Step 1.5**.
   - If no, go to **Step 2**.
5. Choose **Components** > **Presto** > **Service Configuration**, and switch **Basic** to **All**, and choose **Presto** > **Worker** Increase the value of **-Xmx** (maximum heap memory) in the **JAVA_OPTS** parameter based on the site requirements.
6. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 2**.

Step 2 Collect fault information.

1. On MRS Manager, choose **System** > **Export Log**.
2. Contact the O&M personnel and send the collected logs.

----End

Reference

None

10.276 ALM-45000 HetuEngine Service Unavailable

Alarm Description

The system checks the HetuEngine service status every 300 seconds. This alarm is generated when the HetuEngine service is unavailable.

This alarm is cleared when the HetuEngine service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45000	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

HetuEngine tasks fail to execute.

Possible Causes

- The KrbServer service is abnormal.
- The ZooKeeper service is abnormal.
- The HDFS service is abnormal.
- The Yarn service is abnormal.
- The DBService service is abnormal.
- The Hive service is abnormal.
- There are no HSBroker instances in HetuEngine.

Handling Procedure

Check the KrbServer service status.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarm**.

Step 2 In the alarm list, check whether the "ALM-25500 KrbServer Service Unavailable" alarm is generated.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Clear "ALM-25500 KrbServer Service Unavailable" according to the alarm help.

Step 4 In the alarm list, check whether the alarm "ALM-45000 HetuEngine Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the ZooKeeper service status.

Step 5 In the alarm list, check whether the alarm "ALM-12007 Process Fault" is generated.

- If yes, go to **Step 6**.
- If no, go to **Step 9**.

Step 6 In the alarm list, click ⌄ in the row that contains the "Process Fault" alarm. Check whether the name of the service for which the alarm is generated is ZooKeeper in **Location Information**.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Clear "ALM-12007 Process Fault" according to the alarm help.

Step 8 In the alarm list, check whether the alarm "ALM-45000 HetuEngine Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check the HDFS service status.

Step 9 In the alarm list, check whether the "ALM-14000 HDFS Service Unavailable" alarm is generated.

- If yes, go to **Step 10**.
- If no, go to **Step 12**.

Step 10 Clear "ALM-14000 HDFS Service Unavailable" according to the alarm help.

Step 11 In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Check the YARN service status.

Step 12 In the alarm list, check whether the "ALM-18000 YARN Service Unavailable" alarm is generated.

- If yes, go to **Step 13**.
- If no, go to **Step 15**.

Step 13 Clear "ALM-18000 YARN Service Unavailable" according to the alarm help.

Step 14 In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Check the DBService service status.

Step 15 In the alarm list, check whether the "ALM-27001 DBService Service Unavailable" alarm is generated.

- If yes, go to **Step 16**.
- If no, go to **20**.

Step 16 Clear "ALM-27001 DBService Service Unavailable" according to the alarm help.

Step 17 In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to **20**.

Check the Hive service status.

Step 18 In the alarm list, check whether the "ALM-16004 Hive Service Unavailable" alarm is generated.

- If yes, go to **Step 19**.
- If no, go to **20**.

Step 19 Clear "ALM-16004 Hive Service Unavailable" according to the alarm help.

Step 20 In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to **20**.

Check whether there are no HSBroker instances in HetuEngine.

Step 21 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HetuEngine**. On the page that is displayed, click the **Instance** tab.

Step 22 Check whether there are no HSBroker instances.

- If yes, click **Add Instance** to add one.
- If no, go to **23**.

Step 23 In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to **23**.

Check the network connection between HetuEngine and ZooKeeper, HDFS, YARN, DBService, and Hive.

Step 24 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **HetuEngine**. On the page that is displayed, click the **Instance** tab.

Step 25 Click the host name in the **HSBroker** row and record the management IP address in the **Basic Information** area.

Step 26 Log in to the host where HSBroker resides as user **omm** using the IP address obtained in **Step 25**.

Step 27 Run the **ping** command to check whether the network connection between the host where HSBroker resides and the hosts where ZooKeeper, HDFS, Yarn, DBService, and Hive reside is in the normal state.

- If yes, go to **Step 30**.
- If no, go to **Step 28**.

Step 28 Contact the network administrator to restore the network.

Step 29 In the alarm list, check whether the "ALM-45000 HetuEngine Service Unavailable" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 30**.

Collect fault information.

Step 30 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 31 Expand the **Service** drop-down list. In the **Services** dialog box that is displayed, select **HetuEngine** under the target cluster name, and click **OK**.

Step 32 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 33 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 34 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault is rectified, the system automatically clears this alarm.

Reference

None

10.277 ALM-45001 Faulty HetuEngine Compute Instances

This alarm applies only to MRS 3.2.0 or later.

Alarm Description

The system checks the HetuEngine compute instance status every 60 seconds. This alarm is generated when a HetuEngine compute instance is faulty.

This alarm is cleared when all faulty HetuEngine compute instances are restored.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45001	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

HetuEngine tasks fail to execute.

Possible Causes

- The HDFS service is abnormal.

- The Yarn service is abnormal.
- Yarn queue resources are insufficient.
- The process of compute instances is faulty.

Handling Procedure

Check the HDFS service status.

Step 1 In the alarm list, check whether the "ALM-14000 HDFS Service Unavailable" alarm is generated.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Clear "ALM-14000 HDFS Service Unavailable" according to the alarm help.

Step 3 In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Check the YARN service status.

Step 4 In the alarm list, check whether the "ALM-18000 YARN Service Unavailable" alarm is generated.

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Clear "ALM-18000 YARN Service Unavailable" according to the alarm help.

Step 6 In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Check the YARN queue resource status.

Step 7 In the alarm list, check whether the "ALM-18022 Insufficient YARN Queue Resources" alarm is generated.

- If yes, go to **8**.
- If no, go to **Step 10**.

Step 8 Clear "ALM-18022 Insufficient YARN Queue Resources" according to the alarm help.

Step 9 In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check the HetuEngine compute instance status.

Step 10 Log in to FusionInsight Manager as an administrator who can access the HetuEngine web UI and choose **Cluster** > **Services** > **HetuEngine**.

Step 11 In the **Basic Information** area on the **Dashboard** tab page, click the link next to **HSConsole WebUI** to access the HSConsole page.

Step 12 On the compute instance page, check whether any compute instances are in the **FAULT** state.

- If yes, go to **Step 13**.
- If no, go to **Step 14**.

Step 13 In the **Operation** column of the target compute instance, click **Start** and wait until the instance is started.

Step 14 In the alarm list, check whether the "ALM-45001 Faulty HetuEngine Compute Instances" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 15**.

Collect fault information.

Step 15 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 16 Expand the **Service** drop-down list. In the **Services** dialog box that is displayed, select **HetuEngine** under the target cluster name, and click **OK**.

Step 17 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 18 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 19 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.278 ALM-45003 HetuEngine QAS Disk Capacity Is Insufficient

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the HetuEngine QAS disk usage every 60 seconds and compares the actual disk usage with the threshold. The disk usage has a default threshold. This alarm is generated if the disk usage exceeds the threshold.

To change the threshold, choose **O&M** > **Alarm** > **Thresholds**. In the service list, choose **HetuEngine** > **Disk** > **QAS Disk Usage (QAS)**.

If the **Trigger Count** is **1**, this alarm is cleared when the usage of the HetuEngine QAS disk is less than or equal to the threshold. If the **Trigger Count** is greater than **1**, this alarm is cleared when the disk usage is less than or equal to 80% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45003	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
PartitionName	Specifies the disk partition for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the disk capacity is insufficient, QAS fails to write data, affecting SQL diagnosis and automatic recommendation of materialized views.

Possible Causes

- The alarm threshold is improperly configured.
- The configuration of the HetuEngine QAS disk cannot meet service requirements. The disk usage reaches the upper limit.

Handling Procedure

Check whether the threshold is set properly.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. In the service list, choose **HetuEngine** > **Disk** > **QAS Disk Usage (QAS)**. Check whether the alarm threshold is set properly. The default threshold is 80% of the disk capacity. You can change the threshold as required.

- If the threshold is set properly, go to **Step 4**.
- If the threshold is not set properly, go to **Step 2**.

Step 2 Click **Modify** in the **Operation** column to modify and save the alarm threshold as required.

Step 3 Wait 2 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm is not cleared, go to **Step 4**.

Check whether the disk usage reaches the upper limit.

Step 4 Expand the alarm information, view the information in the **Location** area, and check the role name and host name of the QAS disk where the alarm is generated.

Step 5 Choose **Cluster** > **Services** > **HetuEngine** and click **Instance**. On the displayed page, click the QAS role name in the alarm information. On the instance page that is displayed, click **Chart** and check whether the QAS disk usage in the **QAS Disk Usage** chart exceeds the threshold (80% of the disk capacity by default).

- If the disk usage reaches the upper limit, go to **Step 6**.
- If the disk usage does not reaches the upper limit, go to **Step 9**.

Step 6 Log in to the host of the node where the QAS instance reporting the alarm is located as the **root** user.

Step 7 Run the following command to go to the QAS data directory and delete temporary files as required:

cd ${BIGDATA_DATA_HOME}/hetuengine/qas

NOTICE

Deleting temporary files affects the latest QAS execution result but does not affect subsequent results.

Step 8 Wait 2 minutes and check whether the alarm is cleared.

- If the alarm is cleared, no further action is required.
- If the alarm fails to be cleared, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, select **HetuEngine** for the target cluster, and click **OK**.

Step 11 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 12 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.279 ALM-45175 Average Time for Calling OBS Metadata APIs Is Greater than the Threshold

Alarm Description

The system checks whether the average duration for calling OBS metadata APIs is greater than the threshold every 30 seconds. This alarm is generated when the number of consecutive times that the average time exceeds the specified threshold is greater than the number of smoothing times.

This alarm is automatically cleared when the average duration for calling the OBS metadata APIs is lower than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45175	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Parameter	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the average time for calling the OBS metadata APIs exceeds the threshold, the upper-layer big data computing services may be affected. To be more specific, the execution time of some computing tasks will exceed the threshold.

Possible Causes

Frame freezing occurs on the OBS server, or the network between the OBS client and the OBS server is unstable.

Handling Procedure

Check the heap memory usage.

Step 1 On the **FusionInsight Manager** homepage, choose **O&M** > **Alarm** > **Alarms** > **Average Time for Calling the OBS Metadata API Exceeds the Threshold**, view the role name in **Location**, and check the instance IP address.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **meta** > **Instance** > **meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Average time of OBS interface calls** from **OBS Meta data Operations**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Figure 10-109 Average duration for calling the OBS metaData API

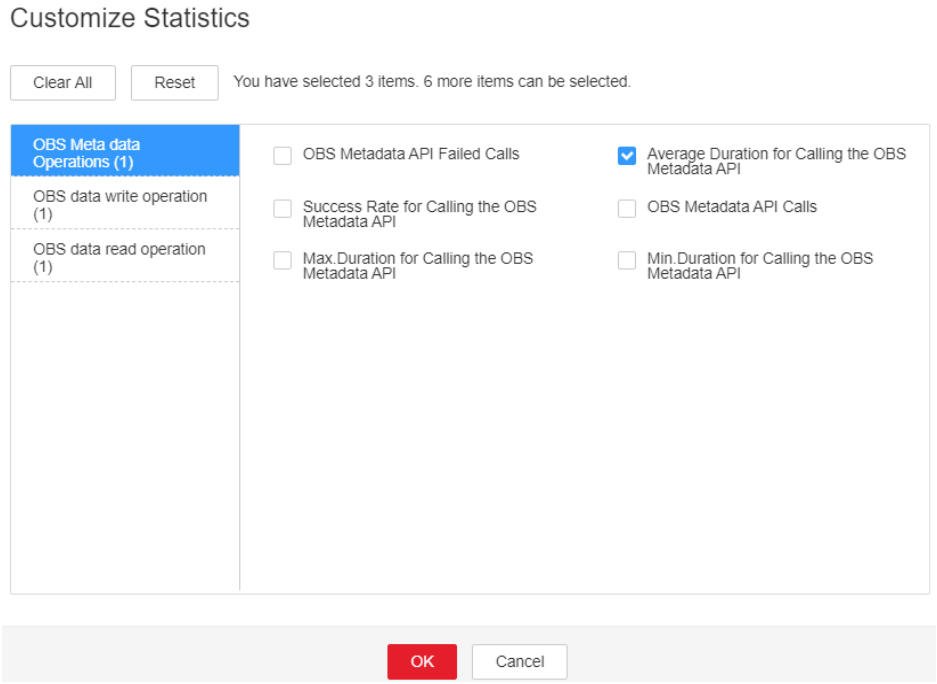

Step 3 Choose **Cluster** > *Name of the desired cluster* > **O&M** > **Alarm** > **Thresholds** > **meta** > **Average Time for Calling the OBS Metadata API**. Increase the threshold or smoothing times as required.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.280 ALM-45176 Success Rate of Calling OBS Metadata APIs Is Lower than the Threshold

Alarm Description

The system checks whether the success rate of calling OBS metadata APIs is lower than the threshold every 30 seconds. This alarm is generated when the success rate is lower than the threshold.

This alarm is automatically cleared when the success rate of calling APIs for writing OBS data is greater than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45176	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the success rate of calling the OBS metadata APIs is less than the threshold, the upper-layer big data computing services may be affected. To be more specific, some computing tasks may fail to be executed.

Possible Causes

An execution exception or severe timeout occurs on the OBS server.

Handling Procedure

Check the heap memory usage.

Step 1 On the **FusionInsight Manager** homepage, choose **O&M** > **Alarm** > **Alarms** > **Success Rate for Calling the OBS Metadata API Is Lower Than the Threshold**, view the role name in **Location**, and check the instance IP address.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **meta** > **Instance** > **meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Success percent of OBS interface calls** from **OBS Meta data Operations**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Figure 10-110 Successful rate for calling the OBS API

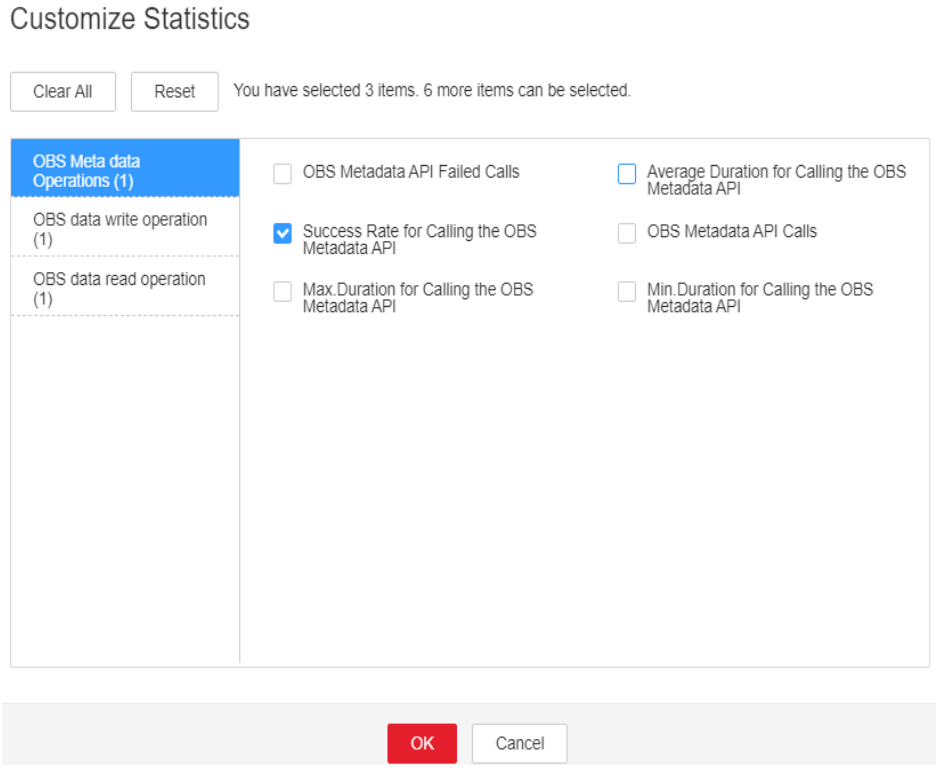

Step 3 Choose **Cluster** > *Name of the desired cluster* > **O&M** > **Alarm** > **Thresholds** > **meta** > **Success Rate for Calling the OBS Metadata API**. Increase the threshold or smoothing times as required.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.281 ALM-45177 Success Rate of Calling OBS Data Read APIs Is Lower than the Threshold

Alarm Description

The system checks whether the success rate of calling APIs for reading OBS data is lower than the threshold every 30 seconds. This alarm is generated when the success rate is lower than the threshold.

This alarm is automatically cleared when the success rate of calling APIs for reading OBS data is greater than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45177	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the success rate of calling the OBS APIs for reading data is less than the threshold, the upper-layer big data computing services may be affected. To be more specific, some computing tasks may fail to be executed.

Possible Causes

An execution exception or severe timeout occurs on the OBS server.

Handling Procedure

Check the heap memory usage.

Step 1 On the **FusionInsight Manager** homepage, choose **O&M** > **Alarm** > **Alarms** > **Success Rate for Calling the OBS Data Read API Is Lower Than the Threshold**, view the role name in **Location**, and check the instance IP address.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **meta** > **Instance** > **meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Success percent of OBS data read operation interface calls** from **OBS data read operation**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Figure 10-111 Success rate for calling the OBS data read API

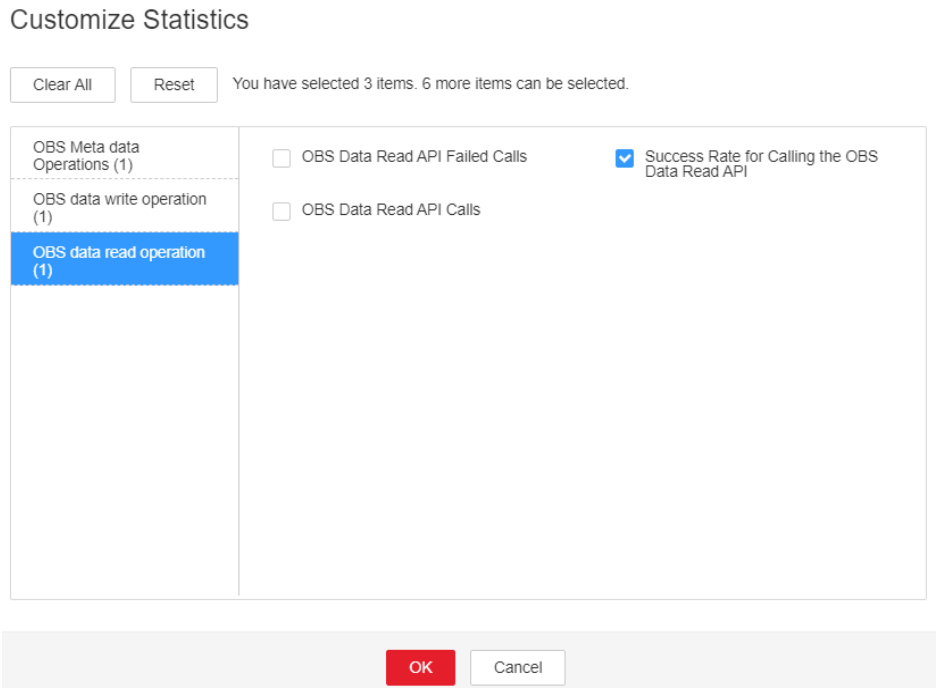

Step 3 Choose **Cluster** > *Name of the desired cluster* > **O&M** > **Alarm** > **Thresholds** > **meta** > **Success Rate for Calling the OBS Data Read API**. Increase the threshold or smoothing times as required.

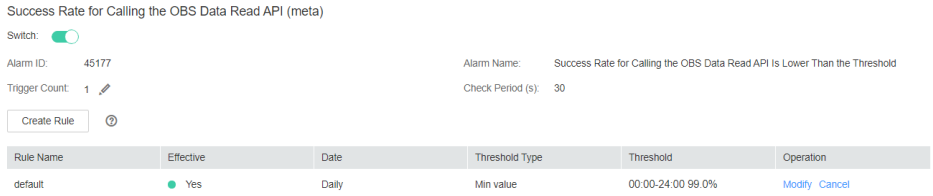

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.282 ALM-45178 Success Rate of Calling OBS Data Write APIs Is Lower Than the Threshold

Alarm Description

The system checks whether the success rate of calling APIs for writing OBS data is lower than the threshold every 30 seconds. This alarm is generated when the success rate is lower than the threshold.

This alarm is automatically cleared when the success rate of calling APIs for writing OBS data is greater than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45178	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the success rate of calling the OBS APIs for writing data is lower than the threshold, the upper-layer big data computing services may be affected. To be more specific, some computing tasks may fail to be executed.

Possible Causes

An execution exception or severe timeout occurs on the OBS server.

Handling Procedure

Check the heap memory usage.

Step 1 On the **FusionInsight Manager** homepage, choose **O&M** > **Alarm** > **Alarms** > **Success Rate for Calling the OBS Data Write API Is Lower Than the Threshold**, view the role name in **Location**, and check the instance IP address.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Services** > **meta** > **Instance** > **meta** (IP address of the instance for which the alarm is generated). Click the drop-down list in the upper right corner of the chart area and choose **Customize**. In the dialog box that is displayed, select **Success percent of OBS data write operation interface calls** from **OBS data write operation**, and click **OK**. Check whether the average time of OBS metadata API calls exceeds the threshold.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Figure 10-112 Success rate for calling the OBS data write API

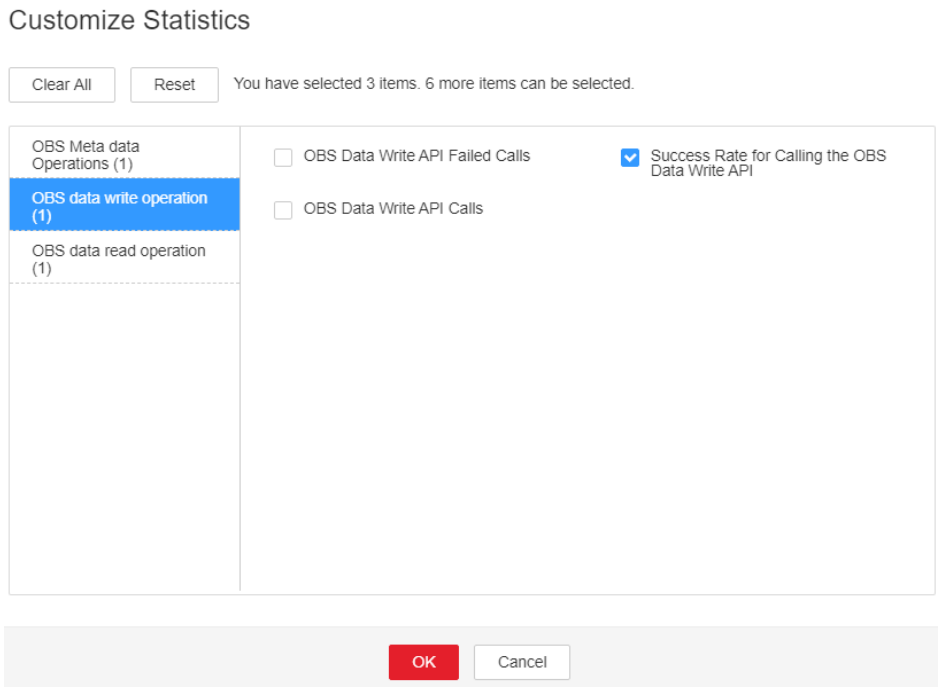

Step 3 Choose **Cluster** > *Name of the desired cluster* > **O&M** > **Alarm** > **Thresholds** > **meta** > **Success Rate for Calling the OBS Data Write API**. Increase the threshold or smoothing times as required.

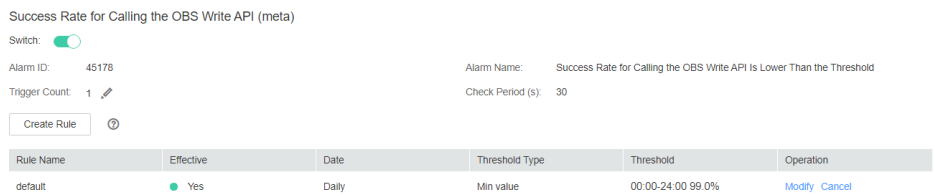

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 In the **Services** area, select **NodeAgent**, **NodeMetricAgent**, **OmmServer**, and **OmmAgent** under OMS.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.283 ALM-45179 Number of Failed OBS readFully API Calls Exceeds the Threshold

Alarm Description

The system checks whether the number of failed OBS readFully API calls exceeds the threshold every 30 seconds. This alarm is generated when the number of failed API calls exceeds the threshold.

This alarm is automatically cleared when the number of failed OBS readFully API calls is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45179	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

Possible Causes

An execution exception or severe timeout occurs on the OBS server.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **meta** > **Number of failed calls to the OBS readFully interface**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.

Step 2 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Contact OBS O&M personnel to check whether the OBS service is normal.

- If yes, go to **Step 4**.
- If no, contact OBS O&M personnel to restore the OBS service.

Collect fault information.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS data read operation** in the **Chart Category** area. In the **Number of failed calls to the OBS readFully interface-All Instances** chart, view the host name of the instance that has the maximum number of failed OBS readFully API calls. For example, the host name is **node-ana-corevpeO003**.

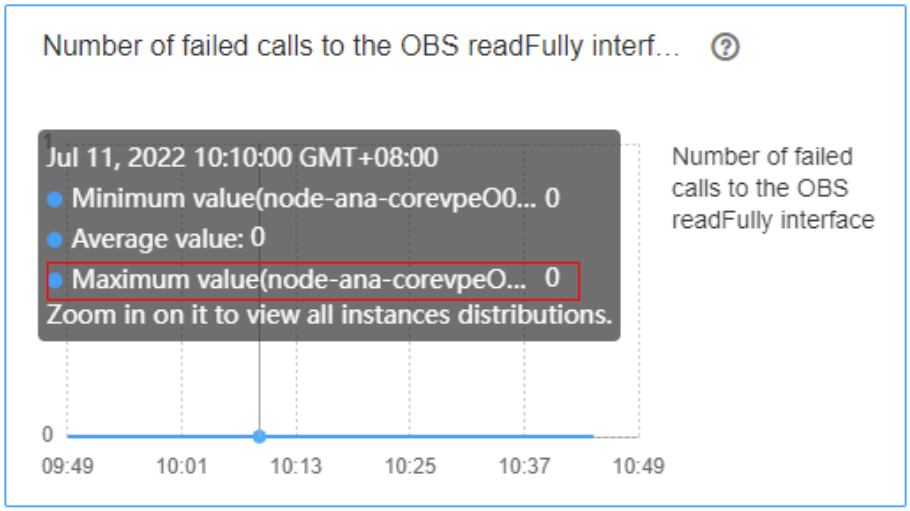

Step 5 Choose **O&M** > **Log** > **Download** and select **meta** and **meta** under it for **Service**.

Step 6 Select the host obtained in **Step 4** for **Hosts**.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.284 ALM-45180 Number of Failed OBS read API Calls Exceeds the Threshold

Alarm Description

The system checks whether the number of failed OBS read API calls exceeds the threshold every 30 seconds. This alarm is generated when the number of failed API calls exceeds the threshold.

This alarm is automatically cleared when the number of failed OBS read API calls is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45180	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

Possible Causes

An execution exception or severe timeout occurs on the OBS server.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **meta** > **Number of failed calls to the OBS read interface**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.

Step 2 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Contact OBS O&M personnel to check whether the OBS service is normal.

- If yes, go to **Step 4**.
- If no, contact OBS O&M personnel to restore the OBS service.

Collect fault information.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS data read operation** in the **Chart Category** area. In the **Number of failed calls to the OBS read interface-All Instances** chart, view the host name of the instance that has the maximum number of failed OBS read API calls. For example, the host name is **node-ana-corevpeO003**.

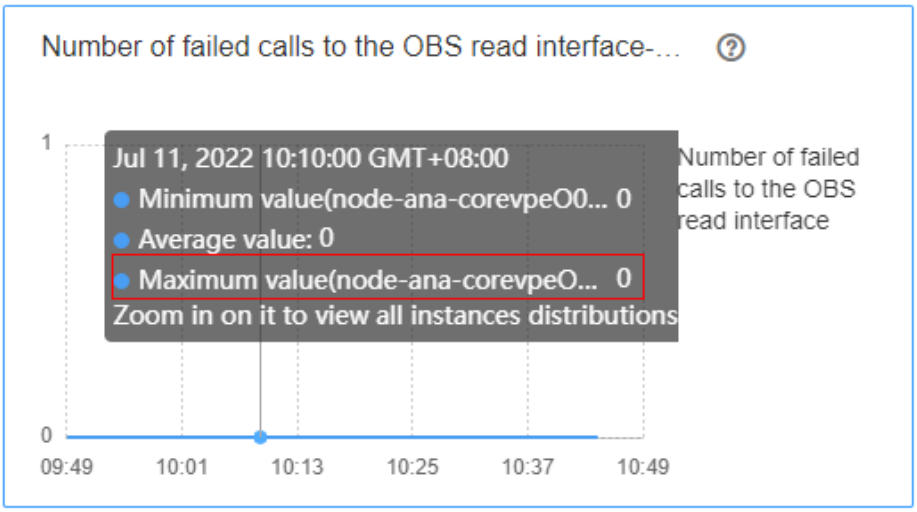

Step 5 Choose **O&M** > **Log** > **Download** and select **meta** and **meta** under it for **Service**.

Step 6 Select the host obtained in **Step 4** for **Hosts**.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.285 ALM-45181 Number of Failed OBS write API Calls Exceeds the Threshold

Alarm Description

The system checks whether the number of failed OBS write API calls exceeds the threshold every 30 seconds. This alarm is generated when the number of failed API calls exceeds the threshold.

This alarm is automatically cleared when the number of failed OBS write API calls is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45181	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

Possible Causes

An execution exception or severe timeout occurs on the OBS server.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **meta** > **Number of failed calls to the OBS write interface**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.

Step 2 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Contact OBS O&M personnel to check whether the OBS service is normal.

- If yes, go to **Step 4**.
- If no, contact OBS O&M personnel to restore the OBS service.

Collect fault information.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS data write operation** in the **Chart Category** area. In the **Number of failed calls to the OBS write interface-All Instances** chart, view the host name of the instance that has the maximum number of failed OBS write API calls. For example, the host name is **node-ana-corevpeO003**.

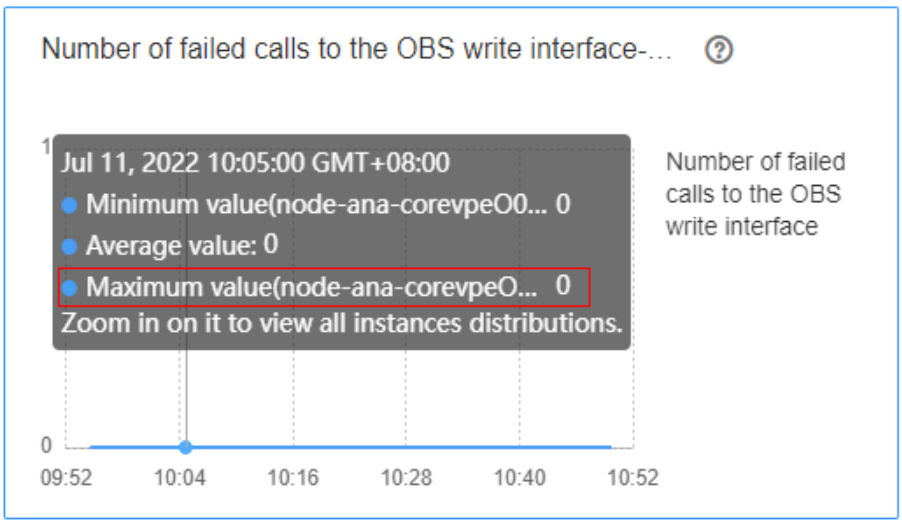

Step 5 Choose **O&M** > **Log** > **Download** and select **meta** and **meta** under it for **Service**.

Step 6 Select the host obtained in **Step 4** for **Hosts**.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.286 ALM-45182 Number of Throttled OBS Operations Exceeds the Threshold

Alarm Description

The system checks whether the number of throttled OBS operations exceeds the threshold every 30 seconds. This alarm is generated when the number of throttled OBS operations exceeds the threshold.

This alarm is automatically cleared when the number of throttled OBS operations is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45182	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Certain upper-layer big data computing tasks will fail to execute.

Possible Causes

The frequency of requesting OBS APIs is too high.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **meta** > **Number of Throttled OBS Operations**. In the right pane, set **Threshold** or **Trigger Count** to a larger value as required.

Step 2 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Step 3 Contact OBS O&M personnel to check whether the OBS service is normal.

- If yes, go to **Step 4**.
- If no, contact OBS O&M personnel to restore the OBS service.

Collect fault information.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **meta**. On the page that is displayed, click the **Chart** tab. On this tab page, select **OBS Throttled** in the **Chart Category** area. In the **Number of Throttled OBS Operations-All Instances** chart, view the host name of the instance that has the maximum number of throttled OBS API calls. For example, the host name is **node-ana-coreUQqJ0002**.

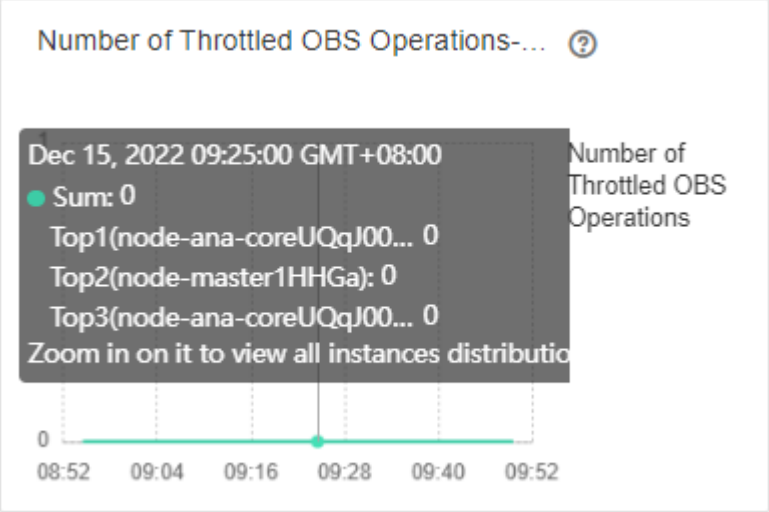

Step 5 Choose **O&M** > **Log** > **Download** and select **meta** and **meta** under it for **Service**.

Step 6 Select the host obtained in **Step 4** for **Hosts**.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.287 ALM-45275 Ranger Service Unavailable

Alarm Description

The alarm module checks the Ranger service status every 180 seconds. This alarm is generated if the Ranger service is abnormal.

This alarm is cleared after the Ranger service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45275	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

When the Ranger service is unavailable, Ranger cannot work properly and the native Ranger UI cannot be accessed.

Possible Causes

- The DBService service on which Ranger depends is abnormal.
- The RangerAdmin role instance is abnormal.

Handling Procedure

Check the DBService process status.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed page, check whether the ALM-27001 DBService Service Unavailable alarm is reported.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 Rectify the DBService service fault by following the handling procedure of ALM-27001 DBService Service Unavailable. After the DBService alarm is cleared, check whether Ranger Service Unavailable alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Check all RangerAdmin instances.

Step 3 Log in to the node where the RangerAdmin instance is located as user **omm** and run the **ps -ef|grep "proc_rangeradmin"** command to check whether the RangerAdmin process exists on the current node.

- If yes, go to **Step 5**.
- If no, restart the faulty RangerAdmin instance or Ranger service and go to **Step 4**.

Step 4 In the alarm list, check whether the alarm "Ranger Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

Related Information

None

10.288 ALM-45276 Abnormal RangerAdmin Status

Alarm Description

The alarm module checks the RangerAdmin service status every 60 seconds. This alarm is generated if RangerAdmin is unavailable.

This alarm is automatically cleared after the RangerAdmin service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45276	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Role for which the alarm is generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

If the status of a RangerAdmin is abnormal, access to the Ranger native UI is not affected. If there are two abnormal RangerAdmin instances, the Ranger native UI cannot be accessed and operations such as creating, modifying, and deleting policies are unavailable.

Possible Causes

The RangerAdmin port is not started.

Handling Procedure

Check the port process.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the name of the host for which the alarm is generated.

Step 2 Log in to the node where the RangerAdmin instance is located as user **omm**. Run the **ps -ef|grep "proc_rangeradmin" | grep -v grep | awk -F ' ' '{print $2}'** command to obtain ***pid*** of the RangerAdmin process, and run the **netstat -anp| grep *pid* | grep LISTEN** command to check whether the RangerAdmin process listens to port 21401 in the security mode and port 21400 in standard mode.

- If yes, go to **Step 4**.
- If no, restart the faulty RangerAdmin instance or Ranger service and go to **Step 3**.

Step 3 In the alarm list, check whether the "Abnormal RangerAdmin Status" alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect the fault information.

Step 4 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 6 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

Related Information

None

10.289 ALM-45277 RangerAdmin Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the RangerAdmin service every 60 seconds. This alarm is generated when the system detects that the heap memory usage of the RangerAdmin instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45277	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Heap memory overflow may cause service breakdown.

Possible Causes

The heap memory usage of the RangerAdmin instance is high or the heap memory is improperly allocated.

Handling Procedure

Check the heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45277 RangerAdmin Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **RangerAdmin Heap Memory Usage**. Click **OK**.

Figure 10-113 RangerAdmin heap memory usage

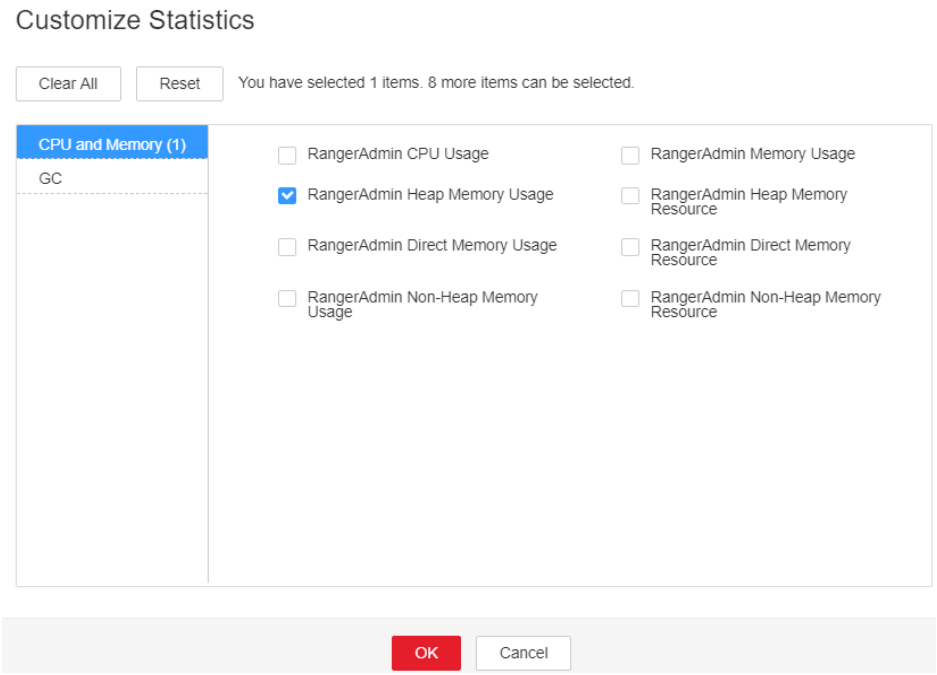

Step 3 Check whether the heap memory used by RangerAdmin reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **RangerAdmin** > **Instance Configuration**. Click **All Configurations**, and choose **RangerAdmin** > **System**. Increase the value of **-Xmx** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for RangerAdmin cannot meet the heap memory required by the RangerAdmin process. You are advised to check the heap memory usage of RangerAdmin and change the value of **-Xmx** in **GC_OPTS** to the twice of the heap memory used by RangerAdmin. The value can be changed based on the actual service scenario. For details, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.290 ALM-45278 RangerAdmin Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the RangerAdmin service every 60 seconds. This alarm is generated when the direct memory usage of the RangerAdmin instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the direct memory usage of RangerAdmin is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45278	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service breakdown.

Possible Causes

The direct memory of the RangerAdmin instance is overused or the direct memory is inappropriately allocated. As a result, the memory usage exceeds the threshold.

Handling Procedure

Check the direct memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45278 RangerAdmin Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **RangerAdmin Direct Memory Usage**. Click **OK**.

Figure 10-114 RangerAdmin direct memory usage

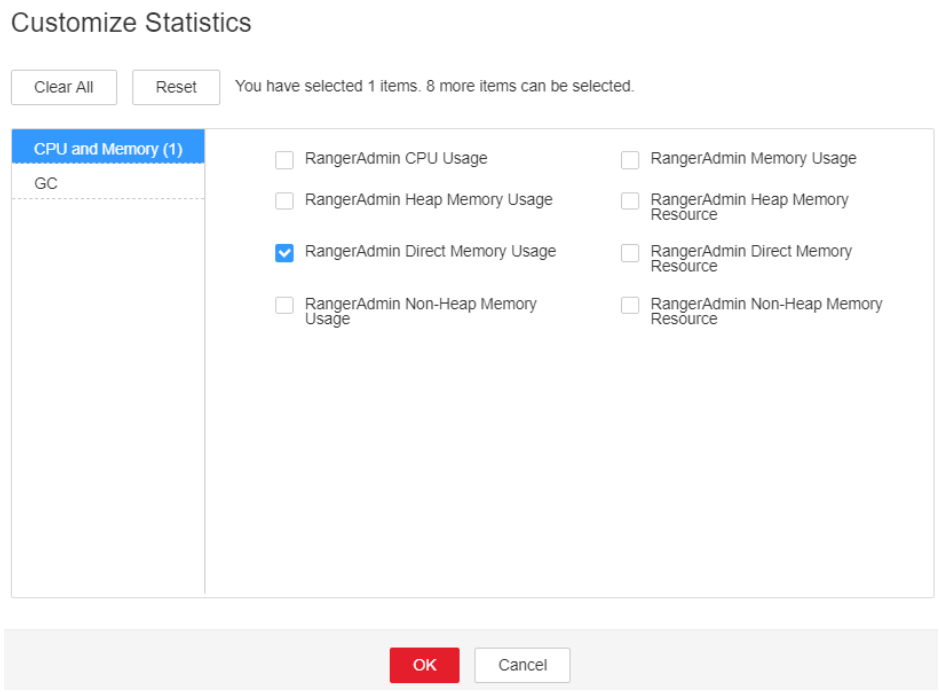

Step 3 Check whether the direct memory used by RangerAdmin reaches the threshold (80% of the maximum direct memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **RangerAdmin** > **Instance Configuration**. Click **All Configurations**, and choose

RangerAdmin > **System**. Increase the value of **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the direct memory configured for RangerAdmin cannot meet the direct memory required by the RangerAdmin process. You are advised to check the direct memory usage of RangerAdmin and change the value of **-XX:MaxDirectMemorySize** in **GC_OPTS** to the twice of the direct memory used by RangerAdmin. You can change the value based on the actual service scenario. For details, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.291 ALM-45279 RangerAdmin Non Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the non-heap memory usage of the RangerAdmin service every 60 seconds. This alarm is generated when the non-heap memory usage of the RangerAdmin instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45279	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Non-heap memory overflow may cause service breakdown.

Possible Causes

The non-heap memory usage of the RangerAdmin instance is high or the non-heap memory is improperly allocated.

Handling Procedure

Check non-heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45279 RangerAdmin Non Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **RangerAdmin Non Heap Memory Usage**. Click **OK**.

Figure 10-115 RangerAdmin non-heap memory usage

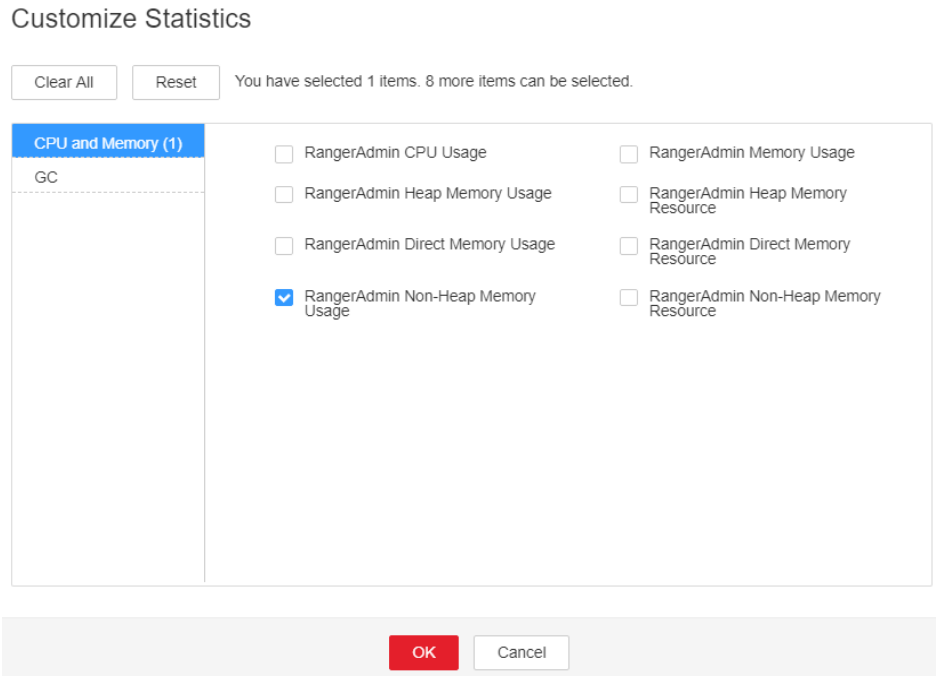

Step 3 Check whether the non-heap memory used by RangerAdmin reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **RangerAdmin** > **Instance Configuration**. Click **All Configurations**, and choose **RangerAdmin** > **System**. Set **-XX: MaxPermSize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the non-heap memory size configured for the RangerAdmin instance cannot meet the non-heap memory required by the RangerAdmin process. You are advised to change the value of **-XX:MaxPermSize** in **GC_OPTS** to the twice of the current non-heap memory usage or change the value based on the site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.292 ALM-45280 RangerAdmin GC Duration Exceeds the Threshold

Alarm Description

The system checks the GC duration of the RangerAdmin process every 60 seconds. This alarm is generated when the GC duration of the RangerAdmin process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45280	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The RangerAdmin responds slowly.

Possible Causes

The heap memory of the RangerAdmin instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Handling Procedure

Check the GC duration.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45280 RangerAdmin GC Duration Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize** > **GC** > **RangerAdmin GC Duration**. Click **OK**.

Figure 10-116 RangerAdmin garbage collection (GC) duration

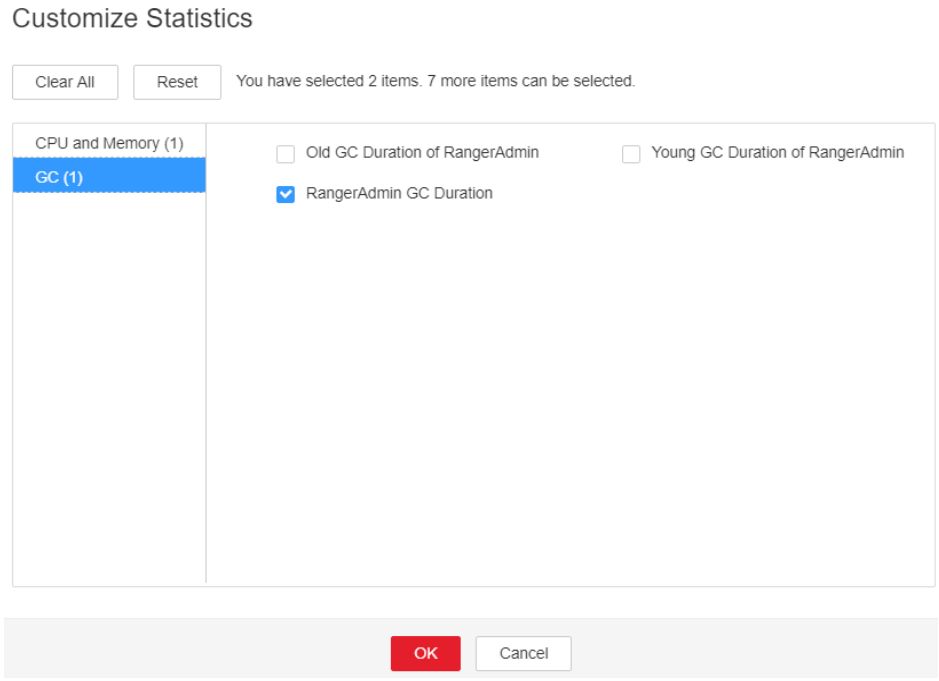

Step 3 Check whether the GC duration of the RangerAdmin process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **RangerAdmin** > **Instance Configuration**. Click **All Configurations**, and choose **RangerAdmin** > **System**. Increase the value of **-Xmx** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for RangerAdmin cannot meet the heap memory required by the RangerAdmin process. You are advised to check the heap memory usage of RangerAdmin and change the value of **-Xmx** in **GC_OPTS** to the twice of the heap memory used by RangerAdmin. The value can be changed based on the actual service scenario. For details, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.293 ALM-45281 UserSync Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the UserSync service every 60 seconds. This alarm is generated when the system detects that the heap memory usage of the UserSync instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45281	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Heap memory overflow may cause service breakdown.

Possible Causes

The heap memory usage of the UserSync instance is high or the heap memory is improperly allocated.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45281 UserSync Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **UserSync Heap Memory Usage**. Click **OK**.

Figure 10-117 UserSync heap memory usage

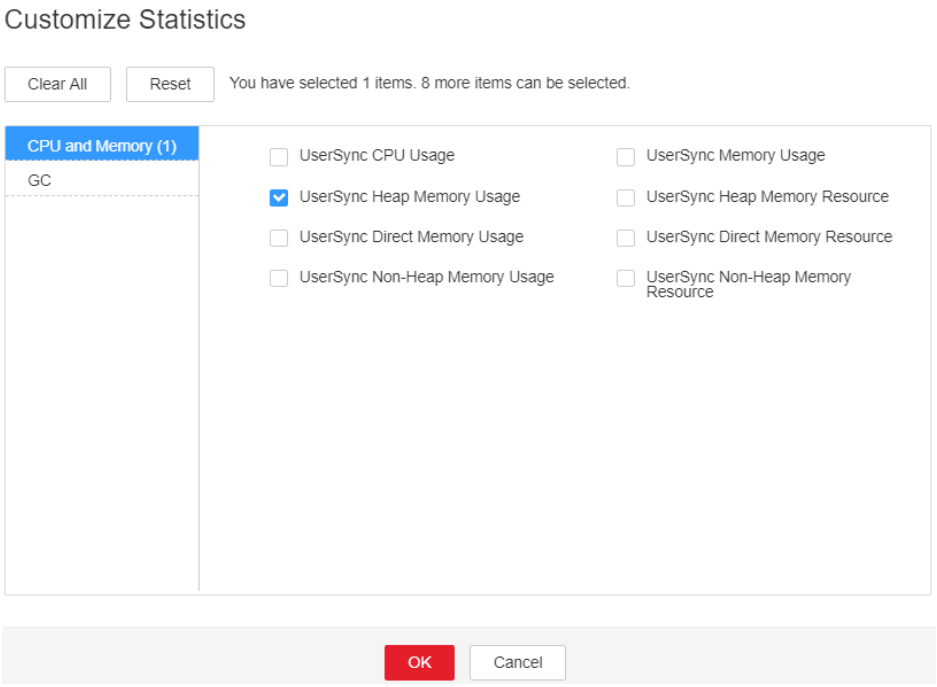

Step 3 Check whether the heap memory used by UserSync reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **UserSync** > **Instance Configuration**. Click **All Configurations**, and choose **UserSync** > **System**. Increase the value of **-Xmx** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for UserSync cannot meet the heap memory required by the UserSync process. You are advised to change the **-Xmx** value of **GC_OPTS** to twice that of the heap memory used by UserSync. You can change the value based on the actual service scenario. For details about how to check the UserSync heap memory usage, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.294 ALM-45282 UserSync Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the UserSync service every 60 seconds. This alarm is generated when the direct memory usage of the UserSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the UserSync direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45282	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service breakdown.

Possible Causes

The direct memory of the UserSync instance is overused or the direct memory is inappropriately allocated. As a result, the memory usage exceeds the threshold.

Handling Procedure

Check the direct memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45282 UserSync Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm. Check the name of the instance host for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **UserSync Direct Memory Usage**. Click **OK**.

Figure 10-118 UserSync direct memory usage

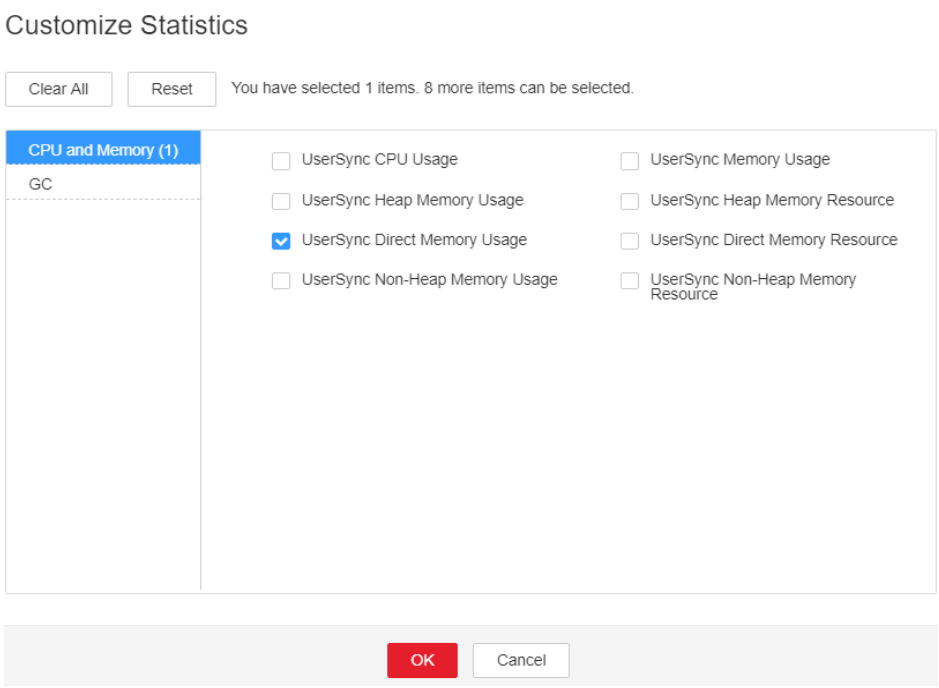

Step 3 Check whether the direct memory used by the UserSync reaches the threshold (80% of the maximum direct memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **UserSync** > **Instance Configuration**. Click **All Configurations**, and choose

UserSync > **System**. Increase the value of **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the direct memory configured for UserSync cannot meet the direct memory required by the UserSync process. You are advised to check the direct memory usage of UserSync and change the value of **-XX:MaxDirectMemorySize** in **GC_OPTS** to the twice of the direct memory used by UserSync. You can change the value based on the actual service scenario. For details, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.295 ALM-45283 UserSync Non Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the non-heap memory usage of the UserSync service every 60 seconds. This alarm is generated when the non-heap memory usage of the UserSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45283	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Non-heap memory overflow may cause service breakdown.

Possible Causes

The non-heap memory of the UserSync process is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45283 UserSync Non Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **UserSync Non Heap Memory Usage**. Click **OK**.

Figure 10-119 UserSync non-heap memory usage

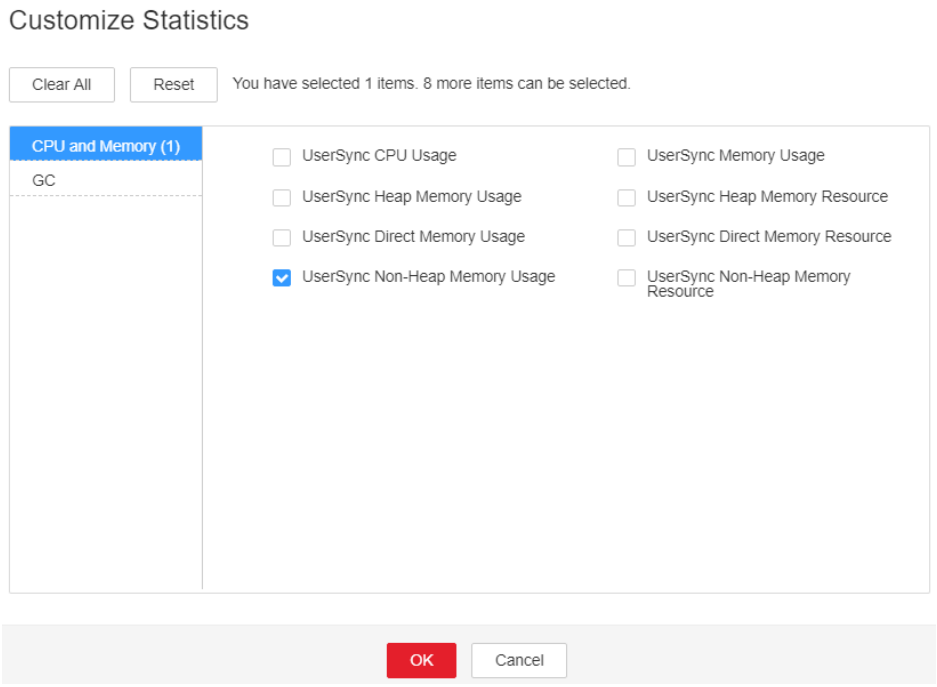

Step 3 Check whether the non-heap memory used by UserSync reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **UserSync** > **Instance Configuration**. Click **All Configurations**, and choose **UserSync** > **System**. Set **-XX: MaxPermSize** in the **GC_OPTS** parameter to a larger value based on site requirements and click **Save** to save the configuration.

NOTE

If this alarm is generated, the non-heap memory size configured for the UserSync instance cannot meet the non-heap memory required by the UserSync process. You are advised to change the **-XX:MaxPermSize** value of **GC_OPTS** to twice that of the current non-heap memory size or change the value based on the site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.296 ALM-45284 UserSync Garbage Collection (GC) Time Exceeds the Threshold

Alarm Description

The system checks the GC duration of the UserSync process every 60 seconds. This alarm is generated when the GC duration of the UserSync process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45284	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Threshold for triggering the alarm.

Impact on the System

UserSync responds slowly.

Possible Causes

The heap memory of the UserSync instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Handling Procedure

Check the GC time.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45284 UserSync Garbage Collection (GC) Time Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize** > **GC** > **UserSync GC Duration**. Click **OK**.

Figure 10-120 UserSync GC Duration

Step 3 Check whether the GC duration of the UserSync process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **UserSync** > **Instance Configuration**. Click **All Configurations**, and choose **UserSync** > **System**. Increase the value of **-Xmx** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for UserSync cannot meet the heap memory required by the UserSync process. You are advised to change the value of **GC_OPTS** to the twice that of the heap memory used by UserSync. You can change the value based on the actual service scenario. For details about how to check the UserSync heap memory usage, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the fault that triggers the alarm is rectified, the alarm is automatically cleared.

Related Information

None

10.297 ALM-45285 TagSync Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the TagSync service every 60 seconds. This alarm is generated when the heap memory usage of the TagSync instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45285	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Heap memory overflow may cause service breakdown.

Possible Causes

The heap memory usage of the TagSync instance is high or the heap memory is improperly allocated.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45285 TagSync Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **TagSync Heap Memory Usage**. Click **OK**.

Figure 10-121 TagSync heap memory usage

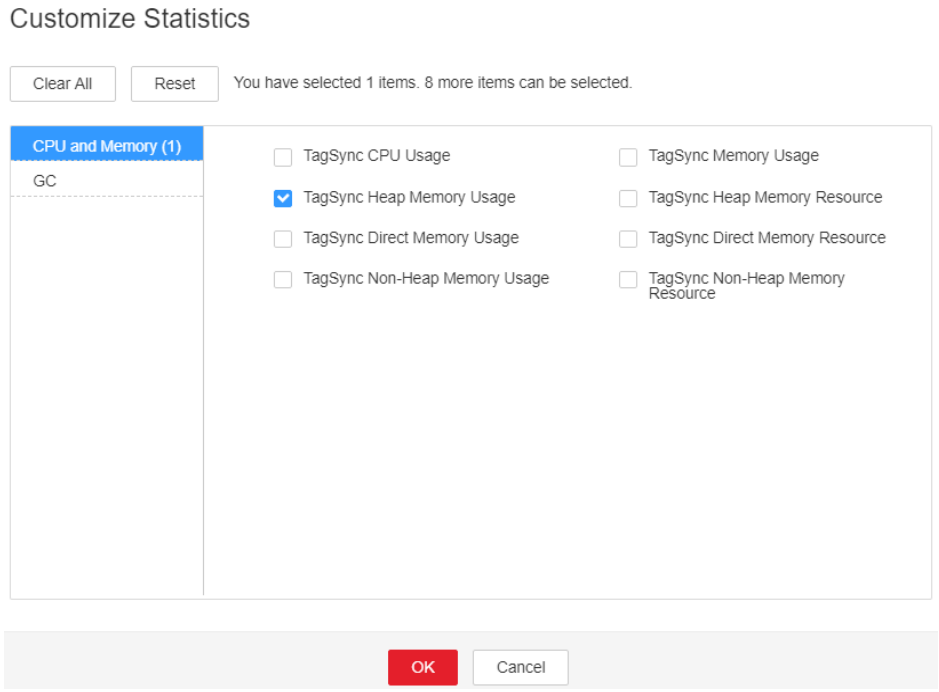

Step 3 Check whether the heap memory used by TagSync reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **TagSync** > **Instance Configuration**. Click **All Configurations** and choose **TagSync** > **System**. Increase the value of **-Xmx** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for TagSync cannot meet the heap memory required by the TagSync process. You are advised to change the **-Xmx** value of **GC_OPTS** to twice that of the heap memory used by TagSync. You can change the value based on the actual service scenario. For details about how to check the TagSync heap memory usage, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.298 ALM-45286 TagSync Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the TagSync service every 60 seconds. This alarm is generated when the direct memory usage of the TagSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the TagSync direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45286	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service breakdown.

Possible Causes

The direct memory of the TagSync instance is overused or the direct memory is inappropriately allocated. As a result, the memory usage exceeds the threshold.

Handling Procedure

Check the direct memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45286 TagSync Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **TagSync Direct Memory Usage**. Click **OK**.

Figure 10-122 TagSync direct memory usage

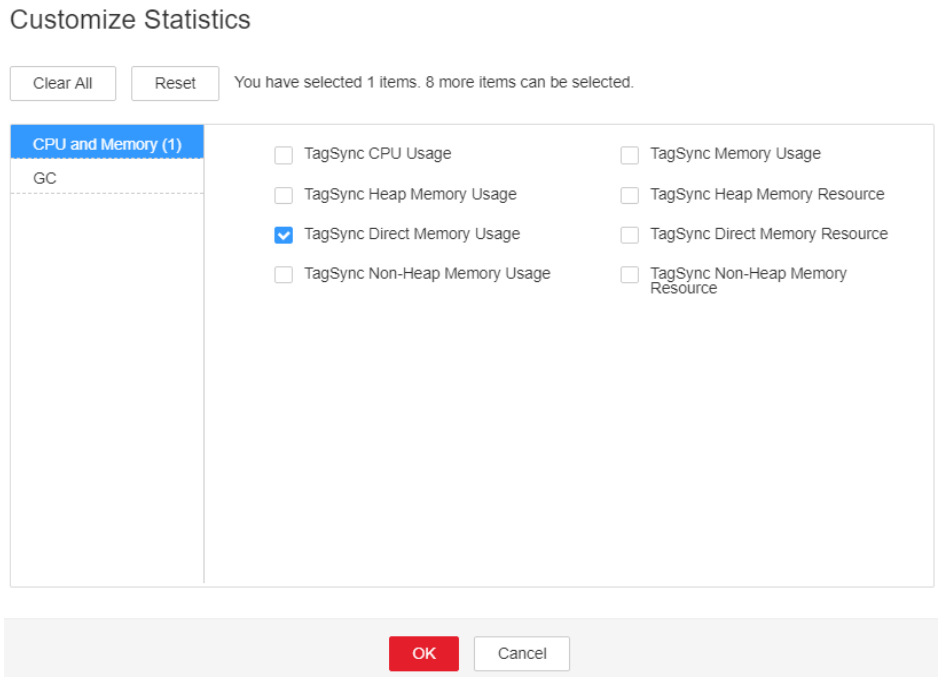

Step 3 Check whether the direct memory used by the TagSync reaches the threshold (80% of the maximum direct memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **TagSync** > **Instance Configuration**. Click **All Configurations** and choose **TagSync**

> **System**. Increase the value of **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the direct memory configured for TagSync cannot meet the direct memory required by the TagSync process. You are advised to check the direct memory usage of TagSync and change the value of **-XX:MaxDirectMemorySize** in **GC_OPTS** to the twice of the direct memory used by TagSync. You can change the value based on the actual service scenario. For details, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.299 ALM-45287 TagSync Non Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the non-heap memory usage of the TagSync service every 60 seconds. This alarm is generated when the non-heap memory usage of the TagSync instance exceeds the threshold (80% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45287	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Non-heap memory overflow may cause service breakdown.

Possible Causes

The non-heap memory of the TagSync process is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45287 TagSync Non Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **TagSync Non Heap Memory Usage**. Click **OK**.

Figure 10-123 TagSync non-heap memory usage

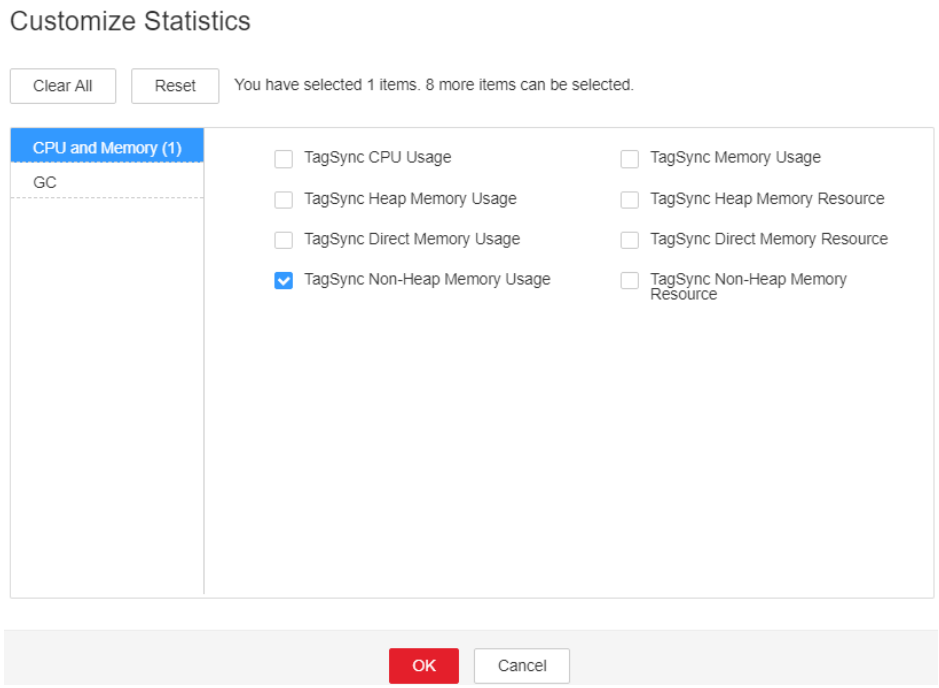

Step 3 Check whether the non-heap memory used by TagSync reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **TagSync** > **Instance Configuration**. Click **All Configurations** and choose **TagSync** > **System**. Set **-XX: MaxPermSize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the non-heap memory size configured for the TagSync instance cannot meet the non-heap memory required by the TagSync process. You are advised to change the **-XX:MaxPermSize** value of **GC_OPTS** to twice that of the current non-heap memory size or change the value based on the site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.300 ALM-45288 TagSync Garbage Collection (GC) Time Exceeds the Threshold

Alarm Description

The system checks the GC duration of the TagSync process every 60 seconds. This alarm is generated when the GC duration of the TagSync process exceeds the threshold (12 seconds by default) for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45288	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

TagSync responds slowly.

Possible Causes

The heap memory of the TagSync instance is overused or the heap memory is inappropriately allocated. As a result, GCs occur frequently.

Handling Procedure

Check the GC duration.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45288 TagSync Garbage Collection (GC) Time Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize** > **GC** > **TagSync GC Duration**. Click **OK**.

Figure 10-124 TagSync GC duration

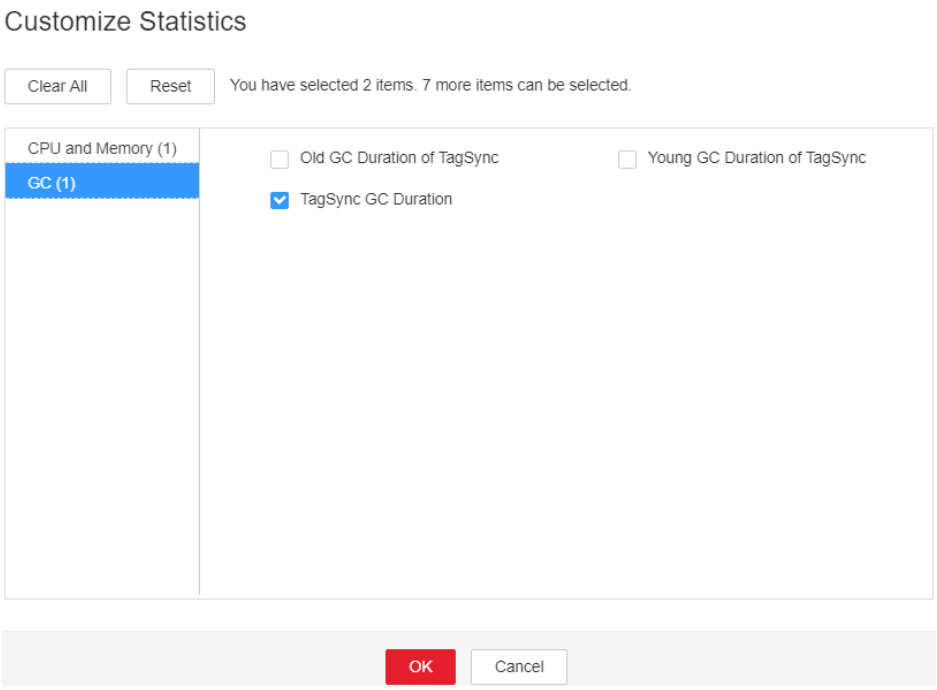

Step 3 Check whether the GC duration of the TagSync process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **TagSync** > **Instance Configuration**. Click **All Configurations** and choose **TagSync**

> **System**. Increase the value of **-Xmx** in the **GC_OPTS** parameter based on the site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for TagSync cannot meet the heap memory required by the TagSync process. You are advised to change the **-Xmx** value of **GC_OPTS** to twice that of the heap memory used by TagSync. You can change the value based on the actual service scenario. For details about how to check the TagSync heap memory usage, see **Step 2**.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.301 ALM-45289 PolicySync Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the PolicySync service every 60 seconds. This alarm is generated when the heap memory usage of the PolicySync instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times. This alarm is cleared when the heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45289	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Heap memory overflow may cause service breakdown.

Possible Causes

The heap memory of the PolicySync instance is overused or the heap memory is inappropriately allocated.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45289 PolicySync Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **PolicySync Heap Memory Usage**. Click **OK**.

Figure 10-125 PolicySync Heap Memory Usage

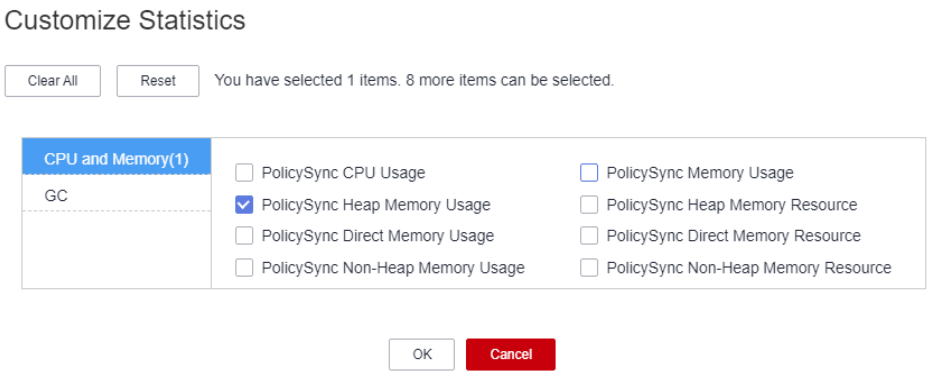

Step 3 Check whether the heap memory used by PolicySync reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **PolicySync**. Click **Instance Configuration** and then **All Configurations**, and choose **PolicySync** > **System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for PolicySync cannot meet the heap memory required by the PolicySync process. You are advised to change the value of **-Xmx** in **GC_OPTS** to twice that of the heap memory used by PolicySync. You can change the value based on the actual service scenario. Refer to **Step 2** to view the PolicySync heap memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.302 ALM-45290 PolicySync Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the PolicySync service every 60 seconds. This alarm is generated when the direct memory usage of the PolicySync instance exceeds the threshold (90% of the maximum memory) for five consecutive times. This alarm is cleared when the PolicySync direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45290	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service breakdown.

Possible Causes

The direct memory of the PolicySync process is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45290 PolicySync Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **PolicySync Direct Memory Usage**. Click **OK**.

Figure 10-126 PolicySync Direct Memory Usage

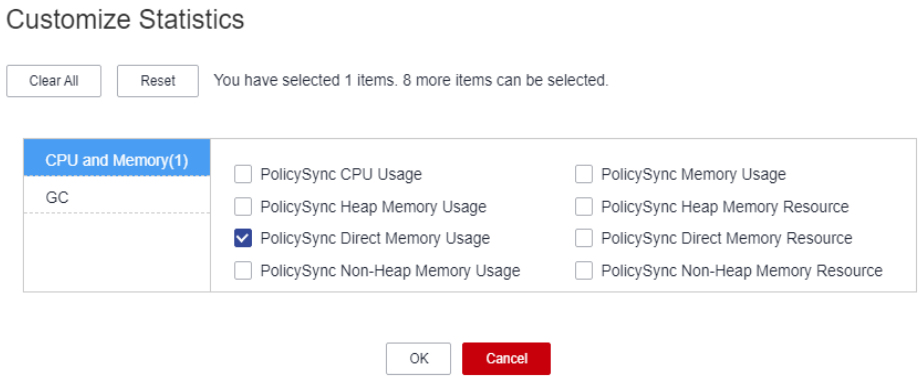

Step 3 Check whether the direct memory used by the PolicySync reaches the threshold (90% of the maximum direct memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **PolicySync**. Click **Instance Configuration** and then **All Configurations**, and choose **PolicySync** > **System**. Set **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the direct memory configured for PolicySync cannot meet the direct memory required by the PolicySync process. You are advised to check the direct memory usage of PolicySync and change the value of **-XX:MaxDirectMemorySize** in **GC_OPTS** to the twice of the direct memory used by PolicySync. You can change the value based on the actual service scenario. Refer to **Step 2** to view the TokenServer direct memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.303 ALM-45291 PolicySync Non-Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the non-heap memory usage of the PolicySync service every 60 seconds. This alarm is generated when the non-heap memory usage of the PolicySync instance exceeds the threshold (90% of the maximum memory) for five consecutive times. This alarm is cleared when the non-heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45291	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Non-heap memory overflow may cause service breakdown.

Possible Causes

The non-heap memory of the PolicySync instance is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45291 PolicySync Non-Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **PolicySync Non-Heap Memory Usage**. Click **OK**.

Figure 10-127 PolicySync Non-Heap Memory Usage

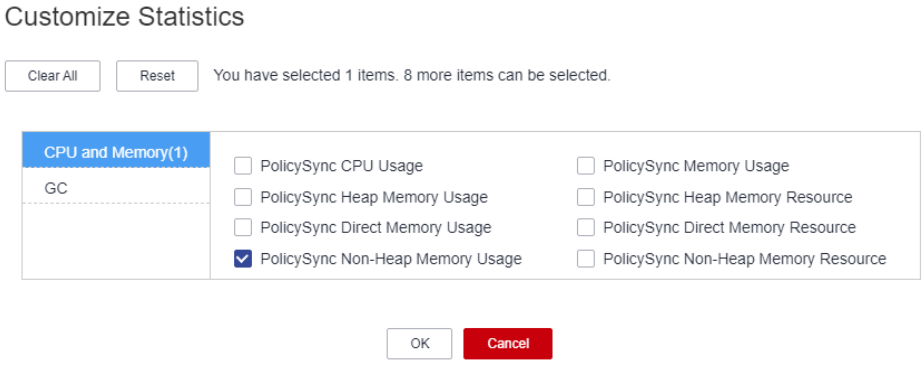

Step 3 Check whether the non-heap memory used by PolicySync reaches the threshold (90% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **PolicySync**. Click **Instance Configuration** and then **All Configurations**, and

choose **PolicySync** > **System**. Set **-XX: MaxPermSize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the non-heap memory size configured for the PolicySync instance cannot meet the non-heap memory required by the PolicySync process. You are advised to change the value of **-XX:MaxPermSize** in **GC_OPTS** to twice that of the current non-heap memory size or change the value based on site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.304 ALM-45292 PolicySync GC Duration Exceeds the Threshold

Alarm Description

The system checks the GC duration of the PolicySync process every 60 seconds. This alarm is generated when the GC duration of the PolicySync process exceeds the threshold for five consecutive times. This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45292	Critical (default threshold: 20000ms) Major (default threshold: 12000ms)	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

PolicySync responds slowly.

Possible Causes

The heap memory of the PolicySync process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45292 PolicySync GC Duration Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance**. Select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize** > **GC** > **PolicySync GC Duration**. Click **OK**.

Figure 10-128 PolicySync GC Duration

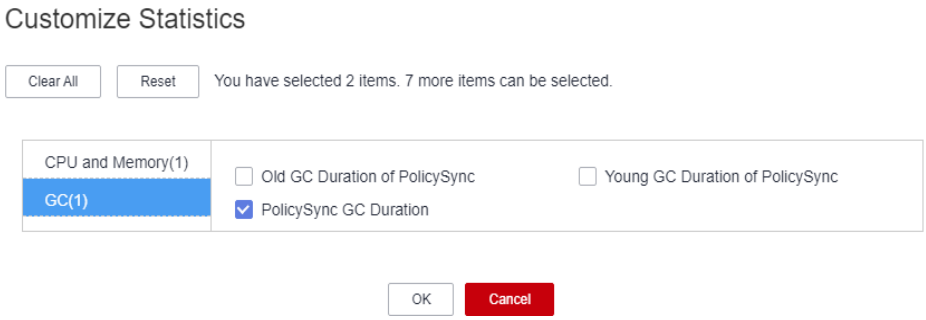

Step 3 Check whether the GC duration of the PolicySync process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Ranger** > **Instance** > **PolicySync**. Click **Instance Configuration** and then **All Configurations**, and choose **PolicySync** > **System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for PolicySync cannot meet the heap memory required by the PolicySync process. You are advised to change the value of **-Xmx** in **GC_OPTS** to twice that of the heap memory used by PolicySync. You can change the value based on the actual service scenario. Refer to **Step 2** to view the PolicySync heap memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Ranger** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.305 ALM-45325 Presto Service Unavailable

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

The system checks the Presto service status every 60 seconds. This alarm is generated when the Presto service is unavailable. This alarm is cleared when the Presto service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45325	Critical	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Presto cannot execute SQL statements.

Possible Causes

- The Presto coordinator or worker process is faulty.
- The network communication between Presto coordinator and worker instances is interrupted.

Handling Procedure

Check the status of the coordinator and worker processes.

Step 1 Log in to FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click the **Instance** tab. In the Presto instance list, check whether the running status of all coordinator or worker instances is **Unknown**.

- If yes, go to 2.
- If no, go to 4.

Step 2 Above the Presto instance list, click **More** and select **Restart Service** to restart the coordinator and worker processes.

Step 3 In the alarm list, check whether **ALM-45325 Presto Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to 4.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 6 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

10.306 ALM-45326 Number of Presto Coordinator Threads Exceeds the Threshold

Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45326	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

None

Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

Handling Procedure

Check the number of concurrent tasks.

Step 1 Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to **2**.
- If no, go to **Step 4**.

Adjust the alarm threshold for the number of threads.

Step 2 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Upgrade coordinator specifications or add worker node groups.

Step 4 Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.307 ALM-45327 Presto Coordinator Process GC Time Exceeds the Threshold

Alarm Description

The system collects the garbage collection (GC) time of the Presto coordinator process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the coordinator process is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45327	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC time of the coordinator process is too long, the coordinator process performance will be adversely affected, and the coordinator process will even become unavailable.

Possible Causes

The heap memory of the coordinator process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click the **Instance** tab and click the **Coordinator** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to **3**.
- If no, go to **6**.

Step 3 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Coordinator** > **JVM**, increase the value of **Xmx** in the **JAVA_OPTS** parameter, and click **Save**.

Step 4 Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, select **Presto** for the target cluster, and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact the O&M engineers and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.308 ALM-45328 Presto Worker Process GC Time Exceeds the Threshold

Alarm Description

The system collects the garbage collection (GC) time of the Presto worker process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time is less than or equal to the alarm threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45328	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC time of the worker process is too long, the worker process performance will be adversely affected, and the worker process will even become unavailable.

Possible Causes

The heap memory of the worker process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to **3**.
- If no, go to **6**.

Step 3 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker** > **JVM**, increase the value of **Xmx** in the **JAVA_OPTS** parameter, and click **Save**.

Step 4 Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, select **Presto** for the target cluster, and click **OK**.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.309 ALM-45329 Presto Coordinator Resource Group Queuing Tasks Exceed the Threshold

Alarm Description

The system queries the number of queuing tasks in a resource group through the JMX interface. This alarm is generated when the system detects that the number of queuing tasks in a resource group exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45329	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A large number of tasks may be in the queuing state and cannot be processed as expected. When the number of queuing tasks in a resource group exceeds the threshold (**maxQueued**), new tasks cannot be executed.

Possible Causes

The resource group configuration is improper or too many tasks in the resource group are submitted.

Handling Procedure

Step 1 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Coordinator** > **Customization**, and change the value of **resourceGroupAlarm** in the **resource-groups** parameter to change the threshold for each resource group.

Step 2 Collect fault information.

1. Log in to the cluster node based on the host name in the fault information and query the number of queuing tasks on the Presto client based on **Resource Group** in the additional information.
2. Log in to the cluster node based on the host name in the fault information, view the **/var/log/Bigdata/nodeagent/monitorlog/monitor.log** file, and search for resource group information to view the monitoring collection information of the resource group.
3. Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.310 ALM-45330 Number of Presto Worker Threads Exceeds the Threshold

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45330	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

None

Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

Handling Procedure

Check the number of concurrent tasks.

Step 1 Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to **2**.
- If no, go to **4**.

Adjust the alarm threshold for the number of threads.

Step 2 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Upgrade coordinator specifications or add worker node groups.

Step 4 Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.311 ALM-45331 Number of Presto Worker1 Threads Exceeds the Threshold

Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45331	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

None

Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

Handling Procedure

Check the number of concurrent tasks.

Step 1 Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to **2**.
- If no, go to **4**.

Adjust the alarm threshold for the number of threads.

Step 2 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Upgrade coordinator specifications or add worker node groups.

Step 4 Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.312 ALM-45332 Number of Presto Worker2 Threads Exceeds the Threshold

Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45332	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.

Impact on the System

None

Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

Handling Procedure

Check the number of concurrent tasks.

Step 1 Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to **2**.
- If no, go to **4**.

Adjust the alarm threshold for the number of threads.

Step 2 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Upgrade coordinator specifications or add worker node groups.

Step 4 Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.313 ALM-45333 Number of Presto Worker3 Threads Exceeds the Threshold

Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45333	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

None

Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

Handling Procedure

Check the number of concurrent tasks.

Step 1 Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to **2**.
- If no, go to **4**.

Adjust the alarm threshold for the number of threads.

Step 2 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Upgrade coordinator specifications or add worker node groups.

Step 4 Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.314 ALM-45334 Number of Presto Worker4 Threads Exceeds the Threshold

Alarm Description

The system checks the number of threads used by Presto coordinator and worker instances. The default threshold is 1024. This alarm is generated when the number of Presto coordinator or worker threads exceeds the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45334	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

None

Possible Causes

- Too many threads are used by Presto instances.
- Too many Presto tasks are concurrently executed.

Handling Procedure

Step 1 Check whether the CPU load of the current cluster is normal and the number of concurrent SQL statements meets the expectation.

- If yes, go to **2**.
- If no, go to **Step 4**.

Adjust the alarm threshold for the number of threads.

Step 2 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Thresholds**. On the **Thresholds** page, choose **Presto** and click **Number of Threads (Coordinator)** or **Number of Threads (Worker)**. Then, locate the row that contains the **default** rule and click **Modify** in the **Operation** column to increase the threshold, for example, increase the threshold by 20%.

Step 3 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Upgrade coordinator specifications or add worker node groups.

Step 4 Check the number of coordinator and worker threads. If the number of coordinator threads is too large, upgrade coordinator node specifications to increase the number of CPU cores. If the number of worker threads is too large, add worker node groups.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.315 ALM-45335 Presto Worker1 Process GC Time Exceeds the Threshold

Alarm Description

The system collects the garbage collection (GC) time of the Presto worker1 process every 30 seconds. This alarm is generated when the GC time exceeds the

threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker1 process is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45335	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC time of the worker1 process is too long, the worker1 process performance will be adversely affected, and the worker1 process will even become unavailable.

Possible Causes

The heap memory of the worker1 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker1** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to **3**.
- If no, go to **6**.

Step 3 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker1** > **JVM**, increase the value of **Xmx** in the **JAVA_OPTS** parameter, and click **Save**.

Step 4 Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.316 ALM-45336 Presto Worker2 Process GC Time Exceeds the Threshold

Alarm Description

The system collects the garbage collection (GC) time of the Presto worker2 process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker2 process is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45336	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC time of the worker2 process is too long, the worker2 process performance will be adversely affected, and the worker2 process will even become unavailable.

Possible Causes

The heap memory of the worker2 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker2** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to **3**.
- If no, go to **6**.

Step 3 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker2** > **JVM**, increase the value of **Xmx** in the **JAVA_OPTS** parameter, and click **Save**.

Step 4 Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.317 ALM-45337 Presto Worker3 Process GC Time Exceeds the Threshold

Alarm Description

The system collects the garbage collection (GC) time of the Presto worker3 process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker3 process is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45337	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC time of the worker3 process is too long, the worker3 process performance will be adversely affected, and the worker3 process will even become unavailable.

Possible Causes

The heap memory of the worker3 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker3** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to **3**.
- If no, go to **6**.

Step 3 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker3** > **JVM**, increase the value of **Xmx** in the **JAVA_OPTS** parameter, and click **Save**.

Step 4 Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.318 ALM-45338 Presto Worker4 Process GC Time Exceeds the Threshold

Alarm Description

The system collects the garbage collection (GC) time of the Presto worker4 process every 30 seconds. This alarm is generated when the GC time exceeds the threshold (exceeds 5 seconds for three consecutive times). This alarm is cleared when the GC time of the worker4 process is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45338	Minor	Yes

Alarm Parameters

Parameter	Description
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

If the GC time of the worker4 process is too long, the worker4 process performance will be adversely affected, and the worker4 process will even become unavailable.

Possible Causes

The heap memory of the worker4 process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. Locate the row that contains this alarm and view the IP address and role name of the instance in **Location**.

Step 2 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click the **Instance** tab and click the **Worker4** instance for which the alarm is generated to access its dashboard. In the upper right corner of the **Dashboard** page, click the drop-down list icon and select **Customize**. On the **Customize Statistics** page, choose **Cluster Status**, select **Presto Garbage Collection Time**, and click **OK** to check whether the GC time is greater than 5 seconds.

- If yes, go to **3**.
- If no, go to **6**.

Step 3 On FusionInsight Manager, choose **Cluster**, click the name of the desired cluster, and choose **Services** > **Presto**. On the page that is displayed, click **Configurations** and **All Configurations**, choose **Worker4** > **JVM**, increase the value of **Xmx** in the **JAVA_OPTS** parameter, and click **Save**.

Step 4 Return to the **Dashboard** page, click **More**, and select **Restart Service** to restart the service.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Presto** for the target cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.319 ALM-45425 ClickHouse Service Unavailable

Alarm Description

The alarm module checks the ClickHouse instance status every 60 seconds. This alarm is generated when the alarm module detects that all ClickHouse instances are abnormal.

This alarm is cleared when the system detects that any ClickHouse instance is restored and the alarm is cleared.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45425	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The ClickHouse service is abnormal. You cannot use FusionInsight Manager to perform cluster operations on the ClickHouse service. The ClickHouse service function is unavailable.

Possible Causes

The configuration information in the **metrika.xml** file in the component configuration directory of the faulty ClickHouse instance node is inconsistent with that of the corresponding ClickHouse instance in the ZooKeeper.

Handling Procedure

Check whether the configuration in metrika.xml of the ClickHouse instance is correct.

Step 1 Log in to FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse** > **Instance**, and locate the abnormal ClickHouse instance based on the alarm information.

- If yes, go to **Step 2**.
- If no, go to **Step 9**.

Step 2 Log in to the host where the ClickHouse service is abnormal and ping the IP address of another normal ClickHouse instance node to check whether the network connection is normal.

- If yes, go to **Step 3**.
- If no, contact the network administrator to repair the network.

Step 3 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- For a cluster with Kerberos authentication enabled (security mode):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**

- For a cluster with Kerberos authentication disabled (normal mode):

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password --port** 9000

Run the following command to query the value of **macros.id**:

select substitution from system.macros where macro='id';

Step 4 Log in to the host where the ZooKeeper client is located and log in to the ZooKeeper client.

Switch to the client installation directory.

Example: **cd /opt/client**

Run the following command to configure environment variables:

source bigdata_env

Run the following command to authenticate the user (skip this step in common mode):

kinit *Component service user*

Run the following command to log in to the client tool:

zkCli.sh -server *service IP address of the node where the ZooKeeper role instance locates*:*client port*

Step 5 Run the following command to check whether the ClickHouse cluster topology information can be obtained.

get /clickhouse/config/*value of* ***macros.id*** *in* **Step 3**/**metrika.xml**

- If yes, go to **Step 6**.
- If no, go to **Step 9**.

Step 6 Log in to the host where the ClickHouse instance is abnormal and go to the configuration directory of the ClickHouse instance.

cd ${BIGDATA_HOME}/**FusionInsight_ClickHouse**_*Version*/x_x_**ClickHouseServer/etc**

cat metrika.xml

Step 7 Check whether the cluster topology information on ZooKeeper obtained in **Step 5** is the same as that in the **metrika.xml** file in the component configuration directory in **Step 6**.

- If yes, check whether the alarm is cleared. If the alarm persists, go to **Step 9**.
- If no, go to **Step 8**.

Step 8 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**, click **More**, and select **Synchronize Configuration**. Then, check whether the service status is normal and whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect the fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 11 Choose the corresponding host form the host list.

Step 12 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.320 ALM-45426 ClickHouse Service Quantity Quota Usage in ZooKeeper Exceeds the Threshold

Alarm Description

The alarm module checks the quota usage of the ClickHouse service in the ZooKeeper every 60 seconds. This alarm is generated when the alarm module detects that the usage exceeds the threshold (90%).

This alarm is cleared when the system detects that the usage is lower than the threshold and the alarm is cleared.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45426	Major (default)	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

After the ZooKeeper quantity quota of the ClickHouse service exceeds the threshold, you cannot perform cluster operations on the ClickHouse service on FusionInsight Manager. As a result, the ClickHouse service cannot be used.

Possible Causes

- When table data is created, inserted, or deleted, the ClickHouse creates znodes on ZooKeeper nodes. As the service volume increases, the number of znodes may exceed the configured threshold.
- No quota limit is set for the metadata directory **/clickhouse** of ClickHouse in ZooKeeper.

Handling Procedure

Check the number of znodes created by ClickHouse on ZooKeeper.

Step 1 Log in to the host where the ZooKeeper client is located and log in to the ZooKeeper client.

Switch to the client installation directory.

Example: **cd /opt/client**

Run the following command to configure environment variables:

source bigdata_env

Run the following command to authenticate the user (skip this step in common mode):

kinit *Component service user*

Run the following command to log in to the client tool:

zkCli.sh -server *service IP address of the node where the ZooKeeper role instance locates*:*client port*

Step 2 Run the following command to check the quota used by the ClickHouse in the ZooKeeper and check whether the quota information is correctly set:

listquota /clickhouse

```
absolute path is /zookeeper/quota/clickhouse
Quota for path /clickhouse does not exist.
```

If the preceding information indicates that the quota configuration is incorrect, go to **Step 3**.

If no, go to **Step 5**.

Step 3 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **ZooKeeper**. On the displayed page, click **Configurations** and click **All Configurations**. On this sub-tab page, search for **quotas.auto.check.enable** to check whether its value is **true**.

If the value is not **true**, change the value to **true** and click **Save**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**, click **More**, and select **Synchronize Configuration**. After the synchronization is successful, go to **Step 1**.

Step 5 Run the following command and check whether the ratio of the **count** value of **Output stat** to the **count** value of **Output quota** in the command output is greater than **0.9**:

listquota /clickhouse

```
absolute path is /zookeeper/quota/clickhouse
Output quota for /clickhouse count=200000,bytes=1000000000
Output stat for /clickhouse count=2667,bytes=60063
```

In the preceding information, the **count** value of **Output stat** is **2667**, and the **count** value of **Output quota** is **200000**.

- If yes, go to **Step 6**.
- If no, check whether the alarm is cleared 5 minutes later. If the alarm persists, go to **Step 8**.

Step 6 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**. Click **Configurations** and then **All Configurations**. Search for **clickhouse.zookeeper.quota.node.count**, set it to a value twice the **count** of **Output stat** in **Step 5**. Do not use a value larger than **6000000**. Otherwise, there will be high risks. Exercise caution when setting this parameter.

Step 7 Restart the ClickHouse instance for which the alarm is generated, and check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, perform **Step 6** again, and check whether the alarm is cleared 5 minutes later. If the alarm persists, go to **Step 8**.

Collect the fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 10 Choose the corresponding host form the host list.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.321 ALM-45427 ClickHouse Service Capacity Quota Usage in ZooKeeper Exceeds the Threshold

Alarm Description

The alarm module checks the quota usage of the ClickHouse service in the ZooKeeper every 60 seconds. This alarm is generated when the alarm module detects that the usage exceeds the threshold (90%).

This alarm is cleared when the system detects that the usage is lower than the threshold and the alarm is cleared.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45427	Major (default)	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

After the ZooKeeper quantity quota of the ClickHouse service exceeds the threshold, you cannot perform cluster operations on the ClickHouse service on FusionInsight Manager. As a result, the ClickHouse service cannot be used.

Possible Causes

- When table data is created, inserted, or deleted, the ClickHouse creates znodes on ZooKeeper nodes. As the service volume increases, the capacity of znodes may exceed the configured threshold.
- No quota limit is set for the metadata directory **/clickhouse** of ClickHouse in ZooKeeper.

Handling Procedure

Check the znode capacity of the ClickHouse in the ZooKeeper.

Step 1 Log in to the host where the ZooKeeper client is located and log in to the ZooKeeper client.

Switch to the client installation directory.

Example: **cd /opt/client**

Run the following command to configure environment variables:

source bigdata_env

Run the following command to authenticate the user (skip this step in common mode):

kinit *Component service user*

Run the following command to log in to the client tool:

zkCli.sh -server *service IP address of the node where the ZooKeeper role instance locates*:*client port*

Step 2 Run the following command to check the quota used by the ClickHouse in the ZooKeeper and check whether the quota information is correctly set:

listquota /clickhouse

```
absolute path is /zookeeper/quota/clickhouse
Quota for path /clickhouse does not exist.
```

- If the preceding information indicates that the quota configuration is incorrect, go to **Step 3**.
- If not, go to **Step 5**.

Step 3 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **ZooKeeper**. On the displayed page, click **Configurations** and click **All Configurations**. On this sub-tab page, search for **quotas.auto.check.enable** to check whether its value is **true**.

If the value is not **true**, change the value to **true** and click **Save**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**, click **More**, and select **Synchronize Configuration**. After the synchronization is successful, go to **Step 1**.

Step 5 Run the following command and check whether the ratio of the **bytes** value of **Output stat** to the **bytes** value of **Output quota** in the command output is greater than **0.9**:

listquota /clickhouse

```
absolute path is /zookeeper/quota/clickhouse
Output quota for /clickhouse count=200000,bytes=1000000000
Output stat for /clickhouse count=2667,bytes=60063
```

In the preceding information, the **bytes** value of **Output stat** is **60063**, and the **bytes** value of **Output quota** is **1000000000**.

- If yes, go to **Step 6**.
- If no, check whether the alarm is cleared 5 minutes later. If the alarm persists, go to **Step 8**.

Step 6 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**. Click **Configurations** and then **All Configurations**. Search for **clickhouse.zookeeper.quota.size**, set it to a value twice the **bytes** of **Output stat** in **Step 5**. Do not use a value larger than **6000000**. Otherwise, there will be high risks. Exercise caution when setting this parameter.

Step 7 Restart the ClickHouse instance for which the alarm is generated, and check whether the alarm is cleared 5 minutes later.

- If yes, no further action is required.
- If no, perform **Step 6** again, and check whether the alarm is cleared 5 minutes later. If the alarm persists, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 10 Choose the corresponding host form the host list.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.322 ALM-45428 ClickHouse Disk I/O Exception

Alarm Description

This alarm is generated when the alarm module detects EIO or EROFS errors during ClickHouse read and write every 60 seconds.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45428	Major (default)	No

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

- ClickHouse fails to read and write data. The INSERT, SELECT, and CREATE operations on the local tables may be abnormal. Distributed tables are not affected.
- Services are affected, and I/Os fail.

Possible Causes

The disk is aged or has bad sectors.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45428 ClickHouse Disk I/O Exception**. Check the role name and the IP address of the host where the alarm is generated in **Location**.

Step 2 Use PuTTY to log in to the node for which the fault is generated as user **root**.

Step 3 Run the **df -h** command to check the mount directory and find the disk mounted to the faulty directory.

Step 4 Run the **smartctl -a /dev/sd***Faulty disk* command to check the disk. In the command, *Faulty disk* indicates the disk obtained in **Step 3**.

- If **SMART Health Status: OK** is displayed, as shown in the following figure, the disk is healthy. In this case, go to **Step 6**.

```
=== START OF READ SMART DATA SECTION ===
SMART Health Status: OK

Current Drive Temperature:     26 C
Drive Trip Temperature:        60 C

Manufactured in week 50 of year 2018
Specified cycle count over device lifetime:  10000
Accumulated start-stop cycles:  25
Specified load-unload count over device lifetime:  300000
Accumulated load-unload cycles:  356
Elements in grown defect list: 0
```


- If the number following **Elements in grown defect list** is not 0, as shown in the following figure, the disk may have bad sectors. If **SMART Health Status: FAILURE** is displayed, the disk is in the sub-health state. In this case, contact O&M personnel.

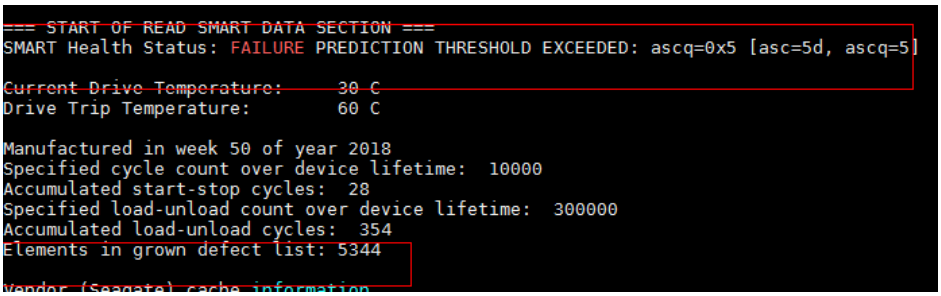

Step 5 After the fault is rectified, manually clear the alarm on FusionInsight Manager and check whether the alarm is generated again during the periodic check.

- If yes, go to **Step 6**.
- If no, no further action is required.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 8 Choose the corresponding host form the host list.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

If the alarm has no impact, manually clear the alarm.

Related Information

None

10.323 ALM-45429 Table Metadata Synchronization Failed on the Added ClickHouse Node

NOTE

This section applies only to MRS 3.1.2 or later.

Alarm Description

This alarm is generated when the local table corresponding to the distributed table fails to be created during ClickHouse capacity expansion.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45429	Major	No

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The distributed table fails to be queried.

Possible Causes

A node is stopped or faulty during capacity expansion.

Handling Procedure

Step 1 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse** > **Instance**.

Step 2 Check whether an instance is stopped, decommissioned, or faulty.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Start the instance or rectify the instance fault until all instances are running properly.

Step 4 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, locate this alarm and the faulty host based on the location information.

Step 5 Log in to the faulty host as user **omm**.

Step 6 Run the following commands to initialize environment variables:

source *Cluster installation directory*/**FusionInsight_ClickHouse_*/*_*_ClickHouseServer/etc/ENV_VARS**

source *Cluster installation directory*/**FusionInsight_ClickHouse_*/*_*_ClickHouseServer/etc/clickhouse-env.sh**

export CLICKHOUSE_CONF_DIR=${CLICKHOUSE_CONF_DIR}

Step 7 Run the following command to run the metadata synchronization tool to synchronize metadata from the existing node to the faulty node:

sh *Cluster installation directory*/**FusionInsight_ClickHouse_*/install/FusionInsight-ClickHouse-*/clickhouse/sbin/clickhouse-create-meta.sh true**

Step 8 Run the following command to view the log information and check whether the metadata has been synchronized:

vim /var/log/Bigdata/clickhouse/clickhouseServer/start.log

- If the synchronization is complete, go to **Step 9**.
- If the synchronization fails, go to **Step 10**.

Step 9 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the **Alarm ID** column, locate the corresponding alarm and click **Clear** in the **Operation** column. In the displayed dialog box, click **OK** to manually clear the alarm.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, select **ClickHouse** for the target cluster, and click **OK**.

Step 12 Choose the corresponding host form the host list.

Step 13 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm needs to be manually cleared after the fault is rectified.

Related Information

None

10.324 ALM-45430 Permission Metadata Synchronization Failed on the Added ClickHouse Node

NOTE

This section applies only to MRS 3.1.2 or later.

Alarm Description

This alarm is generated when user and permission information fails to be synchronized during ClickHouse capacity expansion.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45430	Major	No

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The created user does not have operation permissions on the node.

Possible Causes

A node is stopped or faulty during capacity expansion.

Handling Procedure

Step 1 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse** > **Instance**.

Step 2 Check whether an instance is stopped, decommissioned, or faulty.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Start the instance or rectify the instance fault until all instances are running properly.

Step 4 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, locate this alarm and the faulty host based on the location information.

Step 5 Log in to the faulty host as user **omm**.

Step 6 Run the following commands to initialize environment variables:

source ${BIGDATA_HOME}/FusionInsight_ClickHouse_*/*_*_ClickHouseServer/etc/ENV_VARS

source ${BIGDATA_HOME}/FusionInsight_ClickHouse_*/*_*_ClickHouseServer/etc/clickhouse-env.sh

export CLICKHOUSE_CONF_DIR=${CLICKHOUSE_CONF_DIR}

Step 7 Run the following command to run the metadata synchronization tool to synchronize metadata from the existing node to the faulty node:

sh ${BIGDATA_HOME}/FusionInsight_ClickHouse_*/install/FusionInsight-ClickHouse-*/clickhouse/sbin/clickhouse-create-meta.sh true

Step 8 Run the following command to view the log information and check whether the metadata has been synchronized:

vim /var/log/Bigdata/clickhouse/clickhouseServer/start.log

If the synchronization is complete, go to **Step 9**.

If the synchronization fails, go to **Step 10**.

Step 9 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the **Alarm ID** column, locate the corresponding alarm and click **Clear** in the **Operation** column. In the displayed dialog box, click **OK** to manually clear the alarm.

Collect the fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, select **ClickHouse** for the target cluster, and click **OK**.

Step 12 Choose the corresponding host form the host list.

Step 13 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm needs to be manually cleared after the fault is rectified.

Related Information

None

10.325 ALM-45431 Improper ClickHouse Instance Distribution for Topology Allocation

Alarm Description

The ClickHouseServer instance distribution does not meet the topology allocation requirements.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45431	Critical	No

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Some ClickHouseServer instances are unavailable.

Possible Causes

During installation or capacity expansion, the number of instances or allocation mode does not meet the topology requirements.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**, locate the row that contains the alarm, and analyze the cause based on **Location** and **Additional Information**.

Step 2 Handle the alarm based on the additional alarm information and handling method in the following table.

Additional Information	Remarks	Handling Method
n ClickHouseServer instances should be added to other AZ.	This alarm is generated when a single cluster is deployed in cross-AZ DR mode. The ClickHouseServer instance deployment does not meet the cross-AZ DR topology allocation requirements. As a result, some instances cannot work properly.	1. On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**, click the **Instance** tab, locate the row that contains the alarm, view the host name in **Location**, and find the AZ for which the alarm is generated in the AZ column based on the host name. 2. On the **Instance** page, click **Add Instance** to add *n* ClickHouseServer instances to other AZs except the AZ where the alarm is generated.
n ClickHouseServer instances should be added.	This alarm is generated when a non-single-cluster is deployed in the default deployment mode of cross-AZ DR. The number of ClickHouseServer instances in the cluster is less than an even number. As a result, some instances cannot work properly.	1. Determine the number (*n*) of ClickHouseServer instances to be added based on the alarm information. 2. On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**, click the **Instance** tab, and add *n* ClickHouseServer instances to the cluster.

----End

Alarm Clearance

This alarm needs to be manually cleared after the fault is rectified.

Related Information

None

10.326 ALM-45432 ClickHouse User Synchronization Process Fails

Alarm Description

The system checks the status of the ClickHouse user role synchronization process every 5 minutes. This alarm is generated when the system detects that the ClickHouse user role synchronization process is faulty or the user role synchronization fails.

This alarm is automatically cleared when the ClickHouse user role synchronization process or function becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45432	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Some ClickHouseServer instances are unavailable.

Possible Causes

- The ClickHouse user role synchronization process is not started properly or exits abnormally.
- The ClickHouse user role synchronization process fails to synchronize user role information because the LdapServer service is faulty.

Handling Procedure

Check whether the ClickHouse user role synchronization process is normal.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, search for **ALM-45432 ClickHouse User Synchronization Process Fails**.

Step 2 Check the host name and additional information in the alarm details.

- If the additional information is "Process clickhouse-ugsync is not exit.", go to **Step 3**.
- If the additional information is "Process clickhouse-ugsync sync user failed.", go to **Step 6**.

Step 3 Log in to the faulty host as user **omm** and run the following command to check whether the ClickHouse user role synchronization process is normal:

ps -ef | grep 'clickhouse-ugsync'

Abnormal result of the synchronization process:

```
[omm@server-2110081635-0001 ~]$ ps -ef | grep 'clickhouse-ugsync'
omm      20104 13146  0 15:57 pts/7    00:00:00 grep --color=auto clickhouse-ugsync
```

- If yes, the alarm is automatically cleared. If the alarm is cleared, no further action is required. If the alarm persists, go to **Step 8**.
- If no, go to **Step 4**.

Step 4 Log in to the faulty host as user **omm** and run the following command to check whether the crontab daemon task is correctly configured:

crontab -l

Normal setting of the crontab daemon task:

```
*/5 * * * * bash /xxxxx/clickhouse_ugsync_check.sh >/dev/null 2>&1
```

- If yes, check whether the alarm is cleared 5 minutes later. If the alarm is cleared, no further action is required. If the alarm persists, go to **Step 8**.
- If no, go to **Step 5**.

Step 5 Log in to FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse**. On the page that is displayed, click the **Instance** tab. On this tab page, find the abnormal ClickHouseServer instance based on the fault information, and restart it. Wait for 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the LdapServer service is normal.

Step 6 Log in to FusionInsight Manager, choose **Cluster** > **Services**, and check whether **Running Status** of LdapServer is **Normal**.

- If yes, go to **Step 8**.
- If no, go to **Step 7**.

Step 7 Handle the LdapServer service unavailable alarm according to ALM-25000 LdapServer Service Unavailable.

After **Running Status** of LdapServer becomes **Normal**, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 9 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **ClickHouseServer** for the target cluster.

Step 10 Expand the **Hosts** list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.327 ALM-45433 ClickHouse AZ Topology Exception

Alarm Description

If the cross-AZ HA function is enabled for a cluster where ClickHouse has been deployed, the ClickHouse topology remains unchanged. This alarm is generated when the cross-AZ HA does not take effect if backup nodes of the same shard are in the same AZ.

This alarm is automatically cleared when the system detects that all shards meet the cross-AZ HA deployment requirements.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45433	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The current deployment of the ClickHouse service does not support cross-AZ HA.

Possible Causes

After cross-AZ HA is enabled, all backup nodes of a shard are in the same AZ.

Handling Procedure

Modify the AZ of backup nodes.

Step 1 Log in to the node where the client is installed as the client installation user. Run the following command to switch to the client installation directory:

cd *{Client installation path}*

Step 2 Run the following command to configure environment variables:

source bigdata_env

Step 3 Run the following command to authenticate the user (skip this step in normal mode):

kinit *Component service user*

Step 4 Run the following command to log in to the client tool:

zkCli.sh -server*Service IP address of the node where the ZooKeeper instance resides***:***Client port*

Step 5 Run the following command to view the current topology:

get /clickhouse/topo

NOTE

If the ClickHouse is installed with multiple services, run the **get /clickhouse{-n}/topo** command. For example, if the ClickHouse-1 is installed, run the **get /clickhouse-1/topo** command.

```
[zk: 192.168.20.36:24002(CONNECTED) 0] get /clickhouse/topo

<topo>
 <mcluster>
  <shard id="14" index="1">
   <server id="15">
    <replica>1</replica>
    <az>AZ1</az>
    <host>192-168-20-205</host>
    <port>21427</port>
   </server>
   <server id="16">
    <replica>2</replica>
    <az>AZ1</az>
    <host>192-168-20-2205</host>
    <port>21427</port>
   </server>
  </shard>
 </mcluster>
</topo>
```

Step 6 Select a host from the desired shard and deploy the host in another AZ.

Step 7 Log in to FusionInsight Manager, click **Host**, select the host you have deployed in **Step 6** and choose **More** > **Reinstall** to reinstall the host.

Step 8 Choose **Cluster** > **Cross-AZ HA**, click **Configure AZ and Policy** and change the AZ information of the reinstalled host to the AZ planned in the **Step 6**.

Step 9 Wait for five minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 11 Expand the drop-down list next to the **Service** field. In the **Services** dialog box that is displayed, select **ClickHouseServer** for the target cluster.

Step 12 Expand the **Hosts** list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 13 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.328 ALM-45434 A Single Replica Exists in the ClickHouse Data Table

Description

This alarm is generated when a single replica is detected in a custom logical cluster after the custom logical cluster is enabled for ClickHouse.

This alarm is automatically cleared when the system detects that the custom logical cluster uses multiple replicas.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45434	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If a hardware fault occurs, data cannot be restored.

Possible Causes

The **metrika.xml** file in the ClickHouse configuration directory contains single-replica configuration.

Procedure

Check whether the configuration in metrika.xml of the ClickHouse instance is correct.

Step 1 In the alarm list on FusionInsight Manager, locate the row that contains the alarm, and click ⌄ to view the name of the host for which the alarm is generated. On the **Hosts** page, view the host IP address based on the host name.

Step 2 Log in to the host where the ClickHouse instance is abnormal, go to the configuration directory of the ClickHouse instance, and run the following commands:

cd ${BIGDATA_HOME}**/FusionInsight_ClickHouse_***Version*/x_x_**ClickHouseServer/etc**

cat metrika.xml

Step 3 View the number of shards in each custom logical cluster and check that a single replica exists. Then, go to **Step 4**.

NOTE

If a shard contains only one node, a single replica exists in a logical cluster, as shown in the following:

```
<shard>
 <internal_replication>true</internal_replication>
 <replica>
  <host>host-name1</host>
  <port>port</port>
  <user>clickhouse</user>
  <password/>
 </replica>
</shard>
```

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 6 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearing

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.329 ALM-45435 Inconsistent Metadata of ClickHouse Tables

Alarm Description

This alarm is generated when the metadata in a distributed table or in the local table of the distributed table has been inconsistent for 180 min.

This alarm is automatically cleared when the metadata in the distributed table or in the local table of the distributed table becomes consistent.

Metadata consistency includes:

- Consistent quantity, name, sequence, and type of each column in the table
- Consistent partition keys
- Consistent sorting keys
- Consistent primary keys
- Consistent sampling keys

NOTE

If this alarm exists, table metadata is inconsistent in the ClickHouse cluster to which the current node belongs. The inconsistency may be caused by multiple reasons, not limited to those mentioned in additional information.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45435	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Table	Specifies the database name and table name for which the alarm is generated.

Impact on the System

Subsequent operations such as INSERT and ALTER on the table may fail.

Possible Causes

Table metadata modification fails or is not executed on one or more ClickHouseServer nodes.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- For a cluster with Kerberos authentication enabled (security mode):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**

- For a cluster with Kerberos authentication disabled (normal mode):

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password --port** 9000

Step 3 Check whether any task is being executed for the table to which the alarm is generated.

Run the following command to check whether any SQL task is being executed:

select * from system.processes where current_database='*Database name***' and query like '%***Table name***%'**

Run the following command to check whether a mutation task is being executed:

select * from system.mutations where database=*'Database name'* **and table=***'Table name'***;**

- If the query result is empty, go to **Step 4**.
- If the query result contains error information, rectify the fault accordingly. If the fault cannot be rectified based on the error information, go to **Step 6**.
- If the query result contains information about an on-going task with no error, the SQL/mutation task is being executed.

  Wait for 5 minutes. If the alarm is cleared, no further action is required. If the alarm persists, go to **Step 4**.

Step 4 Modify the table structure, delete a table, or add a table based on service requirements until the table metadata of all nodes in the cluster is consistent. After 5 minutes, check whether this alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 5**.

Step 5 If the table has been deleted, manually clear the alarm and check whether the alarm is reported again.

- If yes, go to **Step 6**.
- If no, no further action is required.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.330 ALM-45436 Skew ClickHouse Table Data

Alarm Description

This alarm is generated when data skew occurs in the local table of a distributed table between ClickHouse nodes. This alarm is automatically cleared when data becomes balanced.

Data skew check method:

- If **min_table_check_data_bytes** is set to **0**, data skew check is disabled.
- If **min_table_check_data_bytes** is greater than **0**, data skew check is enabled.

After data skew check is enabled, if the data volume in a table is less than the **min_table_check_data_bytes** value, no alarm will be reported due to data skew. When the data volume is greater than the **min_table_check_data_bytes** value and the data volume difference between the same table on different nodes is greater than the percentage specified in **min_table_data_varies_rate**, data skew occurs and this alarm is reported.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45436	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Table	Specifies the database name and table name for which the alarm is generated.

Impact on the System

SQL execution efficiency may be lowered.

Possible Causes

The data write policy is improper, causing unbalanced data among nodes.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- Security mode (with Kerberos enabled):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**

- Normal mode (with Kerberos disabled):

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password --port** 9000

Step 3 View data distribution.

select FQDN(), database, table, sum(data_compressed_bytes) from clusterAllReplicas(*Name of the logical cluster***, system.parts) where database='***Database name***' and table='***Table name***' and active=1 group by (FQDN(), database, table);**

Step 4 Balance data with a few clicks or migrate data based on service requirements.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.331 ALM-45437 Excessive Parts in the ClickHouse Table

Alarm Description

This alarm is generated when the number of parts exceeds the threshold specified by **part_num_threshold**.

This alarm is automatically cleared when the number of parts is less than the **part_num_threshold** value.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45437	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
Table	Specifies the database name and table name for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Service errors may occur.

Possible Causes

The data distribution in the ClickHouse table is improper, or the background merge task is executed slowly.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- Security mode (with Kerberos enabled):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**

- Normal mode (with Kerberos disabled):

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password --port** 9000

Step 3 Run the following command to manually merge parts:

optimize table *Database name*.*Table name* **final;**

Step 4 Check whether the number of parts has decreased.

select FQDN(), database, table, count(1) from clusterAllReplicas(default_cluster, system.parts) where database='*Database name***' and table='***Table name***' and active=1 group by (FQDN(), database, table);**

1. If the number of parts is less than the threshold, wait for 5 minutes and check whether the alarm is cleared.
   - If yes, no further action is required.
   - If no, go to **Step 5**.
2. If the number of parts does not decrease, check whether the partition key of the table is set properly. If the number of partitions is too large, rectify the service logic.
3. If the command output is empty, the table does not exist. This alarm is a historical alarm and can be ignored. Manually clear it.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 7 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.332 ALM-45438 ClickHouse Disk Usage Exceeds 80%

Alarm Description

The system checks the disk capacity of the ClickHouseServer node every 1 minute. This alarm is generated when the usage of the disk where the ClickHouse data directory or metadata directory resides exceeds 80%.

This alarm is automatically cleared when the usage of the disk where the ClickHouse data directory or metadata directory is located is lower than 80%.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45438	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DiskPath	Specifies the path of the disk for which the alarm is generated.

Impact on the System

The ClickHouse write operation may fail.

Possible Causes

The disk capacity of the ClickHouseServer node is too small.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Expand the disk capacity of the node for which the alarm is generated.

Step 3 Go to **Step 4** if the expansion fails or the alarm persists after the expansion.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 6 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.333 ALM-45439 ClickHouse Node Enters the Read-Only Mode

Alarm Description

The system checks the disk capacity of the ClickHouseServer node every 1 minute. This alarm is generated when the system detects that the disk capacity exceeds 90% and the ClickHouseServer node enters the read-only mode.

This alarm is automatically cleared when the system detects that the disk capacity is lower than 90% and the ClickHouseServer node exits the read-only mode.

NOTE

If the ClickHouseServer node is in read-only mode and you need to log in to the client to clear data, you can manually exit the read-only mode using the following method:

Log in to FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse** > **Configurations** > **All Configurations**, search for **profiles.default.readonly**, and change its value to **0**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45439	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
DiskPath	Specifies the path of the disk for which the alarm is generated.

Impact on the System

After the ClickHouseServer node enters the read-only mode, all write, modification, and deletion operations fail.

Possible Causes

The disk usage of the ClickHouse node exceeds 90%.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Expand the disk capacity of the node for which the alarm is generated.

Step 3 Go to **Step 4** if the expansion fails or the alarm persists after the expansion.

NOTE

After the capacity expansion, this alarm can be automatically cleared only when **profiles.default.readonly** is **auto**. If its value has been manually changed, change it back to **auto**. If **profiles.default.readonly** needs to be set to **0** or **1** based on service requirements, manually clear this alarm.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 6 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.334 ALM-45440 Inconsistency Between ClickHouse Replicas

Alarm Description

When the number of ClickHouse replicas is greater than 1, the system periodically checks the replicated table. This alarm is generated if replicated table data is not synchronized. This alarm is cleared when data in all replicated tables between replicas becomes synchronized.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45440	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Table	Specifies the table name for which the alarm is generated.

Impact on the System

The data reliability of the ClickHouse replicated table is affected, causing data differences and affecting the query result of the distributed table.

Possible Causes

- The ClickHouse service is overloaded.
- The connection between the ClickHouse and ZooKeeper is abnormal.

Handling Procedure

Check whether the ClickHouse service load is heavy.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the database name, table name, role name and IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- For a cluster with Kerberos authentication enabled (security mode):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**
- For a cluster with Kerberos authentication disabled (normal mode):

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password --port** 9000

Step 3 Run the following statement to check whether data is frequently written to the system table. If yes, wait until the service execution is complete and check whether the alarm is cleared.

SELECT query_id, user, FQDN(), elapsed, query FROM system.processes ORDER BY query_id;

- If yes, no further action is required.
- If no, go to **Step 4**.

Step 4 Check whether a large amount of data is written. If yes, wait until the task is complete and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Run the following statement to check whether replicas are synchronized:

select table,absolute_delay, queue_size, inserts_in_queue, merges_in_queue from system.replicas where absolute_delay > 0 order by absolute_delay desc limit 10;

- If yes, go to **Step 6**.

- If no, go to **Step 9**.

Step 6 If **inserts_in_queue** contains a large amount of content to be inserted, run the following SQL statement to query the replica synchronization queue and locate the error cause:

SELECT database,table,type,any(last_exception),any(postpone_reason),min(create_time),max(last_attempt_time),max(last_postpone_time),max(num_postponed) AS max_postponed,max(num_tries) AS max_tries,min(num_tries) AS min_tries,countIf(last_exception != '') AS count_err,countIf(num_postponed > 0) AS count_postponed,countIf(is_currently_executing) AS count_executing,count() AS count_all FROM system.replication_queue GROUP BY database,table,type ORDER BY count_all DESC

Check whether an error message similar to the following is displayed:

```
Not executing fetch of part xxx because n fetches already executing, max n
```

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse** > **Configurations** > **All Configurations**, and check whether the value of **background_pool_size** is twice the number of cores on the node.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

Step 8 Set this parameter to twice the number of cores on the node and synchronize the configuration. Wait for a while and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check the connectivity between ClickHouse and ZooKeeper.

Step 9 Log in to the node where the ClickHouseServer instance is located, go to **${BIGDATA_HOME}/FusionInsight_ClickHouse_*/*_*_ClickHouseServer/etc**, and check whether the port numbers of the ClickHouseServer and ZooKeeper in the **config.xml** file are the same, as shown in the following information in bold:

NOTE

To view the ZooKeeper port number, choose **Cluster** > **Services** > **ZooKeeper** > **Configurations** > **All Configurations** on FusionInsight Manager, and check the value of **clientPort**.

```
<zookeeper>
  <session_timeout_ms>10000</session_timeout_ms>
  <node index="1">
    <host>server-2110082001-0019</host>
    <port>24002</port>
  </node>
  <node index="2">
    <host>server-2110082001-0018</host>
    <port>24002</port>
  </node>
  <node index="3">
    <host>server-2110082001-0017</host>
    <port>24002</port>
  </node>
</zookeeper>
```


- If yes, go to **Step 11**.
- If no, go to **Step 10**.

Step 10 Change the port number to the ZooKeeper port number, restart the ClickHouseServer instance, and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 11**.

Collect fault information.

Step 11 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 12 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 13 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 14 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 15 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.335 ALM-45441 Zookeeper Disconnected

Alarm Description

The system checks the connection between ClickHouse and ZooKeeper every minute. This alarm is generated when the connection fails. The alarm is reported because the ZooKeeper connection is abnormal. If the connection fails for three consecutive times, the system generates an alarm.

This alarm is automatically cleared when the system detects that the connection is normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45441	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If ClickHouse is disconnected from ZooKeeper, the ClickHouse service cannot be used.

Possible Causes

- The ZooKeeper service is abnormal.
- The ClickHouse service is overloaded.

Handling Procedure

Check whether ZooKeeper is normal.

Step 1 On FusionInsight Manager, choose **Cluster** > **Services** > **ZooKeeper** > **quorumpeer**.

Step 2 Check whether ZooKeeper instances are normal.

- If yes, go to **Step 6**.
- If no, go to **Step 3**.

Step 3 Select instances whose status is not good and choose **More** > **Restart Instance**.

Step 4 Check whether the instance status is good after restart.

- If yes, go to **Step 5**.
- If no, go to **Step 10**.

Step 5 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the ClickHouse service load is heavy.

Step 6 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 7 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- For a cluster with Kerberos authentication enabled (security mode):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**

- For a cluster with Kerberos authentication disabled (normal mode):

  clickhouse client --host IP address of the ClickHouseServer instance that reports the alarm **--user** *User name* **--password --port** 9440

Step 8 Run the following statement to check whether data is frequently written to the system table. If yes, wait until the service execution is complete and check whether the alarm is cleared.

SELECT query_id, user, FQDN(), elapsed, query FROM system.processes ORDER BY query_id;

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Check whether a large amount of data is written. If yes, wait until the task is complete and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 12 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 13 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 14 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.336 ALM-45442 Too Many Concurrent SQL Statements

Alarm Description

The alarm module checks the number of concurrent ClickHouse requests every 30 seconds. This alarm is generated when the number of concurrent ClickHouse requests exceeds the concurrency threshold configured on the UI.

This alarm is cleared when the system detects that the actual number of concurrent requests is less than concurrency threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45442	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If there are too many concurrent SQL statements, a large number of system resources are consumed. As a result, system response becomes slow.

Possible Causes

The ClickHouse service is overloaded.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Choose **Cluster** > **ClickHouse** > **Instance**, select an instance based on the alarm information. Choose **Chart** > **Concurrency** to check whether the actual number of concurrent SQL statements is greater than SQL concurrency threshold.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Confirm with the user whether a large number of tasks were being executed during the alarming period.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 On FusionInsight Manager, choose **O&M** and click **Alarm** > **Thresholds** in the navigation pane on the left. On the displayed page, click **ClickHouse** > **Concurrency** and adjust the threshold, or wait until the task is complete. Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 7 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.337 ALM-45443 Slow SQL Queries in the Cluster

Alarm Description

The system checks slow SQL queries for ClickHouse every 1 minute. This alarm is generated when the execution time of a SQL statement is longer than or equal to the slow SQL threshold.

This alarm is automatically cleared when the system detects that the execution time of the SQL statement is shorter than the slow SQL threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45443	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

The performance of the ClickHouse service deteriorates, which slows the response of other services. If there are too many slow SQL statements, the service may be unavailable.

Possible Causes

- The ClickHouse service is overloaded.
- The execution of SQL statements takes a long time.

Handling Procedure

Check whether the ClickHouse service load is heavy.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- For a cluster with Kerberos authentication enabled (security mode):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port --secure**

- For a cluster with Kerberos authentication disabled (normal mode):

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password --port**

Step 3 Run the following statement to check whether data is frequently written to the system table. If yes, wait until the service execution is complete and check whether the alarm is cleared.

SELECT query_id, user, FQDN(), elapsed, query FROM system.processes ORDER BY query_id;

- If yes, no further action is required.
- If no, go to **Step 4**.

Checking whether the SQL statements take a long time.

Step 4 Check the logical cluster to which the alarm object belongs. Log in to FusionInsight Manager, click **Cluster**, choose **Services** > **ClickHouse**, and click **Logic Cluster**. On the displayed page, choose **Query Management** > **Ongoing Slow Queries**. Check which SQL statements take a long time on the displayed page, confirm with the user to adjust services, optimize slow SQL statements, and check whether the optimization is successful.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 After the SQL statements are complete, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.338 ALM-45444 Abnormal ClickHouse Process

Alarm Description

The health check module checks ClickHouse instances every 30 seconds. If the number of consecutive failures exceeds the threshold, an alarm is reported. In this case, the ClickHouse process may stop responding and services cannot be properly executed.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45444	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

If the ClickHouse process is abnormal, services cannot run properly.

Possible Causes

The ClickHouse process runs improperly.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Log in to the node where the client is installed as the client installation user and run the following commands:

cd *{Client installation path}*

source bigdata_env

- For a cluster with Kerberos authentication enabled (security mode):

  kinit *Component service user*

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--port** 9440 **--secure**
- For a cluster with Kerberos authentication disabled (normal mode):

  clickhouse client --host *IP address of the ClickHouseServer instance that reports the alarm* **--user** *Username* **--password --port** 9000

Step 3 Run the following statement to check whether the result can be properly returned:

SELECT 1;

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Wait for several minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **ClickHouse** for the target cluster.

Step 7 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.339 ALM-45475 A Single Replica Exists in the Kudu Data Table

Alarm Description

The system checks the replica status of the Kudu data table. This alarm is generated when a single replica is detected in the Kudu data table.

This alarm is cleared when all Kudu data tables have multiple replicas or no data.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45475	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

If a hardware fault occurs, for example, a slow disk or a faulty disk, the Kudu data table may lose data.

Handling Procedure

Step 1 Run the following command to query details about each data table and check for the table with only one replica:

kudu table describe *{Kudu master address (separated by commas)} {Data table name}*

You can view details about the table with only one replica on the Kudu web UI.

Collect fault information.

Step 2 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 3 Expand the **Service** drop-down list, and select **Kudu** for the target cluster.

Step 4 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 5 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.340 ALM-45476 Kudu Failed to Enter the Maintenance Mode

Alarm Description

Kudu will enter the maintenance mode during disk replacement. This alarm is generated when Kudu fails to enter the maintenance mode.

This alarm is cleared when Kudu successfully enters the maintenance mode.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45476	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A slow or faulty disk cannot be replaced.

Possible Causes

- The Kudu instance is abnormal, and the process cannot be stopped.
- The permission on the **etc** configuration directory of Kudu is abnormal, and data cannot be written into the configuration file.

Handling Procedure

Collect fault information.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 2 Expand the **Service** drop-down list, and select **Kudu** for the target cluster.

Step 3 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 4 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.341 ALM-45477 Failed to Restore Data After a Disk of Kudu Is Replaced

Alarm Description

This alarm is generated when Kudu fails to restore data by invoking the script in SetupTool after a disk is replaced.

This alarm is cleared when the data is successfully restored.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45477	Critical	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

Kudu fails to restore data, and historical data is unavailable.

Possible Causes

Kudu fails to restore the UUID or the data from the remote end.

Handling Procedure

Collect fault information.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 2 Expand the **Service** drop-down list, and select **Kudu** for the target cluster.

Step 3 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

NOTE

Collect NodeAgent logs and check whether data is successfully restored after the disk replacement.

The logs are saved in the following directory:

/var/log/Bigdata/nodeagent/scriptlog/disk_manager_notify.log

Figure 10-129 Collecting NodeAgent logs

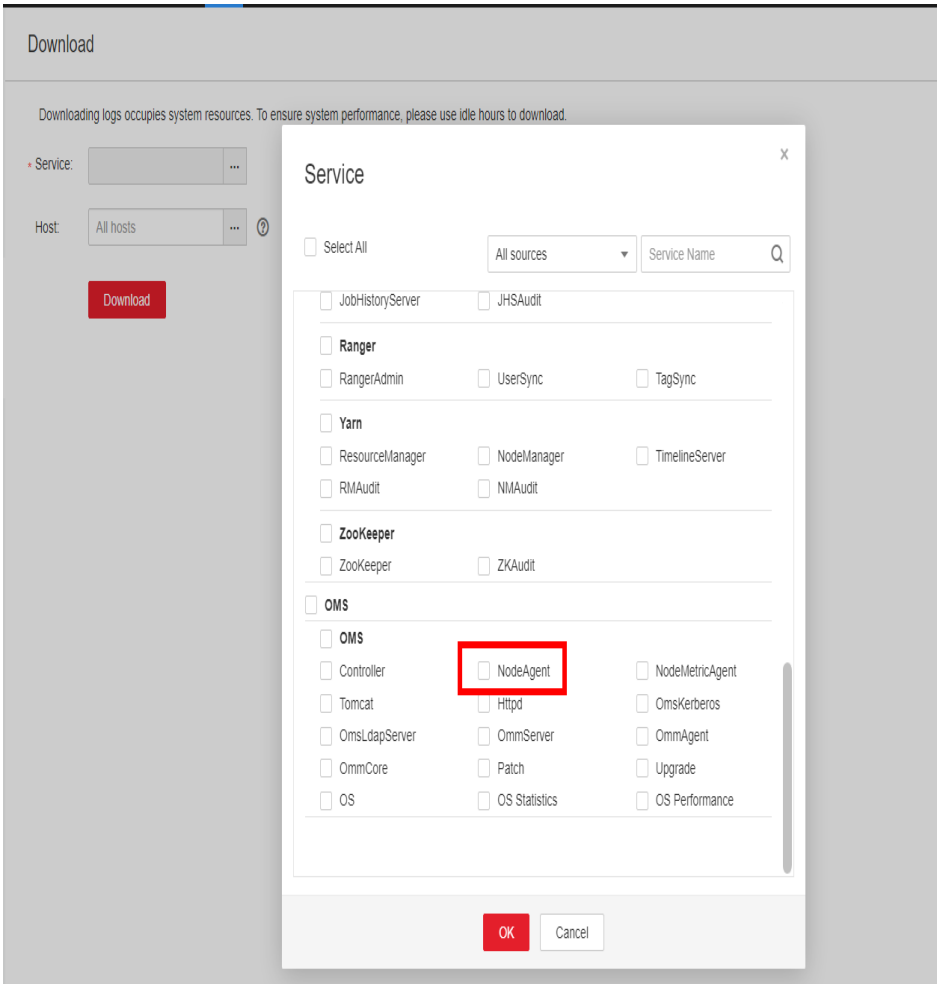

Step 4 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.342 ALM-45478 Kudu Failed Data Balancing

Alarm Description

The system periodically balances data among Kudu data tables. This alarm is generated when the system detects that the data balancing API returns a failure response.

This alarm is cleared when the Kudu data balancing API is successfully called.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45478	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

Failed to call the data balancing API. The Kudu component may be abnormal.

Possible Causes

The segment or replica in the Kudu data table is lost.

Handling Procedure

Handle the Kudu instance exception.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. On the displayed page, locate alarm **ALM-45478 Kudu Failed Data Balancing**.

Step 2 In the **Location** field, record the host name and role name.

Step 3 Choose **Cluster** > **Services** > **Kudu** > **Instances**. Click the role name for the host name obtained in **Step 2** and restore the instance. Then, check whether the alarm is cleared.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 On the Kudu web UI, check whether any segment or replica is missing.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Kudu** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.343 ALM-45479 Number of Tablets of the Tserver Process Exceeds the Threshold

Alarm Description

The system checks the Kudu service status every 60 seconds. This alarm is generated when the number of tablets of the Tserver process exceeds the threshold.

This alarm is cleared when the number of tablets of the Tserver process becomes normal and the system considers that the Kudu service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45479	Minor	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

If the number of tablets exceeds the threshold, the query performance of the Kudu engine deteriorates.

Possible Causes

The Tserver usage is too high or the Tserver load is unbalanced.

Handling Procedure

Handle the Kudu instance exception.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and check whether alarm **ALM-45479 Number of Tablets of the Tserver Process Exceeds the Threshold** exists.

- If yes, go to **Step 2**.
- If no, go to **Step 6**.

Step 2 Choose **O&M** > **Alarm** > **Thresholds** > **Kudu**, locate the threshold of this alarm, compare the threshold with the monitored number of tablets in the Tserver process of the cluster, and handle the alarm accordingly.

- If the threshold is not properly set, change the threshold and go to **Step 5**.
- If the Tserver load is unbalanced, go to **Step 3**.
- If the Tserver usage is too high, delete obsolete tables or add Tserver nodes, and go to **Step 3**.

Step 3 Log in to the Kudu node where the number of tablets exceeds the threshold.

Step 4 Run the following commands to balance the load of the cluster (during off-peak hours, recommended):

su omm

cd /opt/Bigdata/FusionInsight_Kudu_xxx/install/FusionInsight-Kudu-xxx/kudu/bin

./kudu cluster rebalance <master_addresses> [-tables=<tables>]

The preceding parameters can be obtained from the KuduMaster web UI. The **tables** parameter is optional.

Step 5 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Kudu** for the target cluster.

Step 8 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.344 ALM-45480 Tablet Leaders of a Tserver Process Are Unevenly Distributed

Alarm Description

The system checks the Kudu service status every 60 seconds. This alarm is generated when the ratio of tablet leaders of a Tserver process to total tablet leaders in the cluster exceeds the threshold.

This alarm is cleared when the ratio of tablet leaders of a Tserver process to total tablet leaders in the cluster becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45480	Minor	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.

Impact on the System

If Tserver tablet leaders are unevenly distributed, the query performance of the Kudu engine deteriorates.

Possible Causes

- Tserver instances are added or restarted.
- The threshold is not properly set.

Handling Procedure

Handle the uneven distribution of tablet leaders.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and check whether alarm **ALM-45480 Tablet Leaders of a Tserver Process Are Unevenly Distributed** exists.

- If yes, go to **Step 2**.
- If no, go to **Step 8**.

Step 2 Log in to the Kudu node where the number of tablet leaders exceeds the threshold.

Step 3 Run the following commands to manually adjust the number of tablet leaders:

su omm

cd /opt/Bigdata/FusionInsight_Kudu_*xxx*/install/FusionInsight-Kudu-xxx/kudu/bin

./kudu tablet leader_step_down *<master_addresses> <tablet_id>* **[-new_leader_uuid=***<new_tablet_server_uuid>***]**

NOTE

- **master_addresses**: The value is in the format of **KuduMaster service IP address 1:7051,KuduMaster service IP address 2:7051,KuduMaster service IP address 3:7051**.

  KuduMaster service IP address: You can log in to FusionInsight Manager and choose **Cluster** > **Services** > **Kudu** > **Instances** to view the service IP address of the KuduMaster instance.
- **tablet_id**: tablet ID to be adjusted.

  You can log in to Manager, choose **Cluster** > **Services** > **Kudu**, and click **KuduMaster(KuduMaster)** next to **KuduMaster WebUI** to access the KuduMaster web UI. Choose **Tables** on the menu bar, click the ID in the **Table Id** column, and obtain the target tablet ID in **Detail**.
- **new_tablet_server_uuid**: ID of the target tablet servers.

  You can log in to the KuduMaster web UI, choose **Tablet Servers** on the menu bar, and view the UUID.
- **new_leader_uuid**: You are advised to set the UUID of the node with a small number of tablet leaders.

Step 4 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Change the threshold.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and check whether alarm **ALM-45480 Tablet Leaders of a Tserver Process Are Unevenly Distributed** exists.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 Choose **Cluster** > **Kudu** > **Configurations** > **All Configurations** > **KuduMaster**, locate the alarm threshold parameter **TABLET_LEADER_UNBALANCE_SCALE**, change the value, and rolling-restart all Kudu Masters to make the configuration take effect.

Step 7 Choose **O&M** > **Alarm** > **Alarms** and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Kudu** for the target cluster.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.345 ALM-45481 KuduTserver Has Full Disks

The system checks Kudu disk metrics every 60 seconds. This alarm is generated when the number of fully occupied disks for a Tserver is not 0.

This alarm is automatically cleared when the number of fully occupied disks for the Tserver becomes 0.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45481	Major	Yes

Alarm Parameters

Type	Parameter	Description
Location Information	Source	Specifies the cluster for which the alarm was generated.
	ServiceName	Specifies the service for which the alarm was generated.
	RoleName	Specifies the role for which the alarm was generated.
	HostName	Specifies the host for which the alarm was generated.
Additional Information	Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Users cannot use the Kudu service properly.

Possible Causes

The disk configuration cannot meet service requirements, and the disk space used by the Kudu service is insufficient.

Handling Procedure

Check the disk capacity and delete unnecessary files and data tables.

Step 1 Log in to FusionInsight Manager. In the alarm list, click the drop-down arrow in the row that contains the alarm, and view the role name and the IP address of the hostname in **Location**.

Step 2 Log in to the node for which the alarm is generated as user **root** and run the **df -h** command to check the mount directory where the disk usage is 100%.

Step 3 Run the **find / -xdev -size +500M;** command to view files larger than 500 MB in the directory obtained in **Step 2**. Check whether such files are mistakenly written to the directory.

- If yes, go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Delete the mistakenly written files and check whether the alarm is cleared 2 minutes later.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Check whether Kudu contains unnecessary data tables.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 Use **kudu table delete <master_addresses> <table_name> [-nomodify_external_catalogs]** to delete unnecessary tables.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Step 8 Contact the disk administrator to expand the disk capacity.

Step 9 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, select **Kudu** for the target cluster, and click **OK**.

Step 12 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.346 ALM-45585 IoTDB Service Unavailable

Alarm Description

The system checks the IoTDB service status every 300 seconds. This alarm is generated when the IoTDB service is unavailable. This alarm is cleared when the IoTDB service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45585	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Users cannot use the IoTDB service properly.

Possible Causes

- The KrbServer service is abnormal.
- More than 50% of IoTDBServer instances are faulty.

Handling Procedure

Check whether the KrbServer service on which the IoTDB depends is abnormal.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether ALM-25500 KrbServer Service Unavailable exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by referring to "ALM-25500 KrbServer Service Unavailable."

Step 4 After ALM-25500 is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether IoTDBServer instances are faulty.

Step 5 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **IoTDB** > **Instance**.

Step 6 Check whether the percentage of faulty IoTDBServer instances exceeds 50%. If yes, restart the faulty IoTDBServer instances and check whether the status is restored.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.347 ALM-45586 IoTDBServer Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the IoTDBServer process status every 60 seconds. The alarm is generated when the heap memory usage of the IoTDBServer process exceeds the threshold (90% of the maximum memory).

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45586	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

If the available IoTDBServer process heap memory is insufficient, a memory overflow occurs and the service breaks down.

Possible Causes

The heap memory of the IoTDBServer process is overused or the heap memory is inappropriately allocated.

Handling Procedure

Check the heap memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row

containing the alarm whose **Alarm ID** is **45586**, view the role name in **Location**, and check the instance IP address.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Service** > **IoTDB** > **Instance**. Click the IoTDBServer for which the alarm is generated to go to **Dashboard**. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **Memory**. In the dialog box that is displayed, select **IoTDBServer Heap Memory Resource Percentage**, and click **OK**. Check whether the used non-heap memory of the IoTDBServer process reaches 90% (by default) of the maximum non-heap memory specified for IoTDBServer.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Choose **Cluster** > *Name of the desired cluster* > **Service** > **IoTDB** > **Configuration**, click **All Configurations**, choose **IoTDBServer** > **System**, and increase the value of **-Xmx** in the **GC_OPTS** parameter.

NOTE

- The default value of -**Xmx** is **2G**.
- If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.348 ALM-45587 IoTDBServer GC Duration Exceeds the Threshold

Alarm Description

The system checks the GC duration of the IoTDBServer process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default) for three consecutive times. You can choose **O&M** > **Alarm** > **Threshold Configuration** > *Name of the desired cluster* > **IoTDB** > **GC** > **Total GC duration of IoTDBServer process (IoTDBServer)** to change the threshold. This alarm is cleared when the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45587	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A long GC duration of the IoTDBServer process may interrupt the services.

Possible Causes

The heap memory of the IoTDBServer process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **45587**, view the role name in **Location**, and check the instance IP address.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Service** > **IoTDB** > **Instance**. Click the IoTDBServer for which the alarm is generated to go to **Dashboard**. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **GC**. In the dialog box that is displayed, select **Garbage Collection (GC) Time of IoTDBServer**, and click **OK**. Check whether the GC time of the IoTDBServer process is greater than 12 seconds.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Choose **Cluster** > *Name of the desired cluster* > **Service** > **IoTDB** > **Configuration**, click **All Configurations**, choose **IoTDBServer** > **System**, and increase the value of **-Xmx** in the **GC_OPTS** parameter.

NOTE

- The default value of -**Xmx** is **2G**.
- If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect the fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.349 ALM-45588 IoTDBServer Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the IoTDBServer service every 60 seconds. This alarm is generated when the direct memory usage of the IoTDBServer instance exceeds the threshold (90% of the maximum memory) for five consecutive times. This alarm is cleared when the IoTDBServer direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45588	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service breakdown.

Possible Causes

The direct memory of the IoTDBServer process is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**. On the page that is displayed, locate the row containing the alarm whose **Alarm ID** is **45588**, view the role name in **Location**, and check the instance IP address.

Step 2 Choose **Cluster** > *Name of the desired cluster* > **Service** > **IoTDB** > **Instance**. Click the IoTDBServer for which the alarm is generated to go to **Dashboard**. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **Memory**. In the dialog box that is displayed, select **IoTDBServer Direct Buffer Resource Percentage**, and click **OK**.

Step 3 Check whether the direct memory used by the IoTDBServer reaches the threshold (90% of the maximum direct memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **IoTDB** > **Configuration**, click **All Configurations**, choose **IoTDBServer** > **System**, increase the value of **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter as required, and save the configuration.

NOTE

- If this alarm is generated, the direct memory configured for the IoTDBServer process cannot meet the requirements of the IoTDBServer process.
- You are advised to set **-XX:MaxDirectMemorySize** in **GC_OPTS** to twice the direct memory used by the IoTDBServer process. (You can change the value based on the actual service scenario.)
- To obtain the size of the direct memory used by the IoTDBServer process, choose **Customize** > **Memory** > **IoTDBServer Direct Memory Resource Status**. If **GC_OPTS** does not contain the **-XX:MaxDirectMemorySize** parameter, add it manually.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect the fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **IoTDBServer** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.350 ALM-45589 ConfigNode Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the ConfigNode process every 60 seconds. This alarm is generated when the heap memory usage of the ConfigNode process exceeds the threshold (90% of the maximum memory). This alarm is cleared when the heap memory usage of the ConfigNode process is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45589	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

If the heap memory usage of the ConfigNode process is too high, the performance of the ConfigNode process is affected, and even the ConfigNode process becomes unavailable due to memory overflow.

Possible Causes

The heap memory configured for the node is improper. As a result, the usage exceeds the threshold.

Handling Procedure

Check the heap memory configuration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the real-time alarm list, click in front of this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **Cluster** > **Services** > **IoTDB**. Click **Instance**, click the ConfigNode corresponding to the IP address obtained in **Step 1**, and check whether **ConfigNode Heap Memory Usage** on the **Dashboard** tab page reaches the threshold specified for the ConfigNode process.

If the chart is not displayed, click the drop-down list in the upper right corner of the chart area and choose **Customize** > **Memory**. In the dialog box that is displayed, select **ConfigNode Heap Memory Usage** and click **OK**.

NOTE

You can choose **O&M** > **Alarm** > **Threshold Configuration** > *Name of the desired cluster* > **IoTDB** > **Memory** > **ConfigNode Heap Memory Usage (ConfigNode)** to view the threshold.

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Step 3 Choose **Cluster** > **Services** > **IoTDB**. Click **Configurations** then **All Configurations**, click **ConfigNode**, and choose **System**. Set **-Xmx** in **GC_OPTS** to a larger value and save the configuration.

NOTE

- The default value of -**Xmx** is **2G**.
- If this alarm is occasionally generated, increase the value of **-Xmx** by 0.5 times. If this alarm is frequently generated, double the value of **-Xmx**.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Click **Dashboard**. Click **Restart Service** to restart the IoTDB service for the configuration to take effect.

Step 5 Wait for about 120 seconds and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.351 ALM-45590 ConfigNode GC Duration Exceeds the Threshold

Alarm Description

The system checks the GC duration of the ConfigNode process every 60 seconds. This alarm is generated when the GC duration exceeds the threshold (12 seconds by default) for three consecutive times. This alarm is cleared when the GC duration is less than the threshold.

NOTE

You can choose **O&M** > **Alarm** > **Threshold Configuration** > *Name of the desired cluster* > **IoTDB** > **GC** > **Total GC duration of ConfigNode process (ConfigNode)** to increase the threshold by 20% each time.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45590	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.

Parameter	Description
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A long GC duration of the ConfigNode process may interrupt services.

Possible Causes

The heap memory configured on the node is improper. As a result, GC occurs frequently.

Handling Procedure

Check the heap memory configuration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the real-time alarm list, click in the row containing this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **Cluster** > **Services** > **IoTDB**. Click **Instance**, click the ConfigNode corresponding to the IP address obtained by **Step 1**. Switch to the **Dashboard** tab page, locate the **Total GC Duration of ConfigNode** chart, and check whether the GC duration of the ConfigNode process exceeds the threshold.

If the GC duration of ConfigNode is not displayed, click the drop-down list in the upper right corner of the chart area and choose **Customize** > **GC**. In the displayed dialog box, select **Total GC Duration of ConfigNode** and click **OK**.

NOTE

You can choose **O&M** > **Alarm** > **Threshold Configuration** > *Name of the desired cluster* > **IoTDB** > **GC** > **Total GC duration of ConfigNode process (ConfigNode)** to view the threshold.

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Step 3 Choose **Cluster** > **Services** > **IoTDB**. Click **Configurations** then **All Configurations**, click **ConfigNode**, and choose **System**. Set **-Xmx** in **GC_OPTS** to a larger value and save the configuration.

NOTE

- The default value of -**Xmx** is **2G**.
- If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Click **Dashboard**. Click **Restart Service** to restart the IoTDB service for the configuration to take effect.

Step 5 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 8 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.352 ALM-45591 ConfigNode Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the ConfigNode process every 60 seconds. This alarm is generated when the direct memory usage of the ConfigNode exceeds the threshold for five consecutive times. That is, the direct memory configured for ConfigNode cannot meet service requirements. This alarm is cleared when the direct memory usage of ConfigNode is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45591	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause the IoTDB instance to be unavailable.

Possible Causes

The direct memory configured for the node is improper. As a result, the usage exceeds the threshold.

Handling Procedure

Check the direct memory configuration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the real-time alarm list, click in front of this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **Cluster** > **Services** > **IoTDB**. Click **Instance**, click the ConfigNode corresponding to the IP address obtained in **Step 1**, and check whether **ConfigNode Direct Memory Usage** on the **Dashboard** tab page reaches the threshold specified for the ConfigNode process (90% of the maximum direct memory by default).

If the chart is not displayed, click the drop-down list in the upper right corner of the chart area and choose **Customize** > **Memory**. In the dialog box that is displayed, select **ConfigNode Direct Memory Usage** and click **OK**.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Services** > **IoTDB**. Click **Configurations** then **All Configurations**. Click **ConfigNode** and select **System**. Set **-XX:MaxDirectMemorySize** in **GC_OPTS** to a larger value as required and save the configuration.

NOTE

- You are advised to set **-XX:MaxDirectMemorySize** in **GC_OPTS** to twice the direct memory used by the ConfigNode process. (You can change the value based on the actual service scenario.)
- To obtain the size of the direct memory used by the ConfigNode process, choose **Customize** > **Memory** > **ConfigNode Direct Memory Resource Status**.
- If **GC_OPTS** does not contain the **-XX:MaxDirectMemorySize** parameter, add it.

Step 4 Restart the affected IoTDB service or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **ConfigNode** for the destination cluster.

Step 7 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.353 ALM-45592 IoTDBServer RPC Execution Duration Exceeds the Threshold

Alarm Description

The system checks the RPC execution duration of the IoTDBServer process every 60 seconds. This alarm is generated when the execution duration exceeds the

threshold. This alarm is cleared when the RPC execution time of the IoTDBServer process is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45592	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Running performance of the IoTDBServer process is affected.

Possible Causes

The processing duration of an IoTDBServer RPC request exceeds the threshold. Logs need to be further analyzed to locate the cause.

Handling Procedure

Collect fault information.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the real-time alarm list, click in front of this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **O&M** > **Log** > **Download**.

Step 3 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 4 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in **Step 1**, and click **OK**.

Step 5 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 6 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.354 ALM-45593 IoTDBServer Flush Execution Duration Exceeds the Threshold

Alarm Description

This alarm is generated when the data flush duration exceeds the threshold. This alarm is cleared when the flush duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45593	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Data write is blocked and the write operation performance is affected.

Possible Causes

The IoTDB flushing on the node is slow. You need to further analyze logs.

Handling Procedure

Collect fault information.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the real-time alarm list, click  in front of this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **O&M** > **Log** > **Download**.

Step 3 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 4 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in **Step 1**, and click **OK**.

Step 5 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 6 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.355 ALM-45594 IoTDBServer Intra-Space Merge Duration Exceeds the Threshold

Alarm Description

This alarm is generated when the merge duration in the space exceeds the threshold. This alarm is cleared when the merge duration in the space is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45594	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Data write is blocked and the write operation performance is affected.

Possible Causes

The merge task in the IoTDB space of the node is slow. You need to further analyze logs.

Handling Procedure

Collect fault information.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the real-time alarm list, click in front of this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **O&M** > **Log** > **Download**.

Step 3 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 4 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in **Step 1**, and click **OK**.

Step 5 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 6 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.356 ALM-45595 IoTDBServer Cross-Space Merge Duration Exceeds the Threshold

Alarm Description

This alarm is generated when the cross-space merge duration exceeds the threshold. This alarm is cleared when the cross-space merge duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45595	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Data write is blocked and the write operation performance is affected.

Possible Causes

The IoTDB cross-space merge task on the node is slow. You need to further analyze logs.

Handling Procedure

Collect fault information.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the real-time alarm list, click in front of this alarm and view the role name and instance IP address in **Location**.

Step 2 Choose **O&M** > **Log** > **Download**.

Step 3 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 4 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in **Step 1**, and click **OK**.

Step 5 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 6 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.357 ALM-45596 Procedure Execution Failed

Alarm Description

Procedures are the tasks managed and executed by the ConfigNode leader. This alarm is generated when a procedure fails to be executed. This alarm is cleared when the procedure is successfully executed.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45596	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
ProcedureInformation	Specifies the procedure-related information.

Impact on the System

Functions associated with the procedure are adversely affected.

Possible Causes

- The task for adding IoTDB replicas fails to be executed.
- The task for deleting the storage group fails to be executed.

Handling Procedure

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, locate this alarm, and click .

Step 2 Check the value of **ProcedureInformation** in **Location**. The value starts with the procedure type and contains main information about the procedure.

Check whether the task for adding replicas fails.

Step 3 Check whether the value of **ProcedureInformation** starts with **AddRegionProcedure** or **ReJoinDataNodeProcedure**.

- If yes, the task fails. Go to **Step 4**.
- If no, go to **Step 5**.

Step 4 Wait for half an hour. If the region is successfully added, the alarm is automatically cleared. Otherwise, go to **Step 5**.

Check whether the task for deleting the storage group fails.

Step 5 Check whether the value of **ProcedureInformation** starts with **DeleteStorageGroupProcedure**.

- If yes, the storage group fails to be deleted. Go to **Step 6**.
- If no, go to **Step 7**.

Step 6 Delete the storage group displayed in **ProcedureInformation** again on the IoTDB client. If the deletion is successful, the alarm is automatically cleared. Otherwise, go to **Step 7**.

Collect fault information.

Step 7 Choose **Cluster** > **Services** > **IoTDB** > **Instance** to view the hosts where all IoTDBServer and ConfigNode instances are located.

Step 8 Choose **O&M** > **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, select **IoTDB** for the target cluster, and click **OK**.

Step 10 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the host queried in **Step 7**, and click **OK**.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.358 ALM-45615 CDL Service Unavailable

Alarm Description

The system checks the CDL health status every 60 seconds. This alarm is generated when the CDL health status is **DOWN**. This alarm is cleared when the system detects that the CDL health status is **UP**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45615	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.

Parameter	Description
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The CDL service is abnormal. You cannot use FusionInsight Manager to perform cluster operations on the CDL service. The CDL service function is unavailable.

Possible Causes

All CDLService or CDLConnector instances of the CDL service are abnormal, and the Kafka service is unavailable.

Handling Procedure

Check whether the Kafka service on which the CDL service depends is abnormal.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether ALM-38000 Kafka Service Unavailable exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by referring to "ALM-38000 Kafka Service Unavailable".

Step 4 After the alarm is cleared, wait a few minutes and check whether the alarm HetuServer Service Unavailable is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether CDL instances are faulty.

Step 5 On FusionInsight Manager, choose **Cluster** > *Name of the desired cluster* > **Service** > **CDL** > **Instance**.

Step 6 Check whether all CDLService and CDLConnector instances are faulty.

- If yes, restart the CDL service and choose **Cluster** > *Name of the desired cluster* > **Services** > **CDL** > **More** > **Restart Service**. If the fault persists after the restart, go to **Step 7** and contact O&M personnel to check CDL logs.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **CDL** for the target cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the service is restored, the system automatically clears the alarm.

Related Information

None

10.359 ALM-45616 CDL Job Execution Exception

Alarm Description

The system checks whether a CDL job is normal every 60 seconds. This alarm is reported when the CDL job is abnormal. This alarm is cleared when the job is restored or stopped.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45616	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.
Username	Specifies the username of the job for which the alarm was generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

The CDL task fails to be executed due to incorrect parameter settings or other reasons. On the **Job Management** page of the CDL web UI, locate the row where the job is located and click **Failed**/**Abnormal running** in the **Status** column to view the failure cause, or view the failure cause in the logs.

Handling Procedure

Step 1 Log in to FusionInsight Manager as a user who has the CDL job creation or administrator permission.

Step 2 Choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**, click ⌄ in the row where **Alarm ID** is **45616**, and view the name of the job for which this alarm is generated in **Location**.

Step 3 Choose **Cluster** > **Services** > **CDL** and click the link next to **CDLService UI** to go to the CDL web UI.

Step 4 Locate the row where the failed job is located based on the job name obtained in **Step 2**, and click **Abnormal running** or **Failed** in the **Status** column.

Step 5 On the page that is displayed, view the error information and rectify the fault. For example, **Figure 10-130** shows that the task running on Yarn is manually killed. For details, see trace error information, as shown in **Figure 10-131**.

Figure 10-130 CDL job exception

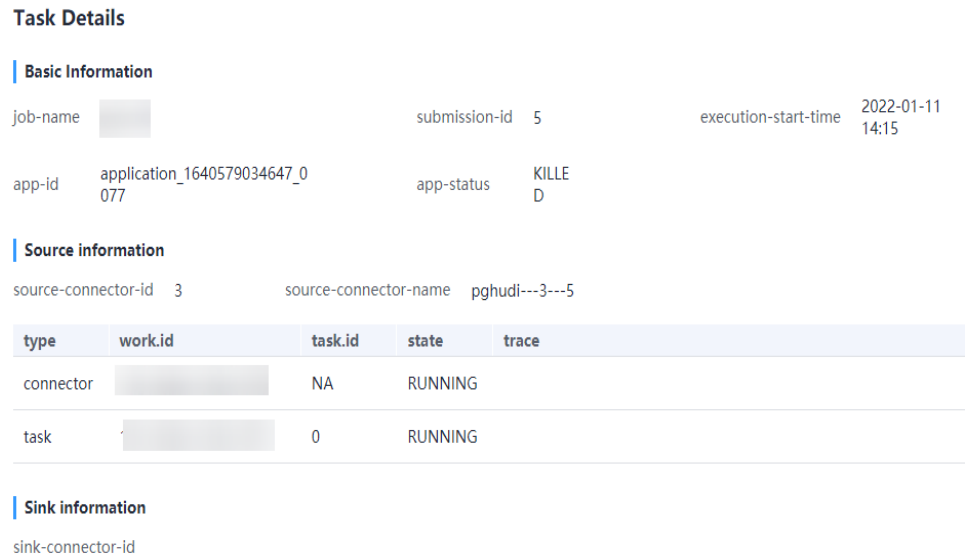

Figure 10-131 Trace error information

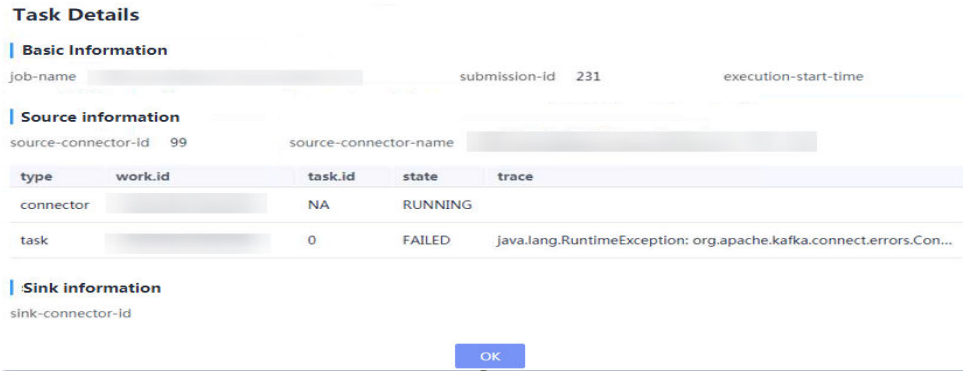

Step 6 Rectify the fault based on the error information, execute the task again, and check whether the task can be executed successfully.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect the fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Select **CDL** in the required cluster for **Service**.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

After the job is successfully restored or stopped, the alarm is cleared if it has been reported.

Related Information

None

10.360 ALM-45617 Data Queued in the CDL Replication Slot Exceeds the Threshold

Alarm Description

If too many WALs are stacked in PostgreSQL or OpenGauss (applicable to MRS 3.3.0 or later), the PostgreSQL or OpenGauss disk space may be used up. The system checks whether the amount of data queued in the replication slot configured for a CDL job exceeds the threshold every 5 minutes. This alarm is generated when the amount of data queued in the replication slot exceeds the

threshold. This alarm is cleared when the number of data queued in the replication slot falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45617	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.
DBName	Specifies the database for which the alarm was generated.
SlotName	Specifies the database replication slot for which the alarm was generated.
Lag	Specifies the data queued in the slot.

Impact on the System

WALs are continuously accumulated in the source PostgreSQL or OpenGauss database, causing the disk space of the source database to be used up.

Possible Causes

The CDL job is abnormal, and data processing stops; the source database is updated quickly, and CDL data processing is slow.

Handling Procedure

Step 1 Log in to FusionInsight Manager as a user who has the CDL job creation or administrator permission.

Step 2 Choose **O&M**. In the navigation pane on the left, choose **Alarm** > **Alarms**, click ⌄ in the row where **Alarm ID** is **45617**, and view the name of the job for which this alarm is generated in **Location**.

Step 3 Check whether **ALM-45616 CDL Job Execution Exception** is displayed in the alarm list.

- If yes, handle the alarm by performing operations provided for **ALM-45616 CDL Job Execution Exception**.
- If no, go to **Step 4**.

Step 4 Choose **Cluster** > **Services** > **CDL**. Click the link next to **CDLService UI** to go to the CDL web UI and check whether the job is displayed in the job list based on its name obtained in **Step 2**.

- If yes, check whether the job is abnormal.
  - If it is abnormal, go to **Step 5**.
  - If it is not, data processing is slow. Contact O&M personnel.
- If no, go to **Step 7**.

Step 5 Click **Abnormal** or **Failed** in the row where the job is located and rectify the fault based on the error information displayed on the page.

Step 6 After rectifying the fault, run the job again and check whether the job can be executed successfully.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **CDL** for the target cluster.

Step 9 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is cleared when the amount of data queued in the replication slot is less than the threshold. You do not need to manually clear the alarm.

Related Information

None

10.361 ALM-45635 FlinkServer Job Execution Failure

This section applies to MRS 3.1.2 or later.

Alarm Description

The system checks whether FlinkServer jobs fail to be executed every 10 seconds. This alarm is generated when a FlinkServer job fails. This alarm is cleared when the job is successfully restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45635	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

You can view failure causes in specific logs.

Handling Procedure

Step 1 Log in to Manager as a user who has the FlinkServer management permission.

Step 2 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

Step 3 Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 10-132 Application ID of a job

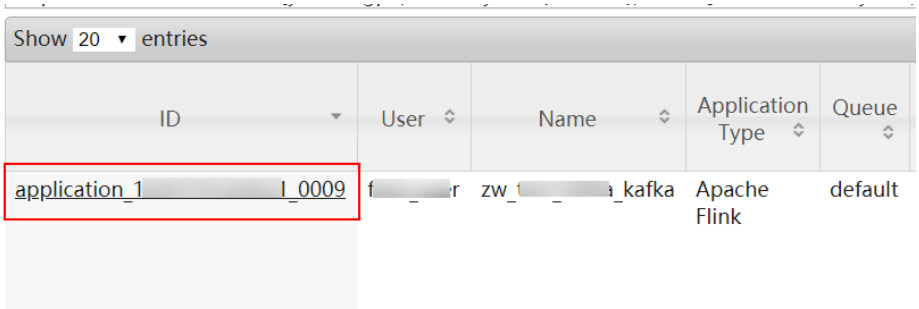

If yes, go to **Step 4**.

If no, go to **Step 6**.

Step 4 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-133 Clicking Logs

   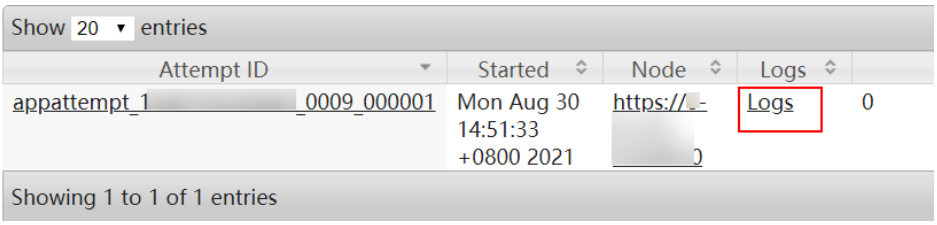

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

   Figure 10-134 Clicking the ID in the Attempt ID column

   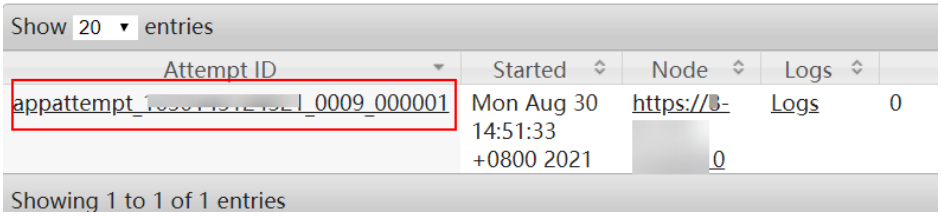

   Figure 10-135 Clicking Logs

   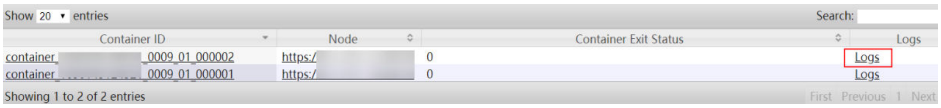

   NOTE

   You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster** > **Services** > **Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More** > **Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***User name*/**logs/***Application ID of the failed job* directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearance

After the job is successfully restarted, the alarm is cleared if it has been reported.

Related Information

None

10.362 ALM-45636 FlinkServer Job Checkpoints Keep Failing

This section applies to MRS 3.1.2 or a version between 3.1.2 and 3.2.0.

Alarm Description

The system checks the number of consecutive checkpoint failures based on the configured alarm checking interval. This alarm is generated when the number of consecutive checkpoint failures of a FlinkServer job reaches the configured threshold. This alarm is cleared when checkpoints are recovered or the job is successfully restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45636	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

You can view failure causes in specific logs.

Handling Procedure

Step 1 Log in to Manager as a user who has the FlinkServer management permission.

Step 2 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

Step 3 Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 10-136 Application ID of a job

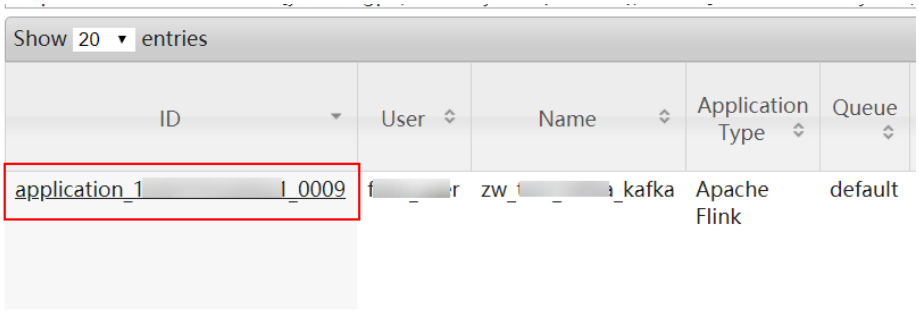

If yes, go to **Step 4**.

If no, go to **Step 6**.

Step 4 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-137 Clicking Logs

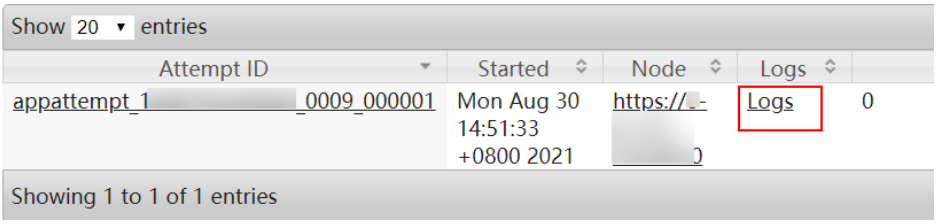

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 10-138 Clicking the ID in the Attempt ID column

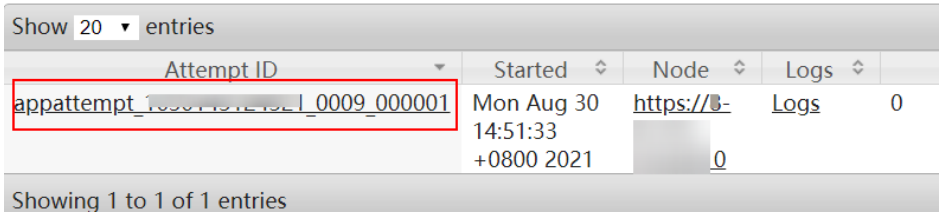

Figure 10-139 Clicking Logs

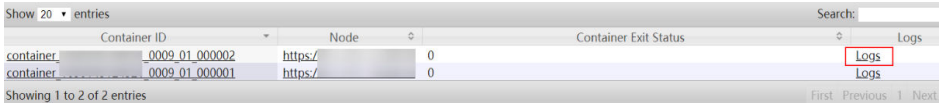

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster** > **Services** > **Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More** > **Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***User name*/**logs/***Application ID of the failed job* directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is cleared when FlinkServer job checkpoints are recovered or the job is successfully restarted.

Related Information

None

10.363 ALM-45636 Flink Job Checkpoints Keep Failing

This section applies to MRS 3.2.0 or later.

Description

The system checks the number of consecutive checkpoint failures based on the configured alarm checking interval. This alarm is generated when the number of consecutive checkpoint failures of a FlinkServer job reaches the configured threshold. This alarm is cleared when checkpoints are recovered or the job is successfully restarted.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45636	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.
Username	Specifies the username of the job for which the alarm is generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

You can view failure causes in specific logs.

Procedure

Step 1 Log in to Manager as a user who has the FlinkServer management permission.

Step 2 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

Step 3 Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 10-140 Application ID of a job

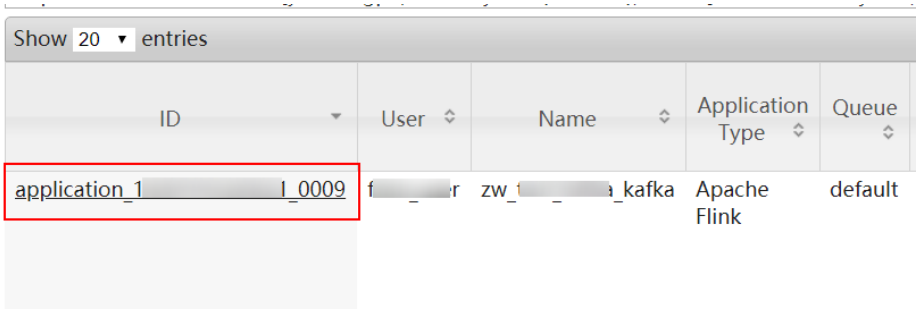

If yes, go to **Step 4**.

If no, go to **Step 6**.

Step 4 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-141 Clicking Logs

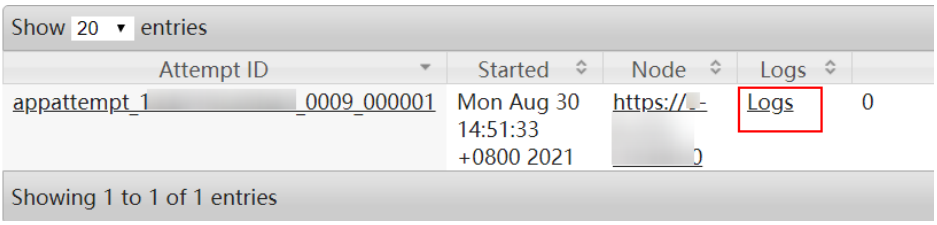

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 10-142 Clicking the ID in the Attempt ID column

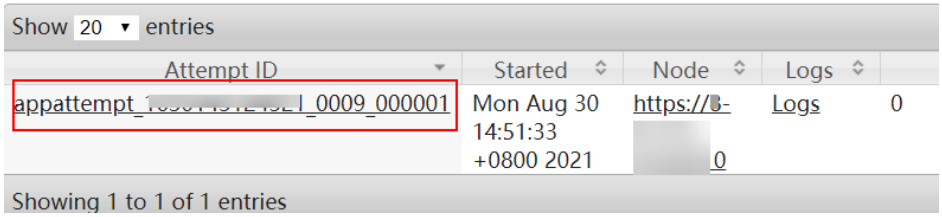

Figure 10-143 Clicking Logs

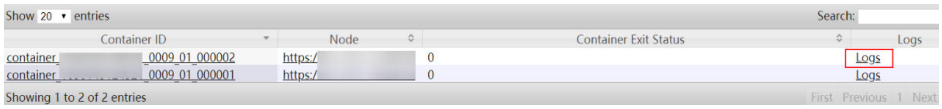

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster** > **Services** > **Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More** > **Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***User name*/**logs/***Application ID of the failed job* directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearing

This alarm is cleared when Flink job checkpoints are recovered or the job is successfully restarted.

Related Information

None

10.364 ALM-45637 FlinkServer Task Is Continuously Under Back Pressure

This section applies to MRS 3.1.2 or later.

Alarm Description

The system checks the back pressure duration of FlinkServer tasks based on the configured alarm checking interval. This alarm is generated when the back pressure duration of a FlinkServer task reaches the configured threshold. This alarm is cleared when the task back pressure is recovered or the job is successfully restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45637	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

You can view the causes in the specific logs.

Handling Procedure

Step 1 Log in to Manager as a user who has the FlinkServer management permission.

Step 2 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

Step 3 Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 10-144 Application ID of a job

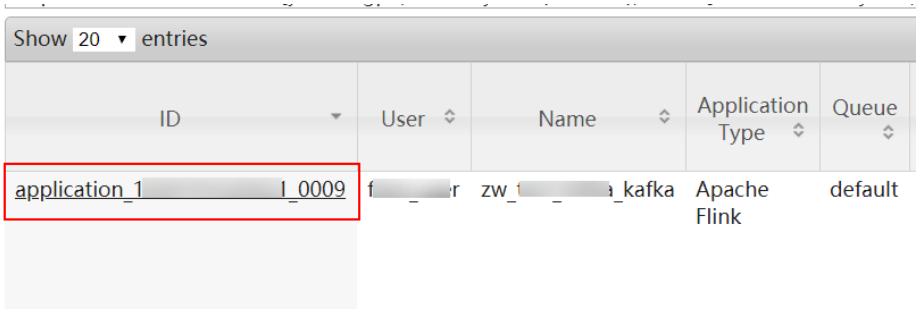

If yes, go to **Step 4**.

If no, go to **Step 6**.

Step 4 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-145 Clicking Logs

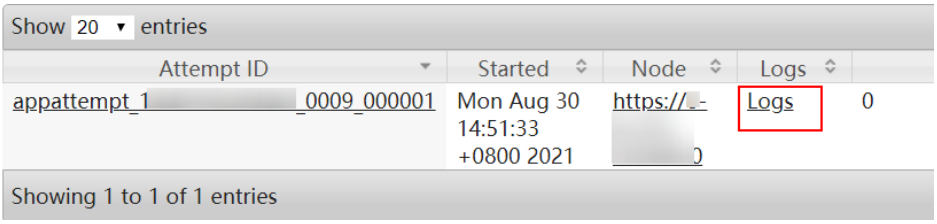

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 10-146 Clicking the ID in the Attempt ID column

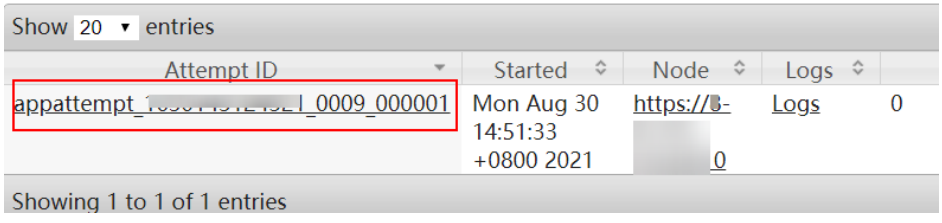

Figure 10-147 Clicking Logs

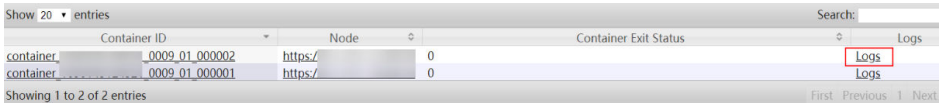

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster** > **Services** > **Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More** > **Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***User name*/**logs/***Application ID of the failed job* directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is cleared when FlinkServer task back pressure is recovered or the job is successfully restarted.

Related Information

None

10.365 ALM-45638 Number of Restarts After FlinkServer Job Failures Exceeds the Threshold

This section applies to MRS 3.1.2 or a version between 3.1.2 and 3.2.0.

Alarm Description

The system checks the number of FlinkServer job restarts based on the alarm checking interval. This alarm is generated when the number exceeds the configured threshold. This alarm is cleared when the job is restarted.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45638	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
JobName	Specifies the job for which the alarm was generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

You can view the causes in the specific logs.

Handling Procedure

Step 1 Log in to Manager as a user who has the FlinkServer management permission.

Step 2 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

Step 3 Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 10-148 Application ID of a job

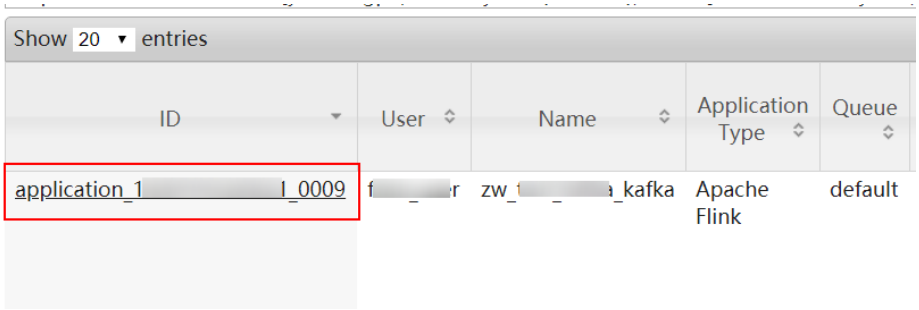

If yes, go to **Step 4**.

If no, go to **Step 6**.

Step 4 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-149 Clicking Logs

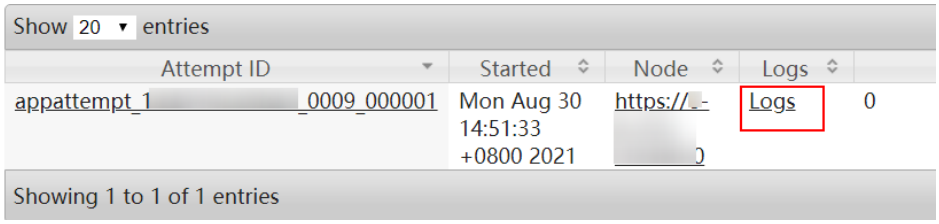

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 10-150 Clicking the ID in the Attempt ID column

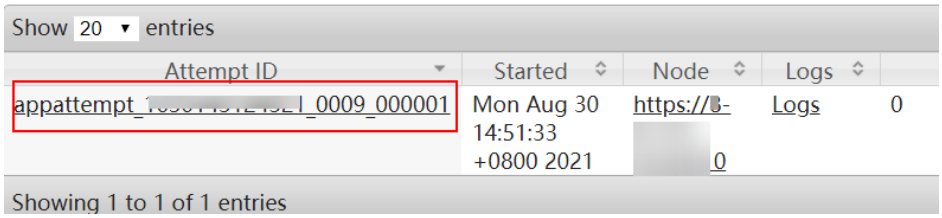

Figure 10-151 Clicking Logs

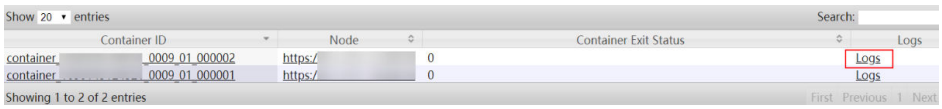

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster** > **Services** > **Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More** > **Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***User name*/**logs/***Application ID of the failed job* directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is cleared when the FlinkServer job is successfully restarted.

Related Information

None

10.366 ALM-45638 Number of Restarts After Flink Job Failures Exceeds the Threshold

This section applies to MRS 3.2.0 or later.

Description

The system checks the number of Flink job restarts based on the alarm checking interval. This alarm is generated when the number exceeds the configured threshold. This alarm is cleared when the job is restarted.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45638	Major	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.
Username	Specifies the username of the job for which the alarm is generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

You can view the causes in the specific logs.

Procedure

Step 1 Log in to Manager as a user who has the FlinkServer management permission.

Step 2 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the Yarn page.

Step 3 Locate the failed job based on its name displayed in **Location**, search for and record the application ID of the failed job, and check whether the job logs are available on the Yarn page.

Figure 10-152 Application ID of a job

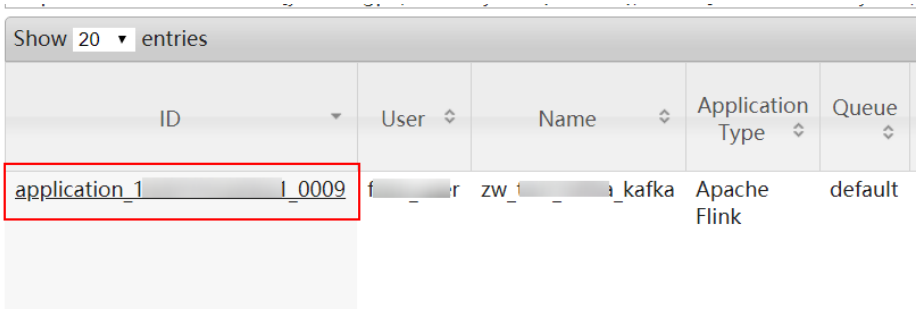

If yes, go to **Step 4**.

If no, go to **Step 6**.

Step 4 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-153 Clicking Logs

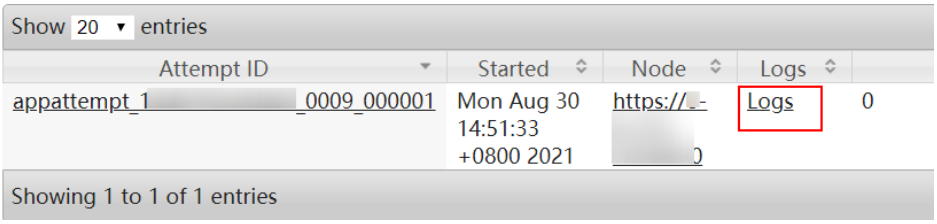

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

Figure 10-154 Clicking the ID in the Attempt ID column

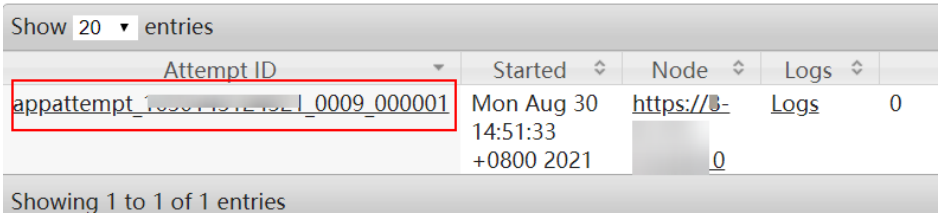

Figure 10-155 Clicking Logs

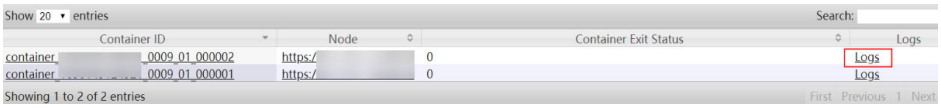

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission, choose **Cluster** > **Services** > **Flink**, click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management** and choose **More** > **Job Monitoring** in the **Operation** column to view the TaskManager logs.

Step 5 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, select **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***User name*/**logs/***Application ID of the failed job* directory.

Step 7 View the logs of the failed job to rectify the fault, or contact the O&M personnel personnel and send the collected fault logs.

----End

Alarm Clearing

This alarm is cleared when the Flink job is successfully restarted.

Related Information

None

10.367 ALM-45639 Checkpointing of a Flink Job Times Out

Description

The system checks the checkpointing timeout of Flink jobs every 30 seconds. This alarm is generated if the checkpointing timeout of a Flink job is longer than the threshold (600 seconds by default). This alarm is cleared when the checkpointing timeout of a job is less than or equal to the threshold.

Attribute

Alarm ID	Alarm Severity	Auto Clear
45639	Minor	Yes

Parameters

Name	Meaning
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName (available in MRS 3.2.1 or later)	Specifies the name of the application for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.
UserName	Specifies the username for which the alarm is generated.

Impact on the System

This alarm has no impact on the system.

Possible Causes

The job may be in the sub-healthy state. The possible causes are as follows:

- The memory for the TaskManager of the job is insufficient.
- The state memory is too large, making checkpointing time-consuming.

Procedure

Step 1 Log in to Manager as a user who has the FlinkServer management permission.

Step 2 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45639 Checkpointing of a Flink Job Times Out**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 3 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 4 Locate the failed task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-156 Application ID of a job

- If yes, go to **Step 5**.
- If no, go to **Step 7**.

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-157 Clicking Logs

   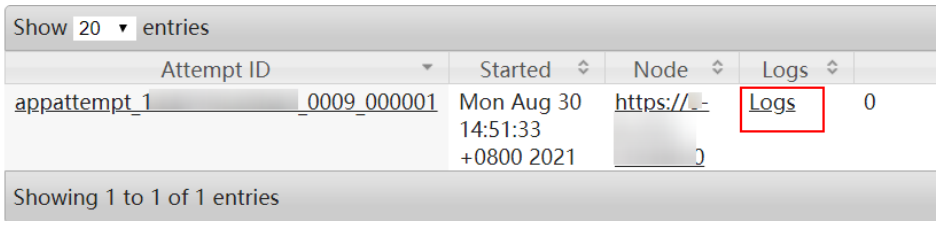

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view TaskManager logs.

   Figure 10-158 Clicking the ID in the Attempt ID column

   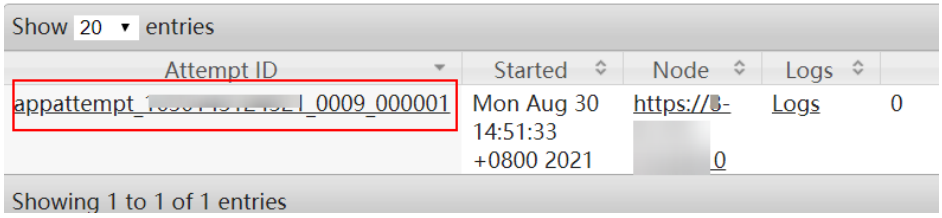

   Figure 10-159 Clicking Logs

   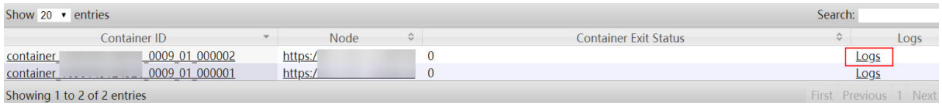

   NOTE

   You can also log in to Manager as a user who has the FlinkServer management permission. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 6 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 7 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**logs/***Application ID of the failed job* directory.

Step 8 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearing

This alarm is cleared when the checkpointing timeout a Flink job is less than or equal to the threshold.

Related Information

None

10.368 ALM-45640 FlinkServer Heartbeat Interruption Between the Active and Standby Nodes

This section applies to MRS 3.2.0 or later.

Alarm Description

This alarm is generated when the FlinkServer active node or standby node does not receive heartbeat messages from the peer for 30 seconds (heartbeat interruption duration configured in keepalive).

This alarm is cleared when the heartbeat recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45640	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

During the FlinkServer heartbeat interruption, only one node can provide the service. If this node is faulty, no standby node is available for failover and the service is unavailable.

Possible Causes

- The active or standby FlinkServer instance is in the stopped state.
- The NIC of the floating IP address of the HA system used by the FlinkServer node is incorrectly configured. FlinkServer fails to be started.
- The link between the active and standby FlinkServer nodes is abnormal.

Handling Procedure

Check the status of the active and standby FlinkServer instances.

Step 1 Log in to FusionInsight Manager, choose **Cluster** > **Services** > **Flink** > **Instance**, and check the state of FlinkServer is normal.

- If yes, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 Select the abnormal FlinkServer instance and start the instance. After the instance is started, check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

NOTICE

During the restart, the FlinkServer instance cannot provide services, but submitted jobs are not affected.

Check whether the link between the standby FlinkServer nodes is normal.

Step 3 Choose **Cluster** > **Services** > **Flink** > **Instance**, and check the two service IP addresses of FlinkServer.

Step 4 Log in to the server where the abnormal FlinkServer instance locates as the **root** user.

Step 5 Run the following command to check whether the server of the other FlinkServer instance is reachable:

ping *IP address of the other FlinkServer instance*

- If yes, go to **Step 8**.
- If no, go to **Step 6**.

Step 6 Ask the network administrator to handle the network exception.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check whether the logs of the node where the abnormal FlinkServer instance locates contains error information.

Step 8 Log in to the server where the abnormal FlinkServer instance locates as the **root** user.

Step 9 Open the log file in the default directory **/var/log/Bigdata/flink/flinkserver/prestart.log** and check whether there is error message **Float ip** *x.x.x.x* **is invalid**.

- If yes, go to **Step 10**.
- If no, go to **Step 12**.

Step 10 On FusionInsight Manager, choose **Cluster** > **Services** > **Flink** > **Configurations** > **All Configurations** and search for **flink.ha.floatip**. Change the parameter value to the correct floating IP address, save the configuration, and restart the Flink service.

NOTICE

- Contact the network engineer to obtain the new floating IP address.
- During the service restart, FlinkServer cannot provide services, but submitted jobs are not affected.
- During the restart, the FlinkServer instance cannot provide services, but submitted jobs are not affected.

Step 11 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Collect the fault information.

Step 12 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 13 Select the Flink service in the required cluster for **Service**.

Step 14 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 15 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 16 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.369 ALM-45641 Data Synchronization Exception Between the Active and Standby FlinkServer Nodes

This section applies to MRS 3.2.0 or later.

Alarm Description

The system checks data synchronization between the active and standby FlinkServer nodes every 60 seconds. This alarm is generated when the standby FlinkServer node fails to synchronize files with the active FlinkServer node.

This alarm is cleared when the standby FlinkServer synchronizes files with the active FlinkServer.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45641	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Because the configuration files on the standby FlinkServer are not updated, some configurations will be lost after an active/standby switchover. FlinkServer and some components may not run properly.

Possible Causes

- The link between the active and standby FlinkServer nodes is interrupted.
- The synchronization file does not exist or the file permission is required.

Handling Procedure

Check whether the network between the active and standby FlinkServer is in normal state.

Step 1 On FusionInsight Manager, choose **Cluster** > **Services** > **ClickHouse** > **Instance**. View and record the IP addresses of active and standby FlinkServer.

Step 2 Log in to the active FlinkServer node as the **root** user.

Step 3 Run the following command to check whether the standby FlinkServer is reachable:

ping *IP address of the standby FlinkServer*

- If yes, go to **Step 6**.
- If no, go to **Step 4**.

Step 4 Contact the network administrator to check whether the network is faulty.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Rectify the network fault and check whether the alarm is cleared from the alarm list.

- If yes, no further action is required.
- If no, go to **Step 6**.

Check whether the storage space of the /srv/BigData/LocalBackup directory is insufficient.

Step 6 Run the following command to check whether the storage space of the **/srv/BigData/LocalBackup** directory is insufficient:

df -hl /srv/BigData/LocalBackup

- If yes, go to **Step 7**.
- If no, go to **Step 10**.

Step 7 Run the following command to clear unnecessary backup files:

rm -rf *Directory to be cleared*

The following are two examples:

rm -rf /srv/BigData/LocalBackup/0/default-oms_20191211143443

Step 8 On FusionInsight Manager, choose **O&M** > **Backup and Restoration** > **Backup Management**.

In the **Operation** column of the backup task, click **Configure** and change the value of **Maximum Number of Backup Copies** to reduce the number of backup file sets.

Step 9 Wait for 1 minute and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 10**.

Check whether the synchronization file exists and whether the file permission is valid.

Step 10 Run the following command to check whether the synchronization file exists:

find /srv/BigData/ -name "sed*"

find /opt -name "sed*"

- If yes, go to **Step 11**.

- If no, go to **Step 12**.

Step 11 Run the following command to check the synchronization file information and permission queried in **Step 10**:

ll *Path of the file you want to search for*

- If the file size is 0 and all values in the permission column are -, the file is a junk file. Run the following command to delete it:

  rm -rf *Files to be deleted*

  Wait for several minutes and check whether the alarm is cleared. If the alarm persists, go to **Step 12**.
- If the file size is not 0, go to **Step 12**.

Step 12 View the log file generated when the alarm is reported.

1. Run the following command to go to the HA run log file path of the current cluster:

   cd /var/log/Bigdata/flink/flinkserver/ha/runlog
2. Decompress log file and view the logs generated when the alarm is reported.

   For example, if the name of the file is **ha.log.2021-03-22_12-00-07.gz**, run the following command:

   gunzip *ha.log.2021-03-22_12-00-07.gz*

   vi *ha.log.2021-03-22_12-00-07*

   Check whether error information is displayed before and after the alarm generation time in the logs.

   – If it is displayed, rectify the fault based on the error information. Go to **Step 13**.

   For example, if the following error information is displayed, the directory permission is required. In this case, obtain the directory permission that is the same as the permission on a normal node.

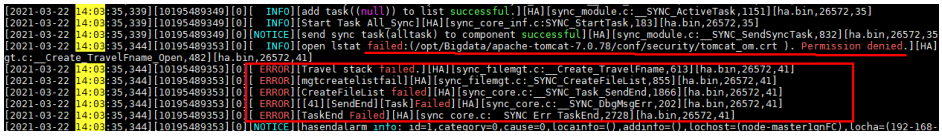

   – If no, go to **Step 14**.

Step 13 Wait for about 10 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 14**.

Collect fault information.

Step 14 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 15 Select FlinkServer information from **Services** and click **OK**.

Step 16 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the hosts to which the role belongs, and click **OK**.

Step 17 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 18 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.370 ALM-45642 RocksDB Continuously Triggers Write Traffic Limiting

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when RocksDB for a job continuously triggers write traffic limiting, that is, the RocksDB write rate is not 0. This alarm is cleared when the RocksDB write rate of the job becomes 0.

The **rocksdb.actual-delayed-write-rate** parameter specifies the RocksDB write rate of a job. Value **0** indicates that the rate is not limited, and other values indicate traffic limiting.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45642	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

When the rate at which Flink jobs write data to RocksDB is not 0, write traffic limiting is triggered. The possible causes are as follows:

- There are too many MemTables. As a result, write traffic is limited or write stops, and **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** is generated.
- The size of SST files at level 0 is too large, and **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The estimated compaction size exceeds the threshold, and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** is generated.

Handling Procedure

Check whether write traffic limiting or write stop is caused due to too many MemTables.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by following the instructions provided in section **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45643 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether write traffic limiting or write stop is caused due to too many SST files at level 0.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 6 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 8 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether write traffic limiting or write stop is caused because the estimated compaction size exceeds the threshold.

Step 9 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 10**.
- If no, go to **Step 12**.

Step 10 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 11 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Collect fault information.

Step 12 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 13 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45642 RocksDB Continuously Triggers Write Traffic Limiting**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 14 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 15 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the native Yarn page.

Figure 10-160 Application ID of a job

- If yes, go to **Step 16**.
- If no, go to **Step 18**.

Step 16 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-161 Clicking Logs

   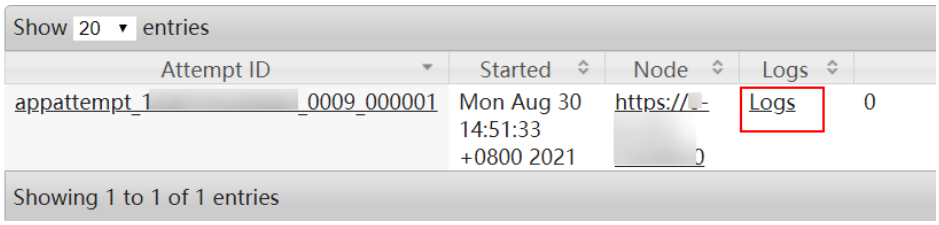

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs.

   Figure 10-162 Clicking the ID in the Attempt ID column

   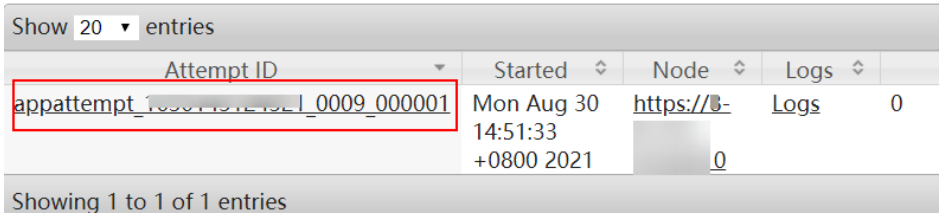

   Figure 10-163 Clicking Logs

   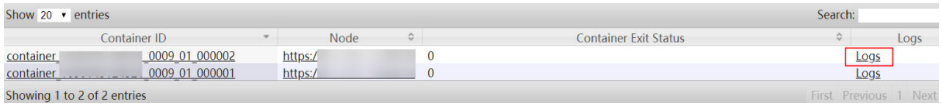

   NOTE

   You can also log in to Manager as a user who has the FlinkServer management permission. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 17 View the job logs to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 18 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Step 19 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.371 ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the MemTable size of RocksDB for a job continuously exceeds the threshold (**metrics.reporter.alarm.job.alarm.rocksdb.get.micros.threshold**, 50000 microseconds by default). This alarm is cleared when the MemTable size of RocksDB for the job is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45643	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

The write pressure of RocksDB is high.

Handling Procedure

Check TaskManager logs for the write pressure of RocksDB and collect logs.

Step 1 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 2 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 3 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 4 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-164 Application ID of a job

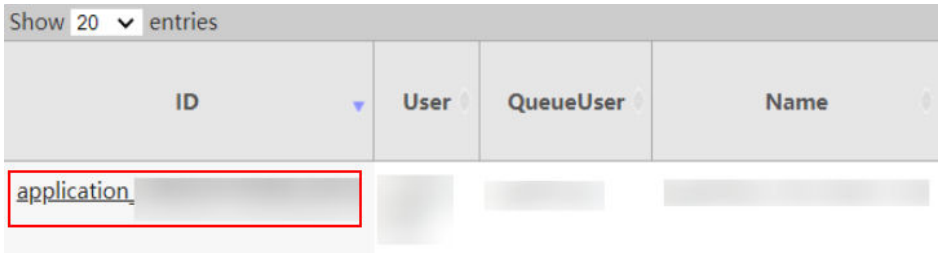

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-165 Clicking Logs

   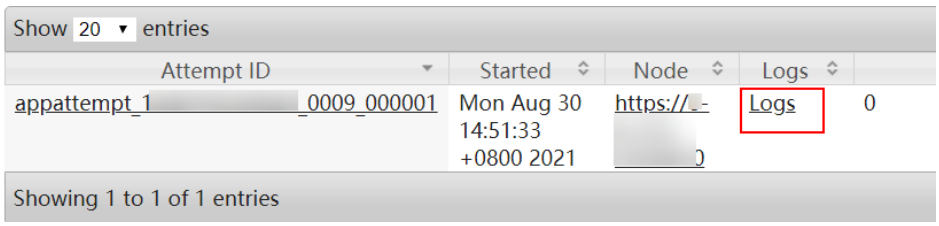

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 7**.

Figure 10-166 Clicking the ID in the Attempt ID column

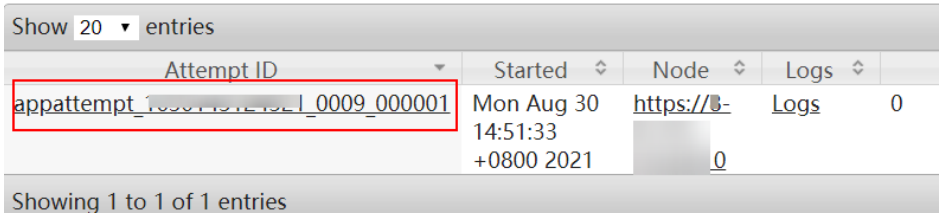

Figure 10-167 Clicking Logs

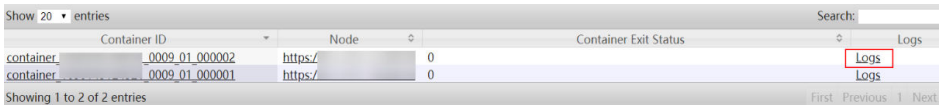

NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Check whether the write pressure of RocksDB is high.

Step 7 Check whether the value of **rocksdb.size-all-mem-tables** (unit: byte) in the TaskManager monitoring logs (keyword **RocksDBMetricPrint**) is greater than or equal to the total write buffer size (Total write buffer = **write_buffer_size** x **max_write_buffer_number**).

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 8**.

Table 10-7 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.writebuffer.count	– **2** – **4**: enables **SPINNING_DISK_OPTIMIZED_HIGH_MEM**.	– Number of buffers – **2** to **10** are recommended. Adjust the value based on service requirements.
state.backend.rocksdb.writebuffer.size	**64MB**	– Buffer size – **64MB** to **256MB** are recommended.

Parameter	Default Value	Description
state.backend.rocksdb.thread.num	– **2** – **4**: enables **SPINNING_DISK_OPTIMIZED_HIGH_MEM**.	– Number of flush threads. Increase the number of threads to quickly flush memory data to disks. – When the number of threads is increased, the number of vCores also needs to be increased. – **2** to **10** are recommended.

- If no, go to **Step 9**.

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.372 ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the number of SST files at level 0 of RocksDB for a job continuously exceeds the threshold (**state.backend.rocksdb.level0_slowdown_writes_trigger**, 20 by default). This alarm is cleared when the number of SST files at level 0 of RocksDB for the job is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45644	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

Possible causes are as follows:

- The compaction pressure of RocksDB is too high, and **ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** are generated.
- There are too many SST files at level 0.

Handling Procedure

Check whether the compaction pressure of RocksDB is too high and ALM-45646 is generated.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether **ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by following the instructions provided in section **ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45646 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the compaction pressure of RocksDB is too high and ALM-45647 is generated.

Step 5 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 6**.
- If no, go to **Step 8**.

Step 6 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 7 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Check TaskManager logs for the number of SST files at level 0 and collect logs.

Step 8 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 9 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 10 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 11 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-168 Application ID of a job

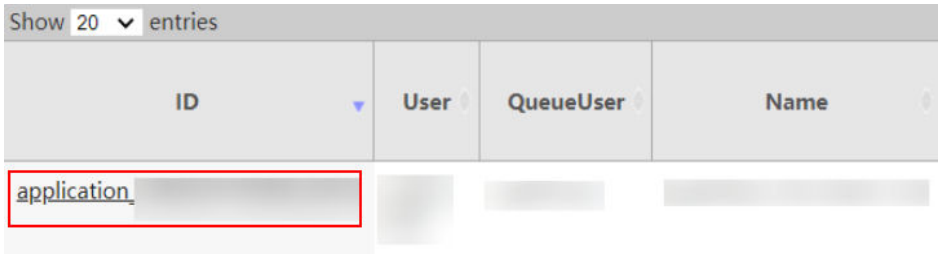

- If yes, go to **Step 12**.

- If no, go to **Step 13**.

Step 12 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-169 Clicking Logs

   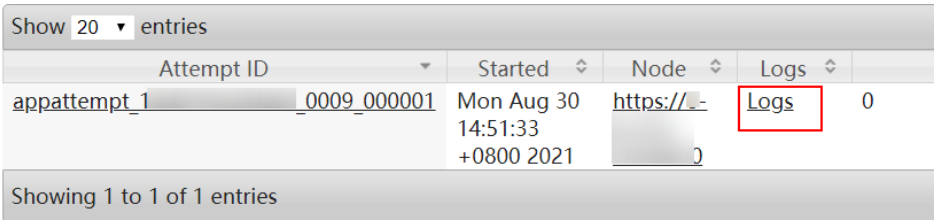

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 14**.

   Figure 10-170 Clicking the ID in the Attempt ID column

   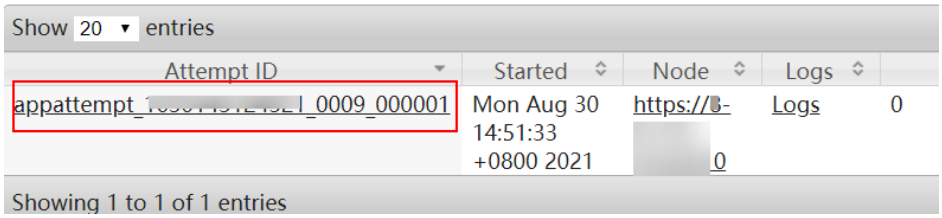

   Figure 10-171 Clicking Logs

   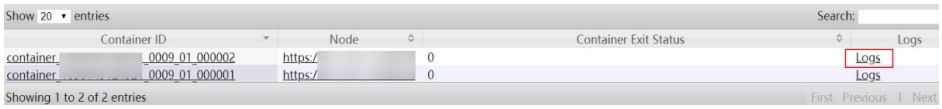

   NOTE

   You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 13 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Check whether the number of SST files at level 0 is too large.

Step 14 Check whether the value of **rocksdb.num-files-at-level0** in TaskManager monitoring logs (keyword **RocksDBMetricPrint**) is greater than or equal to the value of **state.backend.rocksdb.level0_slowdown_writes_trigger** or **state.backend.rocksdb.level0_stop_writes_trigger**.

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 15**.

Table 10-8 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.level0_slowdown_writes_trigger	20	– Number of files that trigger slowdown at level 0 – **20** to **30** are recommended.
state.backend.rocksdb.level0_stop_writes_trigger	36	– Maximum number of files that trigger stop at level 0 – **36** to **46** are recommended.

- If no, go to **Step 16**.

Step 15 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 16**.

Step 16 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.373 ALM-45645 Pending Flush Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the number of pending flush requests of RocksDB for a job continuously reaches *n* times the number of flush/compaction threads. This alarm is cleared when the number of pending flush requests of RocksDB for the job is less than or equal to the threshold.

- The number of flush/compaction threads is the value of **state.backend.rocksdb.thread.num**. The default value is **2**. If **SPINNING_DISK_OPTIMIZED_HIGH_MEM** is enabled, the default value is **4**.
- The **metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier** parameter specifies *n* times the number of flush/compaction threads. The default value is **2**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45645	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

The number of pending flush requests of RocksDB for the Flink job is too large.

Handling Procedure

Check TaskManager logs for the number of pending flush requests and collect logs.

Step 1 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 2 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45645 Pending Flush Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 3 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 4 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-172 Application ID of a job

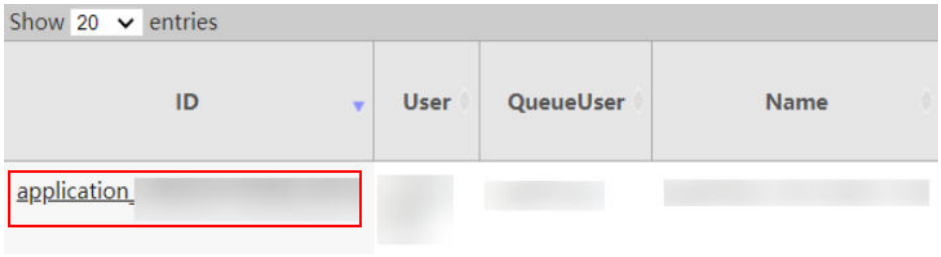

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-173 Clicking Logs

   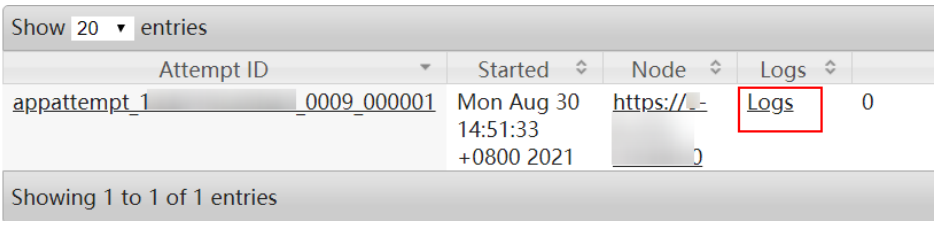

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 7**.

   Figure 10-174 Clicking the ID in the Attempt ID column

   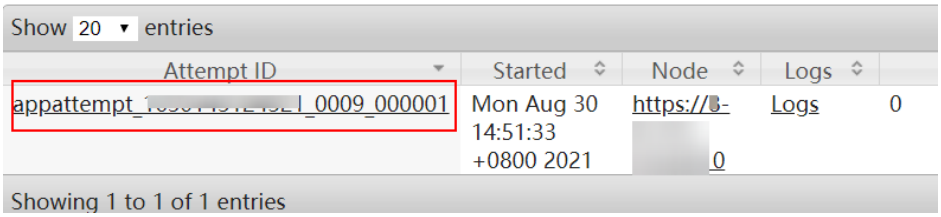

   Figure 10-175 Clicking Logs

   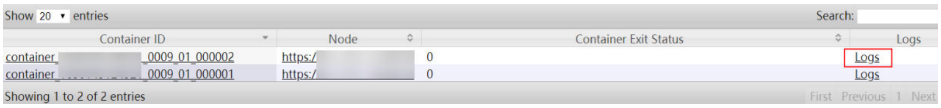

   NOTE

   You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and

download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Check whether there are too many pending flush requests.

Step 7 Check whether the sum of the values of **rocksdb.mem-table-flush-pending** and **rocksdb.compaction-pending** in TaskManager monitoring logs (keyword **RocksDBMetricPrint**) is greater than *n* times the number of RocksDB threads (**metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier**, 2 by default). If it is, you can increase the number of RocksDB threads.

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 8**.

  Table 10-9 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.thread.num	– **2** – **4**: enables **SPINNING_DISK_OPTIMIZED_HIGH_MEM**.	– Number of flush threads. Increase the number of threads to quickly flush memory data to disks. – When the number of threads is increased, the number of vCores also needs to be increased. – **2** to **10** are recommended.

- If no, go to **Step 9**.

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.374 ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the number of pending compaction requests of RocksDB for a job continuously reaches *n* times the number of flush/compaction threads. This alarm is cleared when the number of pending compaction requests of RocksDB for the job is less than or equal to the threshold.

- The number of flush/compaction threads is the value of **state.backend.rocksdb.thread.num**. The default value is **2**. If **SPINNING_DISK_OPTIMIZED_HIGH_MEM** is enabled, the default value is **4**.
- The **metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier** parameter specifies *n* times the number of flush/compaction threads. The default value is **2**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45646	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

The number of pending compaction requests of RocksDB for the Flink job is too large.

Handling Procedure

Check TaskManager logs for the number of pending compaction requests and collect logs.

Step 1 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 2 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45646 Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 3 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 4 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-176 Application ID of a job

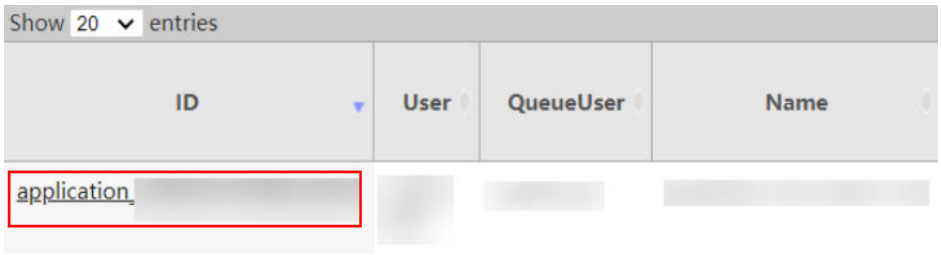

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-177 Clicking Logs

   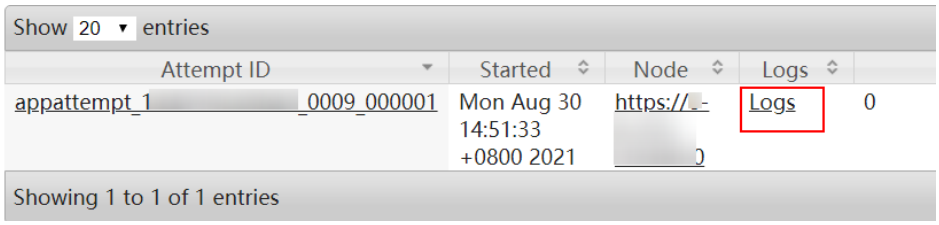

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 7**.

   Figure 10-178 Clicking the ID in the Attempt ID column

   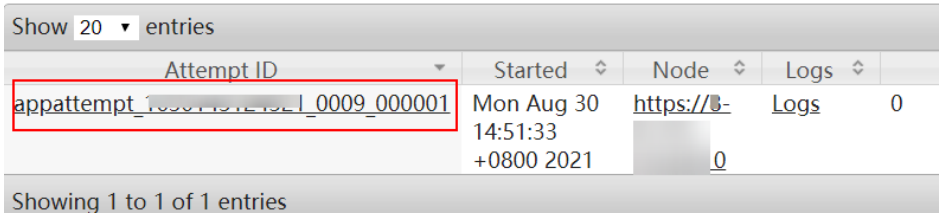

Figure 10-179 Clicking Logs

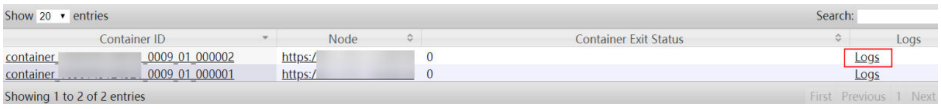

NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Check whether there are too many pending compaction requests.

Step 7 Check whether the sum of the values of **rocksdb.mem-table-flush-pending** and **rocksdb.compaction-pending** in TaskManager monitoring logs (keyword **RocksDBMetricPrint**) is greater than **n** times the number of RocksDB threads (**metrics.reporter.alarm.job.alarm.rocksdb.background.jobs.multiplier**, 2 by default). If it is, you can increase the number of RocksDB threads.

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 8**.

  Table 10-10 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.thread.num	– **2** – **4**: enables **SPINNING_DISK_OPTIMIZED_HIGH_MEM**.	– Number of flush threads. Increase the number of threads to quickly flush memory data to disks. – When the number of threads is increased, the number of vCores also needs to be increased. – **2** to **10** are recommended.

- If no, go to **Step 9**.

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.375 ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the estimated pending compaction size of RocksDB for a job continuously exceeds the threshold. This alarm is cleared when the estimated pending compaction size of RocksDB for the job is less than or equal to the threshold.

The threshold of the estimated pending compaction size is the smaller value of the following two parameters:

- **state.backend.rocksdb.soft-pending-compaction-bytes-limit**. The default value is **64GB**.
- **state.backend.rocksdb.hard-pending-compaction-bytes-limit**. The default value is **256GB**.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45647	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

The estimated compaction data size of RocksDB is too large.

Handling Procedure

Check TaskManager logs for the estimated compaction data size and collect logs.

Step 1 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 2 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 3 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 4 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-180 Application ID of a job

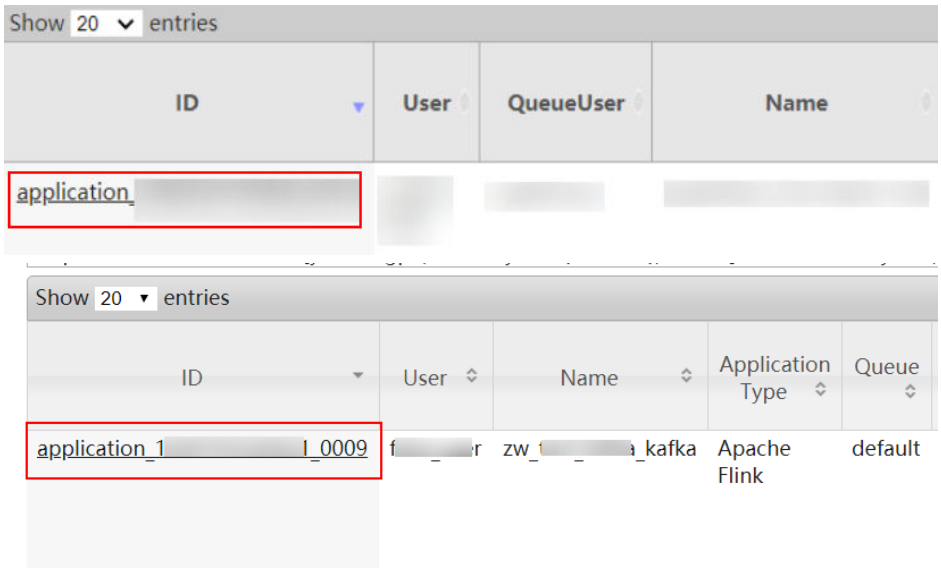

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

   Figure 10-181 Clicking Logs

   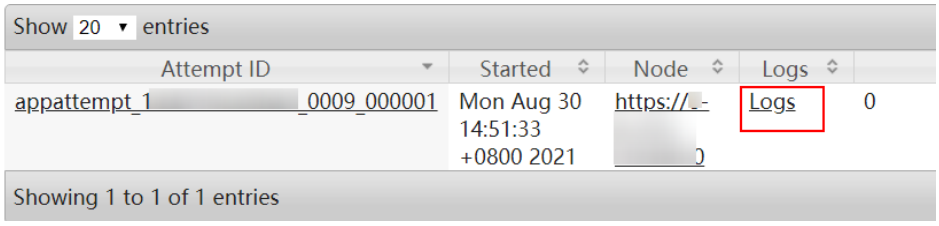

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 7**.

   Figure 10-182 Clicking the ID in the Attempt ID column

   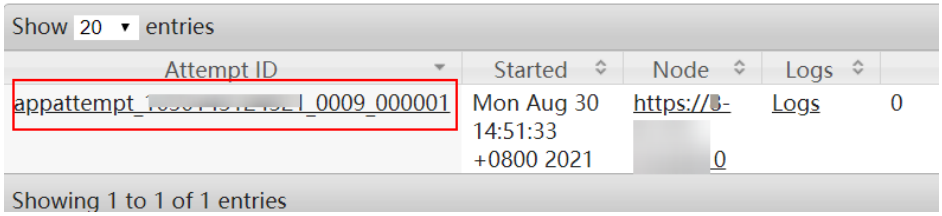

   Figure 10-183 Clicking Logs

   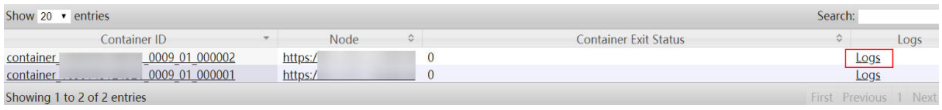

   NOTE

   You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 6 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Check whether the estimated compaction data size of RocksDB is too large.

Step 7 Check whether the value of **rocksdb.estimate-pending-compaction-bytes** (unit: byte) in TaskManager monitoring logs (keyword **RocksDBMetricPrint**) is greater than or equal to the **soft/hard-pending-compaction** size (values of **state.backend.rocksdb.soft-pending-compaction-bytes-limit** and **state.backend.rocksdb.hard-pending-compaction-bytes-limit**).

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 8**.

  Table 10-11 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.soft-pending-compaction-bytes-limit	64GB	– When the pending compaction size exceeds the threshold, the write traffic is limited. – **64GB** to **512GB** are recommended.
state.backend.rocksdb.hard-pending-compaction-bytes-limit	256GB	– When the pending compaction size exceeds the threshold, write operations are stopped. – **64GB** to **512GB** are recommended.

- If no, go to **Step 9**.

Step 8 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Step 9 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.376 ALM-45648 RocksDB Frequently Encounters Write-Stopped

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when RocksDB for a job continuously encounters the **is-write-stopped** state. This alarm is cleared when RocksDB for the job no longer or

does not continuously encounter the **is-write-stopped** state within an alarm reporting interval.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45648	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

The possible causes are as follows:

- There are too many MemTables and **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** is generated.
- There are too many SST files at level 0, and **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The estimated compaction size exceeds the threshold, and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** is generated.

Handling Procedure

Check whether there are too many MemTables.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by following the instructions provided in section **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45643 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the number of SST files at level 0 is too large.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 6 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 8 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether the estimated compaction size exceeds the threshold.

Step 9 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 10**.
- If no, go to **Step 12**.

Step 10 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 11 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Collect fault information.

Step 12 Log in to Manager as a user who has the management permission for the current Flink job.

Step 13 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45648 RocksDB Frequently Encounters Write-Stopped**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 14 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 15 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-184 Application ID of a job

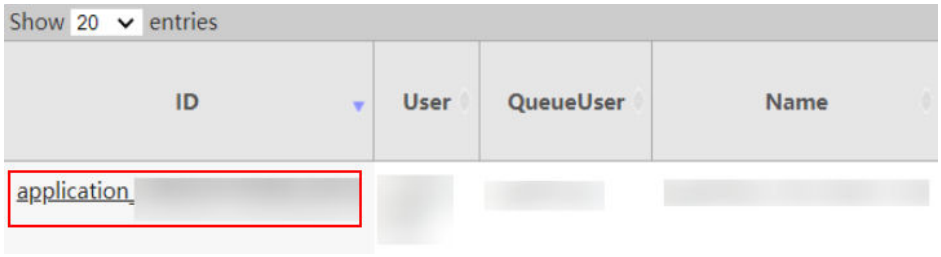

- If yes, go to **Step 16**.
- If no, go to **Step 18**.

Step 16 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-185 Clicking Logs

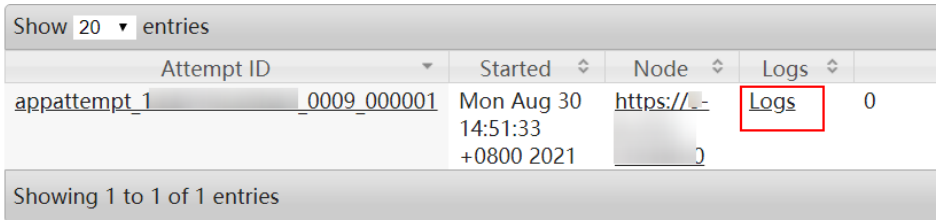

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs.

Figure 10-186 Clicking the ID in the Attempt ID column

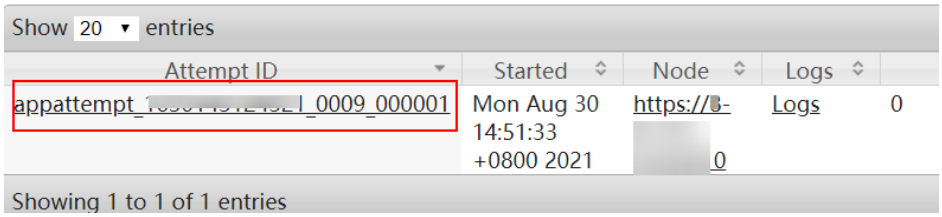

Figure 10-187 Clicking Logs

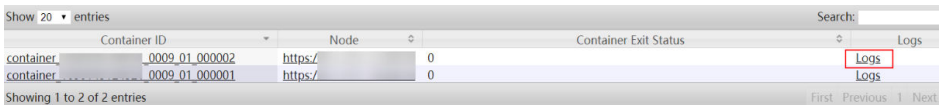

NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 17 View the job logs to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 18 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Step 19 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.377 ALM-45649 P95 Latency of RocksDB Get Requests Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the P95 latency of RocksDB Get requests exceeds the threshold (**metrics.reporter.alarm.job.alarm.rocksdb.get.micros.threshold**, 50000 microseconds by default). This alarm is cleared when the P95 latency of RocksDB Get requests is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45649	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

The possible causes are as follows:

- There are too many SST files at level 0, causing slow queries. In addition, **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The cache hit ratio is lower than 60%, causing frequent swap-ins and swap-outs of the block cache.

Handling Procedure

Check whether the number of SST files at level 0 is too large.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check the cache hit ratio in TaskManager logs and collect logs.

Step 5 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 6 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45649 P95 Latency of RocksDB Get Requests Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 7 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 8 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-188 Application ID of a job

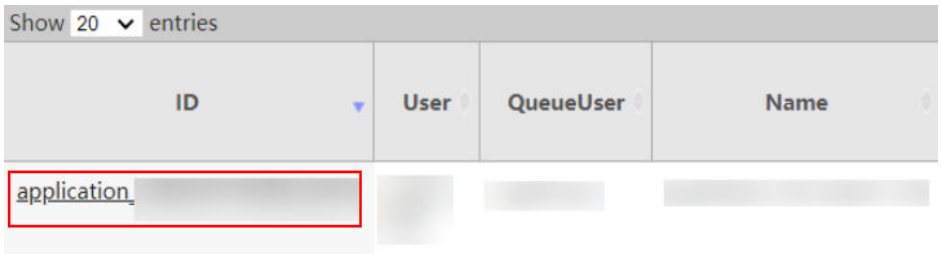

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-189 Clicking Logs

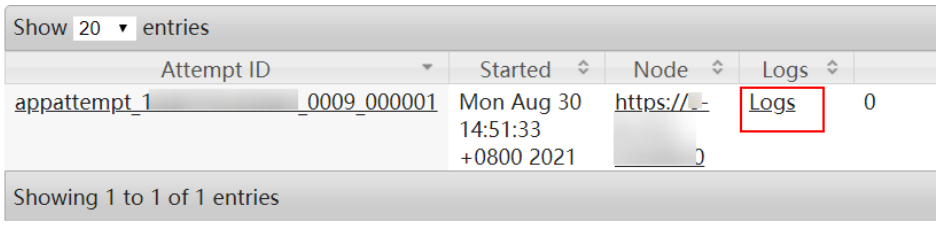

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs. Then go to **Step 11**.

Figure 10-190 Clicking the ID in the Attempt ID column

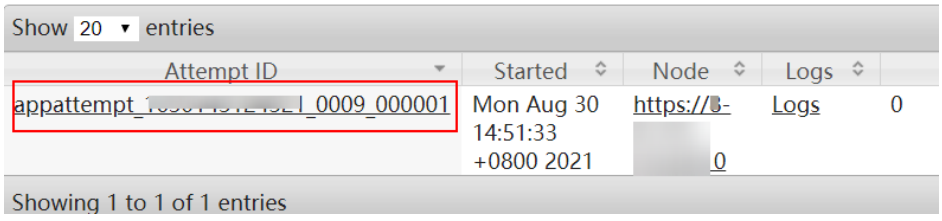

Figure 10-191 Clicking Logs

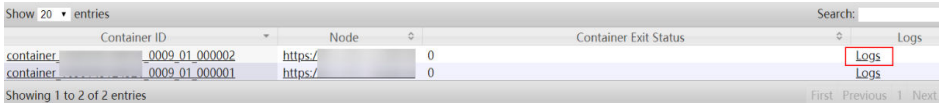

NOTE

You can also log in to Manager as a user who has the management permission for the current Flink job. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 10 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Check whether the cache hit ratio is too low.

Step 11 Check the values of **rocksdb.block.cache.hit** (cache hit) and **rocksdb.block.cache.miss** (cache miss) in TaskManager monitoring logs (keyword **RocksDBMetricPrint**). Calculate the hit ratio using the following formula and check whether it is less than 60%:

rocksdb.block.cache.hit/(rocksdb.block.cache.hit+rocksdb.block.cache.miss)

- If yes, adjust the values of the following custom parameters on the job development page of the Flink web UI, save the settings, and go to **Step 12**.

  Table 10-12 Custom parameters

Parameter	Default Value	Description
state.backend.rocksdb.block.cache-size	– **8MB** – **256MB**: enables **SPINNING_DISK_OPTIMIZED_HIGH_MEM**.	– Cache size – **8MB** to **1GB** are recommended.
state.backend.rocksdb.block.blocksize	– **4KB** – **128KB**: enables **SPINNING_DISK_OPTIMIZED_HIGH_MEM**.	– Block size – **4KB** to **256KB** are recommended.
state.backend.rocksdb.use-bloom-filter	**false**	– Whether to speed up indexing. If it is **true**, each new SST file will contain a Bloom filter. – **true** is recommended.

- If no, go to **Step 13**.

Step 12 Restart the job and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Step 13 Contact O&M personnel and send the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.378 ALM-45650 P95 Latency of RocksDB Write Requests Continuously Exceeds the Threshold

This section applies to MRS 3.3.0 or later.

Alarm Description

The system checks the RocksDB monitoring data of jobs at the user-specified alarm reporting interval (**metrics.reporter.alarm.job.alarm.rocksdb.metrics.duration**, 180s by default). This alarm is generated when the P95 latency of RocksDB write requests exceeds the threshold (**metrics.reporter.alarm.job.alarm.rocksdb.write.micros.threshold**, 50000 microseconds by default). This alarm is cleared when the P95 latency of RocksDB write requests is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45650	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
ApplicationName	Specifies the name of the application for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
JobName	Specifies the job for which the alarm is generated.

Impact on the System

This alarm has no adverse impact on the system.

Possible Causes

The possible causes are as follows:

- There are too many MemTables. As a result, write traffic is limited or write stops, and **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** is generated.
- There are too many SST files at level 0, and **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** is generated.
- The estimated compaction size exceeds the threshold, and **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** is generated.

Handling Procedure

Check whether write traffic limiting or write stop is caused due to too many MemTables.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by following the instructions provided in section **ALM-45643 MemTable Size of RocksDB Continuously Exceeds the Threshold**.

Step 4 After ALM-45643 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the number of SST files at level 0 is too large.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 6 In the alarm list, check whether **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Handle the alarm by following the instructions provided in section **ALM-45644 Number of SST Files at Level 0 of RocksDB Continuously Exceeds the Threshold**.

Step 8 After ALM-45644 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether the estimated compaction size exceeds the threshold.

Step 9 In the alarm list, check whether **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold** exists.

- If yes, go to **Step 10**.
- If no, go to **Step 12**.

Step 10 Handle the alarm by following the instructions provided in section **ALM-45647 Estimated Pending Compaction Size of RocksDB Continuously Exceeds the Threshold**.

Step 11 After ALM-45647 is cleared, wait a few minutes and check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 12**.

Collect fault information.

Step 12 Log in to FusionInsight Manager as a user who has the FlinkServer management permission.

Step 13 Choose **O&M** > **Alarm** > **Alarms** > **ALM-45650 P95 Latency of RocksDB Write Requests Continuously Exceeds the Threshold**, view **Location**, and obtain the name of the task for which the alarm is generated.

Step 14 Choose **Cluster** > **Services** > **Yarn** and click the link next to **ResourceManager WebUI** to go to the native Yarn page.

Step 15 Locate the abnormal task based on its name displayed in **Location**, search for and record the application ID of the job, and check whether the job logs are available on the Yarn page.

Figure 10-192 Application ID of a job

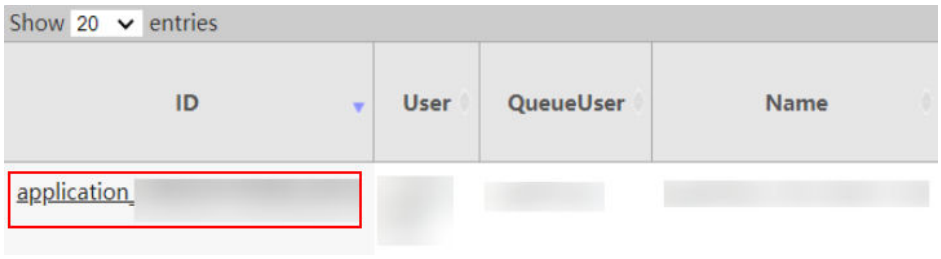

- If yes, go to **Step 16**.
- If no, go to **Step 18**.

Step 16 Click the application ID of the failed job to go to the job page.

1. Click **Logs** in the **Logs** column to view JobManager logs.

Figure 10-193 Clicking Logs

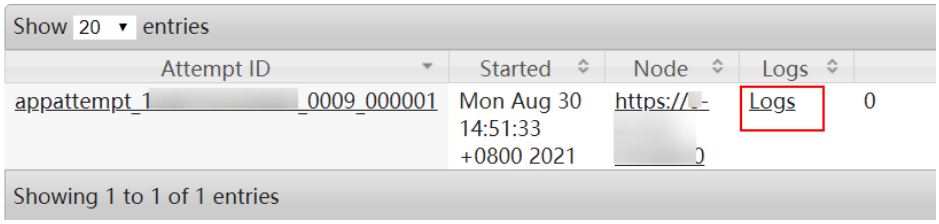

2. Click the ID in the **Attempt ID** column and click **Logs** in the **Logs** column to view and save TaskManager logs.

Figure 10-194 Clicking the ID in the Attempt ID column

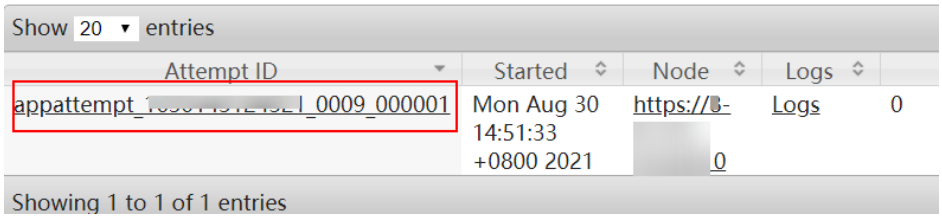

Figure 10-195 Clicking Logs

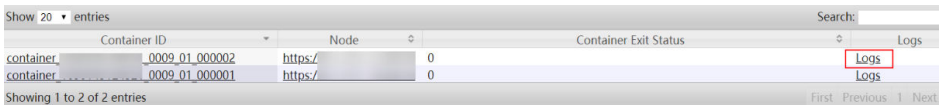

NOTE

You can also log in to Manager as a user who has the FlinkServer management permission. Choose **Cluster** > **Services** > **Flink**, and click the link next to **Flink WebUI**. On the displayed Flink web UI, click **Job Management**, click **More** in the **Operation** column, and select **Job Monitoring** to view TaskManager logs.

Step 17 View the job logs to rectify the fault, or contact the O&M personnel and send the collected fault logs. No further action is required.

If logs are unavailable on the Yarn page, download logs from HDFS.

Step 18 On Manager, choose **Cluster** > **Services** > **HDFS**, click the link next to **NameNode WebUI** to go to the HDFS page, choose **Utilities** > **Browse the file system**, and download logs in the **/tmp/logs/***Username*/**bucket-logs-tfile/***Last four digits of the task application ID/Application ID of the task* directory.

Step 19 View the logs of the failed job to rectify the fault, or contact the O&M personnel and send the collected fault logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.379 ALM-45652 Flink Service Unavailable

This section applies to MRS 3.3.0 or later.

Alarm Description

The alarm module checks the Flink status every 60 seconds. This alarm is generated when the Flink service is unavailable. This alarm is cleared when the Flink service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45652	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the job for which the alarm is generated.

Impact on the System

FlinkServer and the Flink client cannot be used to submit Flink jobs.

Possible Causes

The ZooKeeper, HDFS, Yarn, KrbServer, or DBService service on which Flink depends is unavailable.

Handling Procedure

Check whether the ZooKeeper service on which Flink depends is abnormal.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**.

Step 2 In the alarm list, check whether "ALM-13000 ZooKeeper Service Unavailable" exists.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Handle the alarm by referring to "ALM-13000 ZooKeeper Service Unavailable."

Step 4 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the HDFS service on which Flink depends is abnormal.

Step 5 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 6 In the alarm list, check whether "ALM-14000 HDFS Service Unavailable" exists.

- If yes, go to **Step 7**.
- If no, go to **Step 9**.

Step 7 Handle the alarm by referring to "ALM-14000 HDFS Service Unavailable."

Step 8 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Check whether the Yarn service on which Flink depends is abnormal.

Step 9 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 10 In the alarm list, check whether "ALM-18000 Yarn Service Unavailable" exists.

- If yes, go to **Step 11**.
- If no, go to **Step 13**.

Step 11 Handle the alarm by referring to "ALM-18000 Yarn Service Unavailable."

Step 12 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 13**.

Check whether the KrbServer service on which Flink depends is abnormal.

Step 13 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 14 In the alarm list, check whether "ALM-25500 KrbServer Service Unavailable" exists.

- If yes, go to **Step 15**.
- If no, go to **Step 17**.

Step 15 Handle the alarm by referring to "ALM-25500 KrbServer Service Unavailable."

Step 16 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.

- If no, go to **Step 17**.

Check whether the DBService service on which Flink depends is abnormal.

Step 17 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**.

Step 18 In the alarm list, check whether "ALM-27001 DBService Service Unavailable" exists.

- If yes, go to **Step 19**.
- If no, go to **Step 21**.

Step 19 Handle the alarm by referring to "ALM-27001 DBService Service Unavailable."

Step 20 After the alarm is cleared, wait a few minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 21**.

Collect fault information.

Step 21 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 22 Expand the **Service** drop-down list, and select **Flink** for the target cluster.

Step 23 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time respectively. Then, click **Download**.

Step 24 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.380 ALM-45653 Invalid Flink HA Certificate File

This section applies to MRS 3.3.0 or later.

Alarm Description

Flink checks whether the HA certificate file is valid (whether the certificate exists and whether its format is correct) in the first health check or at 01:00:00 every day. This alarm is generated when the certificate file is invalid. This alarm is automatically cleared when the certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45653	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Some functions are unavailable.

Possible Causes

The HA certificate file is invalid.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45653 Invalid Flink HA Certificate File**, view **Location**, obtain the name of the host for which the alarm is generated, and click the host name to view its IP address.

Check whether the HA certificate file in the system is valid.

Step 2 Log in to the host for which the alarm is generated as user **omm**.

Step 3 Run the **cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/ha/local/cert** command to go to the directory where the HA certificate is stored.

Step 4 Run the **ls -l** command to check whether the **server.crt** file exists.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Run the **openssl x509 -in server.crt -text -noout** command and check whether the command output is normal.

- If yes, go to **Step 9**.
- If no, go to **Step 6**.

Step 6 Run the **cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/flink/sbin** command to go to the Flink script directory.

Step 7 Run the **sh proceed_ha_ssl_cert.sh** command to generate a new HA certificate. Then, check whether the alarm is cleared 1 minute later.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Flink** for the target cluster.

Step 11 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.381 ALM-45654 Flink HA Certificate Is About to Expire

This section applies to MRS 3.3.0 or later.

Alarm Description

Flink checks whether the HA certificate file is about to expire in the first health check or at 01:00:00 every day. This alarm is generated when the remaining validity period is less than or equal to 30 days. This alarm is automatically cleared when the remaining validity period is greater than 30 days.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45654	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Currently, there is no impact on the system.

Possible Causes

The HA certificate is about to expire.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45654 Flink HA Certificate Is About to Expire**, view **Location**, obtain the name of the host for which the alarm is generated, and click the host name to view its IP address.

Check whether the HA certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the host for which the alarm is generated as user **omm**.

Step 3 Run the **cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/ha/local/cert** command to go to the directory where the HA certificate is stored.

Step 4 Run the **openssl x509 -noout -text -in server.crt** command to query the effective time and due time of the HA certificate.

Step 5 Perform **Step 6** to **Step 7** during off-peak hours to update the certificate file as needed.

Step 6 Run the **cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/flink/sbin** command to go to the Flink script directory.

Step 7 Run the **sh proceed_ha_ssl_cert.sh** command to generate a new HA certificate. Then, check whether the alarm is cleared 1 minute later.

- If yes, go to **Step 9**.
- If no, go to **Step 8**.

Step 8 On the node where the standby FlinkServer instance is located, repeat **Step 6** to **Step 7**. Then, check whether the alarm is cleared 1 minute later.

- If yes, go to **Step 9**.
- If no, go to **Step 10**.

Step 9 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 10**.
- If no, no further action is required.

Collect fault information.

Step 10 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 11 Expand the **Service** drop-down list, and select **Flink** for the target cluster.

Step 12 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 13 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.382 ALM-45655 Flink HA Certificate File Has Expired

This section applies to MRS 3.3.0 or later.

Alarm Description

Flink checks whether the HA certificate file has expired in the first health check or at 01:00:00 every day. This alarm is generated when the HA certificate has expired. This alarm is automatically cleared when the certificate file becomes valid again.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45655	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

Some functions are unavailable.

Possible Causes

The HA certificate file has expired.

Handling Procedure

View alarm information.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** > **ALM-45655 Flink HA Certificate File Has Expired**, view **Location**, obtain the name of the host for which the alarm is generated, and click the host name to view its IP address.

Check whether the HA certificate file in the system is valid. If it is not, generate a new one.

Step 2 Log in to the host for which the alarm is generated as user **omm**.

Step 3 Run the **cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/ha/local/cert** command to go to the directory where the HA certificate is stored.

Step 4 Run the **openssl x509 -noout -text -in server.crt** command to query the effective time and due time of the HA certificate and check whether the HA certificate file is valid.

- If yes, go to **Step 9**.
- If no, go to **Step 5**.

Step 5 Run the **cd ${BIGDATA_HOME}/FusionInsight_Flink_*/install/FusionInsight-Flink-*/flink/sbin** command to go to the Flink script directory.

Step 6 Run the **sh proceed_ha_ssl_cert.sh** command to generate a new HA certificate. Then, check whether the alarm is cleared 1 minute later.

- If yes, go to **Step 8**.
- If no, go to **Step 7**.

Step 7 On the node where the standby FlinkServer instance is located, repeat **Step 5** to **Step 6**. Then, check whether the alarm is cleared 1 minute later.

- If yes, go to **Step 8**.
- If no, go to **Step 9**.

Step 8 Check whether this alarm is generated again during periodic system check.

- If yes, go to **Step 9**.
- If no, no further action is required.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Flink** for the target cluster.

Step 11 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.383 ALM-45736 Guardian Service Unavailable

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

The alarm module checks the Guardian service status every 60 seconds. This alarm is generated if Guardian is unavailable.

This alarm is cleared after Guardian recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45736	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Guardian cannot work properly.

Possible Causes

- The HDFS service on which the Guardian service depends is abnormal.
- The TokenServer role instance is abnormal.

Handling Procedure

Check the HDFS service status.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. On the page that is displayed, check whether "ALM-14000 HDFS Service Unavailable" is reported.

- If yes, go to **Step 2**.
- If no, go to **Step 3**.

Step 2 Clear this alarm according to the alarm help.

After the alarm is cleared, wait a few minutes and check whether the alarm GuardianService Unavailable is cleared.

- If yes, no further action is required.
- If no, go to **Step 3**.

Check all TokenServer instances.

Step 3 Log in to the node where the TokenServer instance resides as user **omm** and run the **ps -ef|grep "guardian.token.server.Server"** command to check whether the TokenServer process exists on the node.

- If yes, go to **Step 5**.
- If no, restart the faulty TokenServer instance and go to **Step 4**.

Step 4 In the alarm list, check whether the alarm "Guardian Service Unavailable" is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.

Step 7 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.384 ALM-45737 TokenServer Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the TokenServer service every 60 seconds. This alarm is generated when the heap memory usage of the TokenServer instance exceeds the threshold (95% of the maximum memory) for 10 consecutive times.

This alarm is automatically cleared when the system detects that the heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45737	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Heap memory overflow may cause service unavailability.

Possible Causes

The heap memory of the TokenServer instance is overused or the heap memory is inappropriately allocated.

Handling Procedure

Check heap memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45737 TokenServer Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **TokenServer Heap Memory Usage**. Then click **OK**.

Step 3 Check whether the heap memory used by TokenServer reaches the threshold (95% of the maximum heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer** >

Instance Configuration. Click **All Configurations**, and choose **TokenServer** > **System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for TokenServer cannot meet the heap memory required by the TokenServer process. You are advised to change the value of **-Xmx** in **GC_OPTS** to twice that of the heap memory used by TokenServer. You can change the value based on the actual service scenario. Refer to **Step 2** to view the TokenServer heap memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.385 ALM-45738 TokenServer Direct Memory Usage Exceeds the Threshold

Alarm Description

The system checks the direct memory usage of the TokenServer service every 60 seconds. This alarm is generated when the direct memory usage of the TokenServer instance exceeds the threshold (80% of the maximum memory) for five consecutive times.

This alarm is automatically cleared when the system detects that the TokenServer direct memory usage is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45738	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Direct memory overflow may cause service unavailability.

Possible Causes

The direct memory of the TokenServer process is overused or the direct memory is inappropriately allocated.

Handling Procedure

Check the direct memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45738 TokenServer Direct Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **TokenServer Direct Memory Usage**. Then click **OK**.

Step 3 Check whether the direct memory used by TokenServer reaches the threshold (80% of the maximum direct memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer** > **Instance Configuration**. Click **All Configurations**, and choose **TokenServer** > **System**. Set **-XX:MaxDirectMemorySize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the direct memory configured for TokenServer cannot meet the direct memory required by the TokenServer process. You are advised to check the direct memory usage of TokenServer and change the value of **-XX:MaxDirectMemorySize** in **GC_OPTS** to the twice of the direct memory used by TokenServer. You can change the value based on the actual service scenario. Refer to **Step 2** to view the TokenServer direct memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.386 ALM-45739 TokenServer Non-Heap Memory Usage Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the TokenServer service every 60 seconds. This alarm is generated when the non-heap memory usage of the TokenServer instance exceeds the threshold (80% of the maximum memory) for five consecutive times.

This alarm is automatically cleared when the system detects that the non-heap memory usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45739	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Non-heap memory overflow may cause service unavailability.

Possible Causes

The non-heap memory of the TokenServer instance is overused or the non-heap memory is inappropriately allocated.

Handling Procedure

Check non-heap memory usage.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45739 TokenServer Non-Heap Memory Usage Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated. Click the drop-down list in the upper right corner of the chart area and choose **Customize** > **CPU and Memory** > **TokenServer Non-Heap Memory Usage**. Then click **OK**.

Step 3 Check whether the non-heap memory used by TokenServer reaches the threshold (80% of the maximum non-heap memory by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer** > **Instance Configuration**. Click **All Configurations**, and choose **TokenServer** > **System**. Set **-XX: MaxPermSize** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the non-heap memory size configured for the TokenServer instance cannot meet the non-heap memory required by the TokenServer process. You are advised to change the value of **-XX:MaxPermSize** in **GC_OPTS** to twice that of the current non-heap memory size or change the value based on site requirements.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.

Step 8 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.387 ALM-45740 TokenServer GC Duration Exceeds the Threshold

Alarm Description

The system checks the GC duration of the TokenServer process every 60 seconds. This alarm is generated when the GC duration of the TokenServer process exceeds the threshold (12 seconds by default) for five consecutive times.

This alarm is automatically cleared when the system detects that the GC duration is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45740	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

TokenServer responds slowly.

Possible Causes

The heap memory of the TokenServer process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms** > **ALM-45740 TokenServer GC Duration Exceeds the Threshold**. Check the location information of the alarm and view the host name of the instance for which the alarm is generated.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, select the role corresponding to the host name of the instance for which the alarm is generated and click the drop-down list in the upper right corner of the chart area. Choose **Customize** > **GC** > **TokenServer GC Duration**. Then click **OK**.

Step 3 Check whether the GC duration of the TokenServer process collected every minute exceeds the threshold (12 seconds by default).

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 On FusionInsight Manager, choose **Cluster** > **Services** > **Guardian**. On the page that is displayed, click the **Instance** tab. On this tab page, choose **TokenServer** > **Instance Configuration**. Click **All Configurations**, and choose **TokenServer** > **System**. Set **-Xmx** in the **GC_OPTS** parameter to a larger value based on site requirements and save the configuration.

NOTE

If this alarm is generated, the heap memory configured for TokenServer cannot meet the heap memory required by the TokenServer process. You are advised to change the value of **-Xmx** in **GC_OPTS** to twice that of the heap memory used by TokenServer. You can change the value based on the actual service scenario. Refer to **Step 2** to view the TokenServer heap memory usage.

Step 5 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M** > **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **Guardian** for the destination cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.388 ALM-45741 Failed to Call the ECS securitykey API

NOTE

This section applies only to MRS 3.2.1 or later.

Alarm Description

Guardian caches the temporary AK/SK of the ECS agency. When the cache does not exist or is about to expire, Guardian calls the securitykey API of ECS to update the AK/SK. This alarm is generated when calling to the API fails.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45741	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The task may fail to obtain the temporary AK/SK for accessing OBS. As a result, the task cannot access OBS.

Possible Causes

- No ECS agency is bound to the cluster.
- An underlying interface of ECS is abnormal.

Handling Procedure

Check whether an agency is bound to the cluster.

Step 1 Log in to the MRS console.

Step 2 In the navigation pane on the left, choose **Clusters** > **Active Clusters**. On the page that is displayed, click the cluster name to go to its overview page. Then, check whether the cluster is bound to an agency in the O&M management area.

- If yes, go to **4**.
- If no, go to **3**.

Step 3 Click **Manage Agency**. On the page that is displayed, rebind the cluster to an agency. Then check whether the alarm is cleared a few minutes later.

- If yes, no further action is required.
- If no, go to **4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Guardian** for the target cluster.

Step 6 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.389 ALM-45742 Failed to Call the ECS Metadata API

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

When Guardian calls an IAM API to obtain the temporary AK/SK, it needs to first obtain related metadata via the ECS Metadata API. This alarm is generated when Guardian fails to call the Metadata API.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45742	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The task may fail to obtain the temporary AK/SK for accessing OBS. As a result, the task cannot access OBS.

Possible Causes

An underlying interface of ECS is abnormal.

Handling Procedure

Collect fault information.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 2 Expand the **Service** drop-down list, and select **Guardian** for the target cluster.

Step 3 Click in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 4 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.390 ALM-45743 Failed to Call the IAM API

NOTE

This section applies only to MRS 3.1.5 or later.

Alarm Description

This alarm is generated when Guardian fails to call the IAM API to obtain a temporary AK/SK.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
45743	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

The task may fail to obtain the temporary AK/SK for accessing OBS. As a result, the task cannot access OBS.

Possible Causes

The IAM service is abnormal.

Handling Procedure

Collect fault information.

Step 1 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 2 Expand the **Service** drop-down list, and select **Guardian** for the target cluster.

Step 3 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 4 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None

10.391 ALM-50201 Doris Service Unavailable

Alarm Description

The alarm module checks the Doris service status every 60 seconds. This alarm is generated when the alarm module detects that all FE and BE instances are abnormal.

This alarm is cleared when any FE or BE instance recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50201	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

FusionInsight Manager cannot be used to perform cluster operations on the Doris service, and Doris service functions are unavailable.

Possible Causes

The FE and BE instances are abnormal.

Handling Procedure

Restart the Doris service.

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Doris**.

Step 2 On the page that is displayed, click **More** and select **Restart Service**. In the displayed dialog box, verify the password and click **OK** to restart the Doris service. After the service is started, go to **Step 3**.

Step 3 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the alarm list, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.392 ALM-50202 FE CPU Usage Exceeds the Threshold

Alarm Description

The system checks the CPU usage of the FE instance every 30 seconds. The CPU usage has a default threshold. This alarm is generated when the CPU usage

exceeds the threshold (**95%** by default) for multiple consecutive times (**3** by default).

This alarm is cleared when **Trigger Count** is **1** and the CPU usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the CPU usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50202	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, select the name of the desired cluster, and choose **Doris** > **CPU and Memory** > **CPU Usage of FE (FE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

NOTE

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.393 ALM-50203 FE Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the FE instance every 30 seconds. This alarm is generated when the memory usage exceeds the threshold (**95%** by default) for multiple consecutive times (**3** by default).

This alarm is cleared when **Trigger Count** is **1** and the memory usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the memory usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50203	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, select the name of the desired cluster, and choose **Doris** > **CPU and Memory** > **Memory Usage of FE (FE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

NOTE

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.394 ALM-50205 BE CPU Usage Exceeds the Threshold

Alarm Description

The system checks the CPU usage of the BE instance every 30 seconds. This alarm is generated when the CPU usage exceeds the threshold (**95%** by default) for multiple consecutive times (**3** by default).

This alarm is cleared when **Trigger Count** is **1** and the CPU usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the CPU usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50205	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **CPU and Memory** > **CPU Usage of BE (BE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

NOTE

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.395 ALM-50206 BE Memory Usage Exceeds the Threshold

Alarm Description

The system checks the memory usage of the BE instance every 30 seconds. This alarm is generated when the memory usage exceeds the threshold for multiple consecutive times (**3** by default).

This alarm is cleared when **Trigger Count** is **1** and the memory usage is less than or equal to the threshold. This alarm is cleared when **Trigger Count** is greater than 1 and the memory usage is less than or equal to 85% of the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50206	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

The alarm threshold or alarm trigger count is improperly configured.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **CPU and Memory** > **Memory Usage of BE (BE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

NOTE

Trigger Count specifies how many times the threshold can be hit before an alarm is generated.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.396 ALM-50207 Ratio of Connections to the FE MySQL Port to the Maximum Connections Allowed Exceeds the Threshold

Alarm Description

The system checks the number of MySQL port connections every 30 seconds. This alarm is generated when the ratio of the number of current connections to the maximum number of FE port connections exceeds the threshold (95% by default). The maximum number of FE port connections in the current cluster is specified by the **qe_max_connection** parameter. The default value is **1024**.

This alarm is cleared when the number of MySQL port connections on the FE node is less than or equal to the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50207	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Processes respond slowly or do not work.

Possible Causes

- After the MySQL client is connected to Doris, the connection is not closed.
- A large number of services are concurrently connected to Doris.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M**, and click **Alarm** > **Thresholds** in the navigation pane on the left. Click the name of the desired cluster > **Doris** > **Connection** > **FE MySQL Port Connections (FE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

NOTE

If there are a large number of connections, ensure there are only necessary connections. Otherwise, the service performance may be degraded or even the service may be unavailable.

Step 4 Wait for 2 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.397 ALM-50208 Failures to Clear Historical Metadata Image Files Exceed the Threshold

Alarm Description

The system checks the number of failures to clear historical metadata image files on the FE node every 30 seconds. This alarm is generated when the number of failures exceeds the threshold (**1** by default).

This alarm is cleared when the system detects that the number of failures is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50208	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Doris metadata occupies more and more disk space, which may cause service exceptions.

Possible Causes

The Doris service is abnormal.

Handling Procedure

Check whether the Doris service is normal.

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Doris**.

Step 2 Check whether **Running Status** of the Doris service is **Normal**.

- If yes, go to **Step 4**.
- If no, go to **Step 3**.

Step 3 If the service process is not started, start it first and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Step 4 Check whether other Doris-related alarms are generated in the cluster. If yes, clear them by referring to the alarm help. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

You need to manually clear the alarm after the fault is rectified.

Related Information

None.

10.398 ALM-50209 Failures to Generate Metadata Image Files Exceed the Threshold

Alarm Description

The system checks the number of failures to generate metadata image files on the FE node every 30 seconds. This alarm is generated when the number of failures exceeds the threshold (**1** by default).

This alarm is cleared when the system detects that the number of failures is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50209	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The non-master FE node cannot receive the latest metadata image file. As a result, the system reliability deteriorates.

Possible Causes

The Doris service is abnormal.

Handling Procedure

Check the Doris service status.

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **Doris**.

Step 2 Check whether **Running Status** of the Doris service is **Normal**.

- If yes, go to **Step 4**.
- If no, go to **Step 3**.

Step 3 If the service process is not started, start it first.

Step 4 Check whether other alarms are generated in the cluster. If yes, clear the alarms by referring to the alarm help. Then, check whether this alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

You need to manually clear the alarm after the fault is rectified.

Related Information

None.

10.399 ALM-50210 Maximum Compaction Score of All BE Nodes Exceeds the Threshold

Alarm Description

The system checks the maximum compaction score of all BE nodes every 30 seconds. This alarm is generated when the maximum compaction score exceeds the threshold (10 by default).

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50210	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Query or write may be delayed.

Possible Causes

The number of concurrent service requests is large in the cluster, or the compaction queue is small.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **Performance** > **Maximum compaction score of all BE nodes (BE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Choose **Cluster** > **Services** > **Doris** > **Configurations** > **All Configurations** > **BE(Role)** > **Customization**, add the **max_base_compaction_threads** parameter to **be.conf** with a value of **10**, and add the **max_cumu_compaction_threads** parameter with a value **20**.

Step 6 Click **Save**. Click **Instances**, select the BE instances whose configuration has expired, click **More**, and select **Restart Instance** to restart the Doris BE instances.

Step 7 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.400 ALM-50211 FE Queue Length of BE Periodic Report Tasks Exceeds the Threshold

Alarm Description

The system checks the queue length of each BE periodic report task on FE every 30 seconds. This alarm is generated when the queue length exceeds the threshold (10 by default). This value indicates the number of report tasks waiting on the master FE node. A large value indicates a poor FE processing capability.

This alarm is cleared when the system detects that the queue length of BE periodic report tasks on FE is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50211	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The processing capability of FE is insufficient, affecting the service query speed.

Possible Causes

The processing capability of the master FE node is insufficient due to a large number of concurrent service requests in the Doris cluster or insufficient memory for FE processes.

Handling Procedure

Check the GC duration.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50211**.

Step 2 Choose **Cluster** > **Services** > **Doris** > **Instances**, click the FE instance for which the alarm is generated, and click the **Chart** tab of the instance.

Select **JVM** from **Chart Category** on the left, and check whether **Accumulated GC duration of the old generation** of the FE process is greater than 3 seconds.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Choose **Cluster** > **Services** > **Doris** > **Configurations** > **All Configurations** > **FE(Role)** > **JVM**, and increase the value of **-Xmx** in **FE_GC_OPTS**. The default value is **8GB**.

- If this alarm is generated occasionally, increase the value by 0.5 times. If this alarm is generated frequently, double the parameter value.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 5 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **Queue** > **Queue Length of BE Periodic Report Tasks on the FE (FE)**.

Step 6 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 7 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 8 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.401 ALM-50212 Accumulated Old-Generation GC Duration of the FE Process Exceeds the Threshold

Alarm Description

The system checks the accumulated old-generation GC duration of the FE process every 30 seconds. This alarm is generated when the accumulated GC duration exceeds the threshold (3000 ms by default).

This alarm is cleared when the system detects that the accumulated old-generation GC duration of the FE process is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50212	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A long GC duration of the FE process may interrupt the services.

Possible Causes

The heap memory of the FE process is overused or inappropriately allocated, causing frequent occurrence of the GC process.

Handling Procedure

Check the GC duration.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50212**.

Step 2 Choose **Cluster** > **Services** > **Doris** > **Instances**, click the FE instance for which the alarm is generated, and click the **Chart** tab of the instance.

Select **JVM** from **Chart Category** on the left, and check whether **Accumulated GC duration of the old generation** of the FE process is greater than 3 seconds.

- If yes, go to **Step 3**.
- If no, go to **Step 5**.

Step 3 Choose **Cluster** > **Services** > **Doris** > **Configurations** > **All Configurations** > **FE(Role)** > **JVM**, and increase the value of **-Xmx** in **FE_GC_OPTS**. The default value is **8G**.

- If this alarm is occasionally generated, increase the value by 0.5 times. If this alarm is frequently generated, double the value.
- In the case of large service volume and high service concurrency, you are advised to add instances.

Step 4 Check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 30 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.402 ALM-50213 Number of Tasks Queuing in the FE Thread Pool for Interacting with BE Exceeds the Threshold

Alarm Description

The system checks the number of queuing tasks in the FE thread pool for interacting with BE every 30 seconds. This alarm is generated when the number of queuing tasks exceeds the threshold (10 by default). This FE thread pool is the working thread pool of ThriftServer. It is specified by **rpc_port** in the **fe.conf** file and is used to interact with BE.

This alarm is cleared when the system detects that the number of tasks queuing in the FE thread pool for interacting with BE is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50213	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The read and write of the Doris service slows down.

Possible Causes

There are a large number of concurrent service requests, causing too many queuing tasks.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **Queue** > **Number of tasks that are queuing in the thread pool for interaction between the FE and the BE (FE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.403 ALM-50214 Number of Tasks Queuing in the FE Thread Pool for Task Processing Exceeds the Threshold

Alarm Description

The system checks the number of queuing tasks in the FE thread pool for processing tasks every 30 seconds. This alarm is generated when the number of queuing tasks exceeds the threshold (10 by default). This thread pool is used by the NIO MySQL Server to process tasks.

This alarm is cleared when the system detects that the number of tasks queuing in the FE thread pool for processing tasks is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50214	Minor	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.

Parameter	Description
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The task execution of the entire system becomes slow and blocked.

Possible Causes

Large tasks may block the task execution of the queue.

Handling Procedure

Check the execution status of FE tasks.

Step 1 On FusionInsight Manager, choose **Cluster** > **Services** > **Doris**. Click the **Chart** tab, select **Connection** from **Chart Category** in the left pane, and view the **FE MySQL Port Connections** chart. If the number of connections is large, click **Instances**, select the FE instance, and click the **Chart** tab. Select **CPU and Memory** from **Chart Category** and view the **CPU Usage of FE** chart. If the CPU usage is high, check the **Time** field in FE audit **log /var/log/Bigdata/audit/doris/fe/fe.audit.log** to collect statistics on the average task duration. If the value is also high, the alarm is caused by large concurrent tasks.

Step 2 After connecting to Doris, run the following command to check whether the default value of **queryTimeout** is too large. The default value is **300** seconds.

show variables like 'query_timeout';

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Run the following command to shorten the timeout period based on site requirements to block the tasks that take a long time:

set global query_timeout=*xxx*;

Step 4 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **Queue** > **Queue Length of Query Execution Thread Pool (BE)**.

Step 5 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 6 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 7 Wait 10 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.404 ALM-50215 Longest Duration of RPC Requests Received by Each FE Thrift Method Exceeds the Threshold

Alarm Description

The system checks the longest duration of RPC requests received by each FE Thrift method every 30 seconds. This alarm is generated when the longest duration exceeds the threshold (5000 ms by default).

This alarm is cleared when the longest duration of RPC requests received by each FE Thrift method is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50215	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A longer RPC duration indicates a higher performance load and slower network request processing, which may cause service congestion.

Possible Causes

- The network has a latency.
- There are too many concurrent large SQL tasks.

Handling Procedure

Step 1 Log in to the host where the faulty node is deployed as user **root** and run **ping** *IP addresses of all Doris nodes* to check whether the peer host can be pinged.

- If yes, go to **Step 3**.
- If no, go to **Step 2**.

Step 2 Contact the network administrator to restore the network.

Step 3 On FusionInsight Manager, choose **Cluster** > **Services** > **Doris**. Click the **Chart** tab, select **Connection** from **Chart Category** in the left pane, and view the **FE MySQL Port Connections** chart. If the number of connections is large, click **Instances**, select the FE instance, and click the **Chart** tab. Select **CPU and Memory** from **Chart Category** and view the **CPU Usage of FE** chart. If the CPU usage is high, check the **Time** field in FE audit **log /var/log/Bigdata/audit/doris/fe/fe.audit.log** to collect statistics on the average task duration. If the value is also high, the alarm is caused by large concurrent tasks.

Step 4 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **Performance** > **Longest duration of RPC requests received by each method of the FE thrift interface (FE)**.

Step 5 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 6 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 7 Wait 10 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.405 ALM-50216 Memory Usage of the FE Node Exceeds the Threshold

Alarm Description

The system checks the memory usage of the FE node every 30 seconds. This alarm is generated when the memory usage exceeds the threshold (95% by default).

This alarm is cleared when the memory usage of the FE node falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50216	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Task execution and client connection to the FE are affected.

Possible Causes

The FE heap memory is too small.

Handling Procedure

Check the FE heap memory usage.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, select the name of the desired cluster, and choose **Doris** > **CPU and Memory** > **Memory usage of the FE node (FE)**.

1. Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
2. Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 2 Log in to the FE node for which the alarm is generated as user **omm**, run the **top** command to check the memory usage of processes, locate the process with high memory usage, and check whether the process belongs to the current service and is running properly.

- If yes, go to **Step 3**.
- If no, isolate or stop the process, or adjust the memory size, and check whether the memory is released.

Step 3 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.406 ALM-50217 Heap Memory Usage of the FE Node Exceeds the Threshold

Alarm Description

The system checks the heap memory usage of the FE node every 30 seconds. This alarm is generated when the heap memory usage exceeds the threshold (95% by default).

This alarm is cleared when the heap memory usage of the FE node falls below the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50217	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.

Parameter	Description
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Task execution and client connection to the FE are affected.

Possible Causes

The FE heap memory is too small.

Handling Procedure

Check heap memory usage.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, select the name of the desired cluster, and choose **Doris** > **CPU and Memory** > **Heap memory usage of the FE node (FE)**.

1. Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.
2. Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 2 On FusionInsight Manager, choose **Cluster** > **Services** > **Doris** > **FE** > **Configurations** > **All Configurations**, search for the **FE_GC_OPTS** parameter, increase the value of **-Xmx** as required, click **Save**, and click **OK**.

NOTE

- If this alarm is generated, the heap memory configured for the current Doris instance is not enough for data transmission. You are advised to open the instance monitoring page, display the Doris heap memory resource status monitoring chart, and observe the change trend of the heap memory used by Doris in the monitoring chart. Then change the value of **-Xmx** to twice the current heap memory usage or to another value to meet site requirements.
- When setting the heap memory, you can set **-Xms** and **-Xmx** to approximately the same value to prevent performance deterioration caused by heap size adjustment after each GC.
- The sum of **-Xmx** and **XX:MaxPermSize** cannot be greater than the actual physical memory of the node server.

Step 3 Restart the affected services or instances and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.407 ALM-50219 Length of the Queue in the Thread Pool for Query Execution Exceeds the Threshold

Alarm Description

The system checks the length of the waiting queue in the query execution thread pool every 30 seconds. This alarm is generated when the length exceeds the threshold (20 by default).

This alarm is cleared when the length of the waiting queue in the current query execution thread pool is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50219	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The task execution of the entire system becomes slow and blocked.

Possible Causes

Large tasks may block the task execution of the queue.

Handling Procedure

Check the execution status of tasks.

Step 1 On FusionInsight Manager, choose **Cluster** > **Services** > **Doris**. Click the **Chart** tab, select **Connection** from **Chart Category** in the left pane, and view the **FE MySQL Port Connections** chart. If the number of connections is large, click **Instances**, select the FE instance, and click the **Chart** tab. Select **CPU and Memory** from **Chart Category** and view the **CPU Usage of FE** chart. If the CPU usage is high, check the **Time** field in FE audit **log /var/log/Bigdata/audit/doris/fe/fe.audit.log** to collect statistics on the average task duration. If the value is also high, the alarm is caused by large concurrent tasks.

Step 2 After connecting to Doris, run the following command to check the **queryTimeout** value of the system:

show variables like 'query_timeout';

If the value is too large, run the **set global query_timeout=*xxx*;** command to shorten the timeout interval and block tasks that last for a long time.

Step 3 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **Queue** > **Queue Length of Query Execution Thread Pool (BE)**.

Step 4 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 5 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 6 Wait 10 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 7**.

Collect fault information.

Step 7 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 8 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 9 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 10 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.408 ALM-50220 Error Rate of TCP Packet Receiving Exceeds the Threshold

Alarm Description

The system checks the rate of TCP packet receiving errors every 30 seconds. This alarm is generated when the error rate exceeds the threshold (5% by default).

This alarm is cleared when the error rate is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50220	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The task fails or data is lost.

Possible Causes

The network is faulty, so data cannot be sent.

Handling Procedure

Step 1 Log in to the host where the faulty node is deployed as user **root** and run **ping** *IP addresses of all Doris nodes* to check whether the peer host can be pinged.

- If yes, go to **Step 4**.
- If no, go to **Step 2**.

Step 2 Contact the network administrator to restore the network.

Step 3 Wait for a while and check whether the alarm is cleared in the alarm list.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.409 ALM-50221 BE Data Disk Usage Exceeds the Threshold

Alarm Description

The system checks the usage of BE data disks every 30 seconds. This alarm is generated when the disk usage exceeds the threshold (95% by default).

This alarm is cleared when the system detects that the disk usage is less than the threshold.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50221	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

New data fails to be written, and the task is interrupted.

Possible Causes

- The disk space of the cluster is full.
- Data skew occurs among BE nodes.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms**, and view the role name and the IP address for the hostname in **Location**.

Step 2 Expand the disk capacity of the node for which the alarm is generated.

Step 3 Go to **Step 4** if the expansion fails or the alarm persists after the expansion.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 6 Expand the **Hosts** drop-down list. In the **Select Host** dialog box that is displayed, select the abnormal host, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 1 hour ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.410 ALM-50222 Disk Status of a Specified Data Directory on BE Is Abnormal

Alarm Description

The system checks the disk status of a specified data directory on BE every 30 seconds. This alarm is generated when the disk status is not **1** (**1** indicates the normal state and **0** indicates the abnormal state). This alarm is cleared when the disk status of the specified data directory on BE becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50222	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Parameter	Description
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

Service data may be unavailable.

Possible Causes

- The hard disk is faulty.
- The disk permissions are set incorrectly.

Handling Procedure

Check whether a disk alarm is generated.

Step 1 On FusionInsight Manager, choose **O&M** > **Alarm** > **Alarms** and check whether **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault** exists.

- If yes, go to **Step 2**.
- If no, go to **Step 4**.

Step 2 Rectify the fault by referring to the handling procedure of **ALM-12014 Partition Lost** or **ALM-12033 Slow Disk Fault**. Then, check whether the alarm is cleared.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Modify disk permissions.

Step 4 Choose **O&M** > **Alarm** > **Alarms** and view **Location** and **Additional Information** of the alarm to obtain the location of the faulty disk.

Step 5 Log in to the node for which the alarm is generated as user **root**. Go to the directory where the faulty disk is located, and run the **ll** command to check whether the permission for the faulty disk is **711** and whether the user is **omm**.

- If yes, go to **Step 7**.
- If no, go to **Step 6**.

Step 6 Modify the permission of the faulty disk. For example, if the faulty disk is **data1**, run the following commands:

chown omm:wheel data1

chmod 711 data1

Step 7 Choose **Cluster** > **Services** > **Doris** > **Instances**, select this BE instance, click **More**, and select **Restart Instance**. Wait 5 minutes and check whether an alarm is generated.

- If no, no further action is required.
- If yes, go to **Step 8**.

Collect fault information.

Step 8 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 9 Expand the **Service** drop-down list, and select **Doris** and **OMS** for the target cluster.

Step 10 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 20 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 11 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.411 ALM-50223 Maximum Memory Required by BE Is Greater Than the Remaining Memory of the Machine

Alarm Description

The system checks whether the maximum memory required by BE is greater than the available memory every 30 seconds. This alarm is generated when the value is not **1** (**1** indicates that the maximum required memory is less than or equal to the available memory, and **0** indicates that the maximum required memory is greater than the available memory).

This alarm is cleared when the maximum required memory is less than or equal to the available memory.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50223	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A task may fail to apply for memory when running.

Possible Causes

Too much BE node memory has been occupied by other processes, or the maximum memory set for the BE service is too large.

Handling Procedure

Check whether the maximum memory set for the BE node is proper.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **CPU and Memory** > **Relationship between the maximum memory size of the BE and the remaining memory size of the machine (BE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 5**.

Step 5 Log in to the BE node for which the alarm is generated as user **omm**, run the **top** command to check the memory usage of processes, locate the process with high memory usage, and check whether the process belongs to the current service and is running properly.

- If yes, go to **Step 6**.
- If no, isolate or stop the process, or adjust the memory size, and check whether the memory is released.

Step 6 Log in to the BE node for which the alarm is generated as user **omm** and run the **free -g** command to check the total memory and remaining memory in the system and estimate the memory usage.

Step 7 On FusionInsight Manager, choose **Cluster** > **Services** > **Doris** > **Configurations** > **All Configurations** > **BE(Role)** > **Memory** and decrease the value of **mem_limit**. This parameter specifies the maximum memory allowed for BE. Then save the modification and restart the BE instance.

Step 8 After the BE instance is restarted, wait 5 minutes and check whether the alarm is cleared.

- If yes, no further action is required.
- If no, go to **Step 9**.

Collect fault information.

Step 9 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 10 Expand the **Service** drop-down list, and select **Doris** for the target cluster.

Step 11 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 12 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.412 ALM-50224 Failures a Certain Task Type on BE Are Increasing

Alarm Description

The system checks whether the number of failed tasks of a certain type on BE is increasing every 30 seconds. This alarm is generated when the system detects that the value is not **1** (**1** indicates that the number of failed tasks of a certain type does not increase, and **0** indicates that the failed tasks of a certain type are increasing).

This alarm is cleared when the system detects that the number of failed tasks of a certain type on BE does not increase.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50224	Major	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

A task fails to be executed repeatedly in a certain scenario.

Possible Causes

A BE exception may occur. As a result, the number of failed tasks increases in a certain scenario.

Handling Procedure

Check whether the alarm threshold or alarm trigger count is properly configured.

Step 1 Log in to FusionInsight Manager, choose **O&M** > **Alarm** > **Thresholds**, click the name of the desired cluster, and choose **Doris** > **Exception** > **Check whether the number of failed tasks of a certain type increases (BE)**.

Step 2 Click the edit button next to **Trigger Count**, change the number based on site requirements, and click **OK**.

Step 3 Click **Modify** in the **Operation** column, change the alarm threshold based on site requirements, and click **OK**.

Step 4 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.

- If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 7 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

You need to manually clear the alarm after the fault is rectified.

Related Information

None.

10.413 ALM-50225 FE Instance Fault

Alarm Description

The system checks the FE process status every 30 seconds. This alarm is generated when the value is greater than **0** (**0** indicates that the FE process is normal and **1** indicates that the FE process is abnormal).

This alarm is cleared when the system detects that the FE process becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50225	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The current FE instance is unavailable.

Possible Causes

The FE instance is faulty or restarted.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50225**.

Step 2 Choose **Cluster** > **Services** > **Doris** > **Instances**, click the FE instance for which the alarm is generated, and check whether **Running Status** of the instance is **Restoring**.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.414 ALM-50226 BE Instance Fault

Alarm Description

The system checks the BE process status every 30 seconds. This alarm is generated when the value is greater than **0** (**0** indicates that the BE process is normal and **1** indicates that the BE process is abnormal).

This alarm is cleared when the system detects that the BE process becomes normal.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50226	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.
Trigger Condition	Specifies the threshold for triggering the alarm.

Impact on the System

The current BE instance is unavailable.

Possible Causes

The BE instance is faulty or restarted.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **O&M** > **Alarm** > **Alarms**. In the alarm list, view the role name and obtain the IP address of the instance in **Location** of the alarm whose ID is **50226**.

Step 2 Choose **Cluster** > **Services** > **Doris** > **Instances**, click the BE instance for which the alarm is generated, and check whether **Running Status** of the instance is **Restoring**.

- If yes, go to **Step 3**.
- If no, go to **Step 4**.

Step 3 Wait 2 minutes and check whether the alarm is automatically cleared.

- If yes, no further action is required.
- If no, go to **Step 4**.

Collect fault information.

Step 4 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 5 Expand the **Service** drop-down list, select **Doris** for the target cluster, and click **OK**.

Step 6 Click the edit icon in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 7 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.415 ALM-50401 Number of JobServer Jobs Waiting to Be Executed Exceeds the Threshold

Alarm Description

The system checks the number of jobs submitted to JobServer every 30 seconds. This alarm is generated when the number of jobs to be executed exceeds 800.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50401	Critical (default threshold: 900) Major (default threshold: 800)	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm was generated.
ServiceName	Specifies the service for which the alarm was generated.
RoleName	Specifies the role for which the alarm was generated.
HostName	Specifies the host for which the alarm was generated.

Impact on the System

Too many JobServer jobs are detected in the queue. For details about the queue usage, see the **Additional Information** field of this alarm. The impacts are as follows:

1. When the number of JobServer jobs in the queue reaches the maximum (1000 by default), new jobs cannot be added.
2. Before the number of JobServer jobs in the queue reaches the maximum, new JobServer jobs cannot be submitted quickly. For example, it takes more time (even hours) to submit added jobs or add new jobs to Yarn.

Possible Causes

Too many jobs are submitted instantaneously.

Handling Procedure

Step 1 Log in to FusionInsight Manager and choose **Cluster** > **Services** > **JobGateway**.

Step 2 Click the **Instances** tab, click **Add Instance**, and add JobServer instances based on the number of submitted jobs.

Step 3 After the instances are added, restart the JobGateway service.

Step 4 Wait 5 minutes and check whether the alarm is automatically cleared.

If yes, no further action is required.

If no, go to **Step 5**.

Collect fault information.

Step 5 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 6 Expand the **Service** drop-down list, and select **JobGateway** for the target cluster.

Step 7 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 8 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

10.416 ALM-50402 JobGateway Service Unavailable

Alarm Description

The system checks the JobGateway service status every 60 seconds. This alarm is generated when the JobGateway service is abnormal.

This alarm is cleared when the JobGateway service recovers.

Alarm Attributes

Alarm ID	Alarm Severity	Auto Cleared
50402	Critical	Yes

Alarm Parameters

Parameter	Description
Source	Specifies the cluster or system for which the alarm is generated.
ServiceName	Specifies the service for which the alarm is generated.

Parameter	Description
RoleName	Specifies the role for which the alarm is generated.
HostName	Specifies the host for which the alarm is generated.

Impact on the System

No job submission operation can be performed on JobGateway in the cluster. The components that depend on JobGateway in the cluster will become faulty.

Possible Causes

The node where the JobGateway service locates is faulty.

Handling Procedure

Step 1 Log in to FusionInsight Manager, choose **Cluster** > **Services** > **JobGateway**, and click the **Instance** tab. Check for JobServer or JobBalancer instances that are faulty or not started and view the host names of these instances.

Step 2 On the **Alarm** page of FusionInsight Manager, check whether the **NodeAgent Process Is Abnormal** alarm is generated.

- If yes, go to **Step 3**.
- If no, go to **Step 6**.

Step 3 Check whether the host name in the alarm information is the same as the host name in **Step 1**.

- If yes, go to **Step 4**.
- If no, go to **Step 6**.

Step 4 Clear the alarm by following the instructions provided in **ALM-12006 NodeAgent Process Is Abnormal**.

Step 5 In the alarm list, check whether alarm **JobGateway Service Unavailable** is cleared.

- If yes, no further action is required.
- If no, go to **Step 6**.

Collect fault information.

Step 6 On FusionInsight Manager, choose **O&M**. In the navigation pane on the left, choose **Log** > **Download**.

Step 7 Expand the **Service** drop-down list, and select **JobGateway** for the target cluster.

Step 8 Click  in the upper right corner, and set **Start Date** and **End Date** for log collection to 10 minutes ahead of and after the alarm generation time, respectively. Then, click **Download**.

Step 9 Contact O&M personnel and provide the collected logs.

----End

Alarm Clearance

This alarm is automatically cleared after the fault is rectified.

Related Information

None.

11 Security Description

11.1 Security Configuration Suggestions for Clusters with Kerberos Authentication Disabled

The Hadoop community version provides two authentication modes: Kerberos authentication (security mode) and Simple authentication (normal mode). When creating a cluster, you can choose to enable or disable Kerberos authentication.

Clusters in security mode use the Kerberos protocol for security authentication.

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

Therefore, in normal mode, network access permissions must be strictly controlled to ensure cluster security. You are advised to perform the following operations to ensure cluster security.

- Deploy service applications on ECSs in the same VPC and subnet and avoid accessing MRS clusters through an external network.
- Configure security group rules to strictly control the access scope. Do not configure access rules that allow **Any** or **0.0.0.0** for the inbound direction of MRS cluster ports.
- If you want to access the native pages of the components in the cluster from the external, follow instructions in **Creating an SSH Channel for Connecting to an MRS Cluster and Configuring the Browser** for configuration.

11.2 Security Authentication Principles and Mechanisms

Function

For clusters in security mode with Kerberos authentication enabled, security authentication is required during application development.

Kerberos adopts a client/server structure and encryption technologies such as AES, and supports mutual authentication (both the client and server can authenticate each other). Kerberos is used to prevent interception and replay attacks and protect data integrity. It is a system that manages keys by using a symmetric key mechanism.

Architecture

Kerberos architecture is shown in **Figure 11-1** and module description in **Table 11-1**.

Figure 11-1 Kerberos architecture

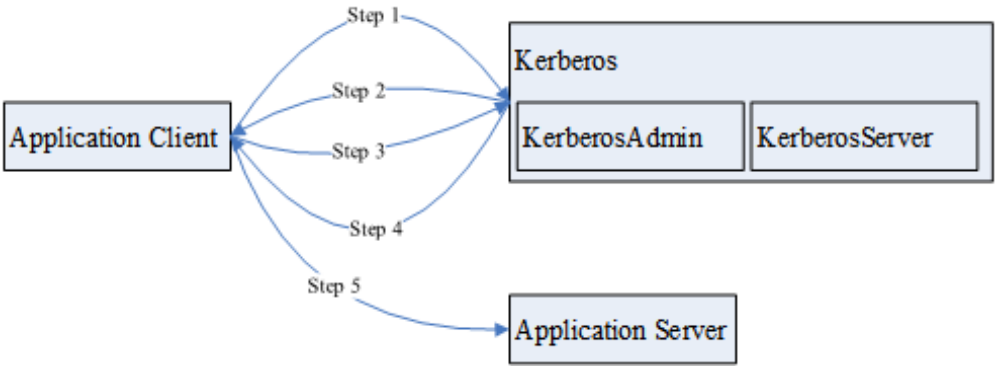

Table 11-1 Module description

Module	Description
Application Client	An application client, which is usually an application that submits tasks or jobs
Application Server	An application server, which is usually an application that an application client accesses
Kerberos	A service that provides security authentication
KerberosAdmin	A process that provides authentication user management

Module	Description
KerberosServer	A process that provides authentication ticket distribution

The process and principle are described as follows:

An application client can be a service in a cluster or a secondary development application of the customer. An application client can submit tasks or jobs to an application service.

1. Before submitting a task or job, the application client needs to apply for a ticket granting ticket (TGT) from the Kerberos service to establish a secure session with the Kerberos server.
2. After receiving the TGT request, the Kerberos service resolves parameters in the request to generate a TGT, and uses the key of the username specified by the client to encrypt the response.
3. After receiving the TGT response, the application client (based on the underlying RPC) resolves the response and obtains the TGT, and then applies for a server ticket (ST) of the application server from the Kerberos service.
4. After receiving the ST request, the Kerberos service verifies the TGT validity in the request and generates an ST of the application service, and then uses the application service key to encrypt the response.
5. After receiving the ST response, the application client packages the ST into a request and sends the request to the application server.
6. After receiving the request, the application server uses its local application service key to resolve the ST. After successful verification, the request becomes valid.

Basic Concepts

The following concepts can help users learn the Kerberos architecture quickly and understand the Kerberos service better. The following uses security authentication for HDFS as an example.

TGT

A TGT is generated by the Kerberos service and used to establish a secure session between an application and the Kerberos server. The validity period of a TGT is 24 hours. After 24 hours, the TGT expires automatically.

The following describes how to apply for a TGT (HDFS is used as an example):

1. Obtain a TGT through an API provided by HDFS.

```
/**
 * login Kerberos to get TGT, if the cluster is in security mode
 * @throws IOException if login is failed
 */
private void login() throws IOException {
// not security mode, just return
  if (! "kerberos".equalsIgnoreCase(conf.get("hadoop.security.authentication"))) {
      return;
  }
```



```
    //security mode
    System.setProperty("java.security.krb5.conf", PATH_TO_KRB5_CONF);

    UserGroupInformation.setConfiguration(conf);
    UserGroupInformation.loginUserFromKeytab(PRNCIPAL_NAME, PATH_TO_KEYTAB);
  }
```

2. Run shell commands on the client in kinit mode.

ST

An ST is generated by the Kerberos service and used to establish a secure session between an application and application service. An ST is valid only once.

In FusionInsight products, the generation of an ST is based on the Hadoop-RPC communication. The underlying RPC submits a request to the Kerberos server and the Kerberos server generates an ST.

Sample Authentication Code

```
package com.huawei.bigdata.hdfs.examples;

import java.io.IOException;

import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.fs.FileStatus;
import org.apache.hadoop.fs.FileSystem;
import org.apache.hadoop.fs.Path;
import org.apache.hadoop.security.UserGroupInformation;

public class KerberosTest {
    private static String PATH_TO_HDFS_SITE_XML = KerberosTest.class.getClassLoader().getResource("hdfs-
site.xml")
            .getPath();
    private static String PATH_TO_CORE_SITE_XML = KerberosTest.class.getClassLoader().getResource("core-
site.xml")
            .getPath();
    private static String PATH_TO_KEYTAB =
KerberosTest.class.getClassLoader().getResource("user.keytab").getPath();
    private static String PATH_TO_KRB5_CONF =
KerberosTest.class.getClassLoader().getResource("krb5.conf").getPath();
    private static String PRNCIPAL_NAME = "develop";
    private FileSystem fs;
    private Configuration conf;

    /**
     * initialize Configuration
     */
    private void initConf() {
        conf = new Configuration();

        // add configuration files
        conf.addResource(new Path(PATH_TO_HDFS_SITE_XML));
        conf.addResource(new Path(PATH_TO_CORE_SITE_XML));
    }

    /**
     * login Kerberos to get TGT, if the cluster is in security mode
     * @throws IOException if login is failed
     */
    private void login() throws IOException {
        // not security mode, just return
        if (! "kerberos".equalsIgnoreCase(conf.get("hadoop.security.authentication"))) {
            return;
        }

        //security mode
        System.setProperty("java.security.krb5.conf", PATH_TO_KRB5_CONF);
```



```
        UserGroupInformation.setConfiguration(conf);
        UserGroupInformation.loginUserFromKeytab(PRNCIPAL_NAME, PATH_TO_KEYTAB);
    }

    /**
     * initialize FileSystem, and get ST from Kerberos
     * @throws IOException
     */
    private void initFileSystem() throws IOException {
        fs = FileSystem.get(conf);
    }

    /**
     * An example to access the HDFS
     * @throws IOException
     */
    private void doSth() throws IOException {
        Path path = new Path("/tmp");
        FileStatus fStatus = fs.getFileStatus(path);
        System.out.println("Status of " + path + " is " + fStatus);
        //other thing
    }


    public static void main(String[] args) throws Exception {
        KerberosTest test = new KerberosTest();
        test.initConf();
        test.login();
        test.initFileSystem();
        test.doSth();
    }
}
```

NOTE

1. During Kerberos authentication, you need to configure the file parameters required for configuring the Kerberos authentication, including the keytab path, Kerberos authentication username, and the **krb5.conf** configuration file of the client for Kerberos authentication.
2. Method **login()** indicates calling the Hadoop API to perform Kerberos authentication and generating a TGT.
3. Method **doSth** indicates calling the Hadoop API to access the file system. In this situation, the underlying RPC automatically carries the TGT to Kerberos for verification and then an ST is generated.
4. The preceding code can be used to create **KerberosTest.java** in the HDFS secondary development sample project in security mode and run and view the commissioning result. For details, see **HDFS Development Guide (Security Mode)**.

12 High-Risk Operations

Forbidden Operations

Table 12-1 lists forbidden operations during the routine cluster operation and maintenance process.

Table 12-1 Forbidden operations

Item	Risk
Delete ZooKeeper data directories.	ClickHouse, HDFS, Yarn, HBase, and Hive depend on ZooKeeper, which stores metadata. This operation has adverse impact on normal operating of related components.
Frequently switch over the active and standby JDBCServer nodes.	This operation may interrupt services.
Delete Phoenix system tables and data (SYSTEM.CATALOG, SYSTEM.STATS, SYSTEM.SEQUENCE, and SYSTEM.FUNCTION).	This operation will cause service operation failures.
Manually modify data in the Hive metabase (hivemeta database).	This operation may cause Hive data parse errors. As a result, Hive cannot provide services.
Manually perform **INSERT** or **UPDATE** operations on Hive metadata tables.	This operation may cause Hive data parse errors. As a result, Hive cannot provide services.
Change permission on the Hive private file directory **hdfs:///tmp/hive-scratch**.	This operation may cause unavailable Hive services.
Modify **broker.id** in the Kafka configuration file.	This operation may cause invalid node data.

Item	Risk
Modify the host names of nodes.	Instances and upper-layer components on the host cannot provide services properly. The fault cannot be rectified.
Reinstall the OS of a node.	This operation will cause MRS cluster exceptions, leaving MRS clusters in abnormal status.
Use private images.	This operation will cause MRS cluster exceptions, leaving MRS clusters in abnormal status.

The following tables list the high-risk operations during the operation and maintenance of each component.

High-Risk Operations on a Cluster

Table 12-2 High-risk operations on a cluster

Operation	Risk	Severity	Workaround	Check Item
Modify the file directory or file permissions of user **omm** without permission.	This operation will lead to MRS service unavailability.	▲ ▲ ▲ ▲ ▲	Do not perform this operation.	Check whether the MRS cluster service is available.
Bind an EIP.	This operation exposes the Master node hosting MRS Manager of the cluster to the public network, increasing the risk of network attacks from the Internet.	▲ ▲ ▲ ▲ ▲	Ensure that the bound EIP is a trusted public IP address.	None

Operation	Risk	Severity	Workaround	Check Item
Enable security group rules for port 22 of a cluster.	This operation increases the risk of exploiting vulnerability of port 22.	▲ ▲ ▲ ▲ ▲	Configure a security group rule for port 22 to allow only trusted IP addresses to access the port. You are not advised to configure the inbound rule to allow 0.0.0.0 to access the port.	None
Delete a cluster or cluster data.	Data will get lost.	▲ ▲ ▲ ▲ ▲	Before deleting the data, confirm the necessity of the operation and ensure that the data has been backed up.	None
Scale in a cluster.	Data will get lost.	▲ ▲ ▲ ▲ ▲	Before scaling in the cluster, confirm the necessity of the operation and ensure that the data has been backed up.	None
Detach or format a data disk.	Data will get lost.	▲ ▲ ▲ ▲ ▲	Before performing this operation, confirm the necessity of the operation and ensure that the data has been backed up.	None

Manager High-Risk Operations

Table 12-3 Manager high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Change the OMS password.	This operation will restart all processes of OMSServer, which has adverse impact on cluster maintenance and management.	▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check whether there are uncleared alarms and whether the cluster management and maintenance are normal.
Import the certificate.	This operation will restart OMS processes and the entire cluster, which has adverse impact on cluster maintenance and management and services.	▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.
Perform an upgrade.	This operation will restart Manager and the entire cluster, affecting management, maintenance, and services of the cluster. Strictly manage the user who is eligible to assign the cluster management permission to prevent security risks.	▲ ▲ ▲	Ensure that there is no other maintenance and management operations when the operation is performed.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.

Operation	Risk	Severity	Workaround	Check Item
Restore the OMS.	This operation will restart Manager and the entire cluster, affecting management, maintenance, and services of the cluster.	▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.
Change an IP address.	This operation will restart Manager and the entire cluster, affecting management, maintenance, and services of the cluster.	▲ ▲ ▲	Ensure that there is no other maintenance and management operations when the operation is performed and that the new IP address is correct.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.
Change log levels.	If the log level is changed to **DEBUG**, Manager responds slowly.	▲ ▲	Before the modification, confirm the necessity of the operation and change it back to the default log level in time.	None

Operation	Risk	Severity	Workaround	Check Item
Replace a control node.	This operation will interrupt services deployed on the node. If the node is a management node, the operation will restart all OMS processes, affecting the cluster management and maintenance.	▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.
Replace a management node.	This operation will interrupt services deployed on the node. As a result, OMS processes will be restarted, affecting the cluster management and maintenance.	▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.
Restart the upper-layer service at the same time during the restart of a lower-layer service.	This operation will interrupt the upper-layer service, affecting the management, maintenance, and services of the cluster.	▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.

Operation	Risk	Severity	Workaround	Check Item
Modify the OLDAP port.	This operation will restart the LdapServer and Kerberos services and all associated services, affecting service running.	▲ ▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	None
Delete the **supergroup** group.	Deleting the **supergroup** group decreases user rights, affecting service access.	▲ ▲ ▲ ▲ ▲	Before the change, confirm the rights to be added. Ensure that the required rights have been added before deleting the **supergroup** rights to which the user is bound, ensuring service continuity.	None
Restart a service.	Services will be interrupted during the restart. If you select and restart the upper-layer service, the upper-layer services that depend on the service will be interrupted.	▲ ▲ ▲	Confirm the necessity of restarting the system before the operation.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.

Operation	Risk	Severity	Workaround	Check Item
Change the default SSH port No.	After the default port (22) is changed, functions such as cluster creation, service/instance adding, host adding, and host reinstallation cannot be used, and results of cluster health check items for node mutual trust, **omm**/**ommdba** user password expiration, and others are incorrect.	▲ ▲ ▲	Before performing this operation, restore the SSH port to the default value.	None

CDL High-risk Operations

Table 12-4 CDL high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Start or stop basic components independently.	This operation has adverse impact on the basic functions of some services. As a result, service failures occur.	▲ ▲ ▲	Do not start or stop basic components such as Kafka, DBService, ZooKeeper, Kerberos, and LDAP separately. To start or stop basic components, select associated services.	Check whether the service status is normal.
Restart or stop services.	This operation may interrupt services.	▲ ▲	Restart or stop services when necessary.	Check whether the service is running properly.

ClickHouse High-Risk Operations

Table 12-5 ClickHouse high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Delete data directories.	This operation may cause service information loss.	▲ ▲ ▲	Do not delete data directories manually.	Check whether data directories are normal.
Remove ClickHouseServer instances.	The ClickHouseServer instance nodes in the same shard must be removed in at the same time. Otherwise, the topology information of the logical cluster is incorrect. Before performing this operation, check the database and data table information of each node in the logical cluster and perform scale-in pre-analysis to ensure that data is successfully migrated during the scale-in process to prevent data loss	▲ ▲ ▲ ▲ ▲	Before scale-in, collect information in advance to learn the status of the ClickHouse logical cluster and instance nodes.	Check the ClickHouse logical cluster topology information, database and data table information in each ClickHouseServer instance, and data volume.
Add ClickHouseServer instances.	When performing this operation, you must check whether a database or data table with the same name as that on the old node needs to be created on the new node. Otherwise, subsequent data migration, data balancing, scale-in, and decommissioning will fail.	▲ ▲ ▲ ▲ ▲	Before scale-out, confirm the function and purpose of new ClickHouseServer instances and determine whether to create related databases and data tables.	Check the ClickHouse logical cluster topology information, database and data table information in each ClickHouseServer instance, and data volume.

Operation	Risk	Severity	Workaround	Check Item
Decommission ClickHouseServer instances.	The ClickHouseServer instance nodes in the same shard must be decommissioned in at the same time. Otherwise, the topology information of the logical cluster is incorrect. Before performing this operation, check the database and data table information of each node in the logical cluster and perform decommissioning pre-analysis to ensure that data is successfully migrated during the decommissioning process to prevent data loss	▲ ▲ ▲ ▲ ▲	Before decommissioning, collect information in advance to learn the status of the ClickHouse logical cluster and instance nodes.	Check the ClickHouse logical cluster topology information, database and data table information in each ClickHouseServer instance, and data volume.
Recommission ClickHouseServer instances.	When performing this operation, you must select all nodes in the original shard. Otherwise, the topology information of the logical cluster is incorrect.	▲ ▲ ▲ ▲ ▲	Before recommissioning, you need to confirm the home information about the shards of the node to be recommissioned.	Check the ClickHouse logical cluster topology information.
Modify data directory content (file and folder creation).	This operation may cause the ClickHouse instance of the node faults.	▲ ▲ ▲	Do not create or modify files or folders in the data directories manually.	Check whether data directories are normal.

Operation	Risk	Severity	Workaround	Check Item
Start or stop basic components independently.	This operation has adverse impact on the basic functions of some services. As a result, service failures occur.	▲ ▲ ▲	Do not start or stop ZooKeeper, Kerberos, and LDAP basic components independently. Select related services when performing this operation.	Check whether the service status is normal.
Restart or stop services.	This operation may interrupt services.	▲ ▲	Restart or stop services when necessary.	Check whether the service is running properly.

DBService High-Risk Operations

Table 12-6 DBService high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Change the DBService password.	The services need to be restarted for the password change to take effect. The services are unavailable during the restart.	▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check whether there are uncleared alarms and whether the cluster management and maintenance are normal.

Operation	Risk	Severity	Workaround	Check Item
Restore DBService data.	After the data is restored, the data generated after the data backup and before the data restoration is lost. After the data is restored, the configuration of the components that depend on DBService may expire and these components need to be restarted.	▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check whether there are uncleared alarms and whether the cluster management and maintenance are normal.
Perform active/ standby DBService switchover.	During the DBServer switchover, DBService is unavailable.	▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	None
Change the DBService floating IP address.	The DBService needs to be restarted for the change to take effect. The DBService is unavailable during the restart. If the floating IP address has been used, the configuration will fail, and the DBService will fail to be started.	▲ ▲ ▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.

Flink High-Risk Operations

Table 12-7 Flink high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Change log levels.	If the log level is modified to DEBUG, the task running performance is affected.	▲ ▲	Before the modification, confirm the necessity of the operation and change it back to the default log level in time.	None
Modify file permissions.	Tasks may fail.	▲ ▲ ▲	Confirm the necessity of the operation before the modification.	Check whether related service operations are normal.

Flume High-Risk Operations

Table 12-8 Flume high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Modify the Flume instance start parameter **GC_OPTS**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.

Operation	Risk	Severity	Workaround	Check Item
Change the default value of **dfs.replication** from **3** to **1**.	This operation will have the following impacts: 1. The storage reliability deteriorates. If the disk becomes faulty, data will be lost. 2. NameNode fails to be restarted, and the HDFS service is unavailable.	▲ ▲ ▲ ▲	When modifying related configuration items, check the parameter description carefully. Ensure that there are more than two replicas for data storage.	Check whether the default replica number is not 1 and whether the HDFS service is normal.

HBase High-Risk Operations

Table 12-9 HBase high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Modify encryption configuration. • hbase.regionserver.wal.encryption • hbase.crypto.keyprovider.parameters.uri • hbase.crypto.keyprovider.parameters.encryptedtext	Services cannot start properly.	▲ ▲ ▲ ▲	Strictly follow the prompt information when modifying related configuration items, which are associated. Ensure that new values are valid.	Check whether services can be started properly.

Operation	Risk	Severity	Workaround	Check Item
Change the value of **hbase.regionserver.wal.encryption** to **false** or switch encryption algorithm from AES to SMS4.	This operation may cause start failures and data loss.	▲ ▲ ▲ ▲	When HFile and WAL are encrypted using an encryption algorithm and a table is created, do not close or switch the encryption algorithm randomly. If an encryption table (ENCRYPTION=>AES/SMS4) is not created, you can only switch the encryption algorithm.	None
Modify HBase instance start parameter **GC_OPTS** and **HBASE_HEAPSIZE**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid. GC_OPTS does not conflict with HBASE_HEAPSIZE.	Check whether services can be started properly.
Use **OfflineMetaRepair** tool	Services cannot start properly.	▲ ▲ ▲ ▲	This tool can be used only when HBase is offline and cannot be used in data migration scenarios.	Check whether HBase services can be started properly.

HDFS High-Risk Operations

Table 12-10 HDFS high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Change HDFS NameNode data storage directory **dfs.namenode.name.dir** and data configuration directory **dfs.datanode.data.dir**.	Services cannot start properly.	▲ ▲ ▲ ▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Use the **-delete** parameter when you run the **hadoop distcp** command.	During DistCP copying, files that do not exist in the source cluster but exist in the destination cluster are deleted from the destination cluster.	▲ ▲	When using DistCP, determine whether to retain the redundant files in the destination cluster. Exercise caution when using the **-delete** parameter.	After DistCP copying is complete, check whether the data in the destination cluster is retained or deleted according to the parameter settings.

Operation	Risk	Severity	Workaround	Check Item
Modify the HDFS instance start parameter **GC_OPTS**, **HADOOP_HEAPSIZE**, and **GC_PROFILE**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid. GC_OPTS does not conflict with HADOOP_HEAPSIZE.	Check whether services can be started properly.
Change the default value of **dfs.replication** from **3** to **1**.	This operation will have the following impacts: 1. The storage reliability deteriorates. If the disk becomes faulty, data will be lost. 2. NameNode fails to be restarted, and the HDFS service is unavailable.	▲ ▲ ▲ ▲	When modifying related configuration items, check the parameter description carefully. Ensure that there are more than two replicas for data storage.	Check whether the default replica number is not 1 and whether the HDFS service is normal.
Change the remote procedure call (RPC) channel encryption mode (**hadoop.rpc.protection**) of each module in Hadoop.	This operation causes service faults and service exceptions.	▲ ▲ ▲ ▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether HDFS and other services that depend on HDFS can properly start and provide services.

Hive High-Risk Operations

Table 12-11 Hive high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Modify the Hive instance start parameter **GC_OPTS**.	This operation may cause Hive instance start failures.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Delete all MetaStore instances.	This operation may cause Hive metadata loss. As a result, Hive cannot provide services.	▲ ▲ ▲	Do not perform this operation unless ensure that Hive table information can be discarded.	Check whether services can be started properly.
Delete or modify files corresponding to Hive tables over HDFS interfaces or HBase interfaces.	This operation may cause Hive service data loss or tampering.	▲ ▲	Do not perform this operation unless ensure that the data can be discarded or that the operation meets service requirements.	Check whether Hive data is complete.

Operation	Risk	Severity	Workaround	Check Item
Delete or modify files corresponding to Hive tables or directory access permission over HDFS interfaces or HBase interfaces.	This operation may cause related service scenarios to be unavailable.	▲ ▲ ▲	Do not perform this operation.	Check whether related service operations are normal.
Delete or modify **hdfs:///apps/templeton/hive-3.1.0.tar.gz** over HDFS interfaces.	WebHCat fails to perform services due to this operation.	▲ ▲	Do not perform this operation.	Check whether related service operations are normal.
Export table data to overwrite the data at the local. For example, export the data of **t1** to **/opt/dir**. **insert overwrite local directory '/opt/dir' select * from t1;**	This operation will delete target directories. Incorrect setting may cause software or OS startup failures.	▲ ▲ ▲ ▲ ▲	Ensure that the path where the data is written does not contain any files or do not use the key word overwrite in the command.	Check whether files in the target path are lost.

Operation	Risk	Severity	Workaround	Check Item
Direct different databases, tables, or partition files to the same path, for example, default warehouse path **/user/hive/warehouse**.	The creation operation may cause disordered data. After a database, table, or partition is deleted, other object data will be lost.	▲ ▲ ▲ ▲ ▲	Do not perform this operation.	Check whether files in the target path are lost.

IoTDB High-Risk Operations

Table 12-12 IoTDB high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Delete data directories.	This operation may cause service information loss.	▲ ▲ ▲	Do not delete data directories manually.	Check whether data directories are normal.
Modify data directory content (file and folder creation).	This operation may cause the IoTDB instance of the node faults.	▲ ▲ ▲	Do not create or modify files or folders in the data directories manually.	Check whether data directories are normal.

Operation	Risk	Severity	Workaround	Check Item
Start or stop basic components independently.	This operation has adverse impact on the basic functions of some services. As a result, service failures occur.	▲ ▲ ▲	Do not start or stop Kerberos, and LDAP basic components independently. Select related services when performing this operation.	Check whether the service status is normal.
Restart or stop services.	This operation may interrupt services.	▲ ▲	Restart or stop services when necessary.	Check whether the service is running properly.

Kafka High-Risk Operations

Table 12-13 Kafka high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Delete Topic	This operation may delete existing topics and data.	▲ ▲ ▲	Kerberos authentication is used to ensure that authenticated users have operation permissions. Ensure that topic names are correct.	Check whether topics are processed properly.
Delete data directories.	This operation may cause service information loss.	▲ ▲ ▲	Do not delete data directories manually.	Check whether data directories are normal.

Operation	Risk	Severity	Workaround	Check Item
Modify data directory content (file and folder creation).	This operation may cause the Broker instance of the node faults.	▲ ▲ ▲	Do not create or modify files or folders in the data directories manually.	Check whether data directories are normal.
Modify the disk auto-adaptation function using the **disk.adapter.enable** parameter.	This operation adjusts the topic data retention period when the disk usage reaches the threshold. Historical data that does not fall within the storage retention may be deleted.	▲ ▲ ▲	If the retention period of some topics cannot be adjusted, add this topic to the value of **disk.adapter.topic.blacklist**.	Observe the data storage period on the Kafka topic monitoring page.
Modify data directory **log.dirs** configuration.	Incorrect operation may cause process faults.	▲ ▲ ▲	Ensure that the added or modified data directories are empty and that the directory permissions are right.	Check whether data directories are normal.
Reduce the capacity of the Kafka cluster.	This operation may cause quantity reduction of backups of some data duplicates of topic. As a result, some topics cannot be accessed.	▲ ▲	Perform backup operation and then reduce the capacity of the Kafka cluster.	Check whether backup nodes where partitions are located are activated to ensure data security.

Operation	Risk	Severity	Workaround	Check Item
Start or stop basic components independently.	This operation has adverse impact on the basic functions of some services. As a result, service failures occur.	▲ ▲ ▲	Do not start or stop ZooKeeper, Kerberos, and LDAP basic components independently. Select related services when performing this operation.	Check whether the service status is normal.
Restart or stop services.	This operation may interrupt services.	▲ ▲	Restart or stop services when necessary.	Check whether the service is running properly.
Modify configuration parameters.	This operation requires service restart for configuration to take effect.	▲ ▲	Modify configuration when necessary.	Check whether the service is running properly.
Delete or modify metadata.	Modifying or deleting Kafka metadata on ZooKeeper may cause the Kafka topic or service unavailability.	▲ ▲ ▲	Do not delete or modify Kafka metadata stored on ZooKeeper.	Check whether the Kafka topics or Kafka service is available.
Delete metadata backup files.	After Kafka metadata backup files are modified and used to restore Kafka metadata, Kafka topics or the Kafka service may be unavailable.	▲ ▲ ▲	Do not delete Kafka metadata backup files.	Check whether the Kafka topics or Kafka service is available.

KrbServer High-Risk Operations

Table 12-14 KrbServer high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Modify the **KADMIN_PORT** parameter of KrbServer.	After this parameter is modified, if the KrbServer service and its associated services are not restarted in a timely manner, the configuration of KrbClient in the cluster is abnormal and the service running is affected.	▲ ▲ ▲ ▲ ▲	After this parameter is modified, restart the KrbServer service and all its associated services.	None
Modify the **kdc_ports** parameter of KrbServer.	After this parameter is modified, if the KrbServer service and its associated services are not restarted in a timely manner, the configuration of KrbClient in the cluster is abnormal and the service running is affected.	▲ ▲ ▲ ▲ ▲	After this parameter is modified, restart the KrbServer service and all its associated services.	None
Modify the **KPASSWD_PORT** parameter of KrbServer.	After this parameter is modified, if the KrbServer service and its associated services are not restarted in a timely manner, the configuration of KrbClient in the cluster is abnormal and the service running is affected.	▲ ▲ ▲ ▲ ▲	After this parameter is modified, restart the KrbServer service and all its associated services.	None

Operation	Risk	Severity	Workaround	Check Item
Modify the domain name of Manager system.	After the domain name is modified, if the KrbServer service and its associated services are not restarted in a timely manner, the configuration of KrbClient in the cluster is abnormal and the service running is affected.	▲ ▲ ▲ ▲ ▲	After this parameter is modified, restart the KrbServer service and all its associated services.	None
Configure cross-cluster mutual trust relationships.	This operation will restart the KrbServer service and all associated services, affecting the management and maintenance and services of the cluster.	▲ ▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.

LdapServer High-Risk Operations

Table 12-15 LdapServer high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Modify the **LDAP_SERVER_PORT** parameter of LdapServer.	After this parameter is modified, if the LdapServer service and its associated services are not restarted in a timely manner, the configuration of LdapClient in the cluster is abnormal and the service running is affected.	▲ ▲ ▲ ▲ ▲	After this parameter is modified, restart the LdapServer service and all its associated services.	None

Operation	Risk	Severity	Workaround	Check Item
Restore LdapServer data.	This operation will restart Manager and the entire cluster, affecting management, maintenance, and services of the cluster.	▲ ▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.
Replace the Node where LdapServer is located.	This operation will interrupt services deployed on the node. If the node is a management node, the operation will restart all OMS processes, affecting the cluster management and maintenance.	▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	Check for uncleared alarms, and check whether the cluster management and maintenance and services are normal.
Change the password of LdapServer.	The LdapServer and Kerberos services need to be restarted during the password change, affecting the management, maintenance, and services of the cluster.	▲ ▲ ▲ ▲	Before performing the operation, ensure that the operation is necessary, and that no other management and maintenance operations are performed at the same time.	None

Operation	Risk	Severity	Workaround	Check Item
Restart the node where LdapServer is located.	Restarting the node without stopping the LdapServer service may cause LdapServer data damage.	▲ ▲ ▲ ▲ ▲	Restore LdapServer using LdapServer backup data	None

Loader High-Risk Operations

Table 12-16 Loader high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Change the floating IP address of a Loader instance (**loader.float.ip**).	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether the Loader UI can be connected properly.
Modify the Loader instance start parameter **LOADER_GC_OPTS**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Clear table contents when adding data to HBase.	This operation will clear original data in the target table.	▲ ▲	Ensure that the contents in the target table can be cleared before the operation.	Check whether the contents in the target table can be cleared before the operation.

Spark2x High-risk Operations

NOTE

Spark high-risk operations apply to MRS 3.x earlier versions.

Table 12-17 Spark2x high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Modify the configuration item **spark.yarn.queue**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Modify the configuration item **spark.driver.extraJavaOptions**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Modify the configuration item **spark.yarn.cluster.driver.extraJavaOptions**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.

Operation	Risk	Severity	Workaround	Check Item
Modify the configuration item **spark.eventLog.dir**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Modify the configuration item **SPARK_DAEMON_JAVA_OPTS**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Delete all JobHistory2x instances.	The event logs of historical applications are lost.	▲ ▲	Reserve at least one JobHistory2x instance.	Check whether historical application information is included in JobHistory2x.
Delete or modify the **/user/spark2x/jars/8.1.0.1/spark-archive-2x.zip** file in HDFS.	JDBCServer2x fails to be started and service functions are abnormal.	▲ ▲ ▲	Delete **/user/spark2x/jars/8.1.0.1/spark-archive-2x.zip**, and wait for 10-15 minutes until the .zip package is automatically restored.	Check whether services can be started properly.

Storm High-Risk Operations

Table 12-18 Storm high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Modify the following plug-in related configuration items: • storm.scheduler • nimbus.authorizer • storm.thrift.transport • nimbus.blobstore.class • nimbus.topology.validator • storm.principal.tolocal	Services cannot start properly.	▲ ▲ ▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that the class names exist and are valid.	Check whether services can be started properly.

Operation	Risk	Severity	Workaround	Check Item
Modify the Storm instance **GC_OPTS** startup parameters, including: NIMBUS_GC_OPTS SUPERVISOR_GC_OPTS UI_GC_OPTS LOGVIEWER_GC_OPTS	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.
Modify the user resource pool configuration parameter **resource.aware.scheduler.user.pools**.	Services cannot run properly.	▲ ▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that resources allocated to each user are appropriate and valid.	Check whether services can be started and run properly
Change data directories.	If this operation is not properly performed, services may be abnormal and unavailable.	▲ ▲ ▲ ▲	Do not manually change data directories.	Check whether data directories are normal.
Restart services or instances.	The service will be interrupted for a short period of time, and ongoing operations will be interrupted.	▲ ▲ ▲	Restart services or instances when necessary.	Check whether the service is running properly and whether interrupted operations are restored.

Operation	Risk	Severity	Workaround	Check Item
Synchronize configurations (by restarting the required service).	The service will be restarted, resulting in temporary service interruption. If Supervisor is restarted, ongoing operations will be interrupted for a short period of time.	▲ ▲ ▲	Modify configuration when necessary.	Check whether the service is running properly and whether interrupted operations are restored.
Stop services or instances.	The service will be stopped, and related operations will be interrupted.	▲ ▲ ▲	Stop services when necessary.	Check whether the services are properly stopped.
Delete or modify metadata.	If Nimbus metadata is deleted, services are abnormal and ongoing operations are lost.	▲ ▲ ▲ ▲ ▲	Do not manually delete Nimbus metadata files.	Check whether Nimbus metadata files are normal.
Modify file permissions.	If permissions on the metadata and log directories are incorrectly modified, service exceptions may occur.	▲ ▲ ▲ ▲	Do not manually modify file permissions.	Check whether the permissions on the data and log directories are correct.
Delete topologies.	Topologies in use will be deleted.	▲ ▲ ▲ ▲	Delete topologies when necessary.	Check whether the topologies are successfully deleted.

Yarn High-Risk Operations

Table 12-19 Yarn high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Delete or change data directories **yarn.nodemanager.local-dirs** and **yarn.nodemanager.log-dirs**	This operation may cause service information loss.	▲ ▲ ▲	Do not delete data directories manually.	Check whether data directories are normal.

ZooKeeper High-Risk Operations

Table 12-20 ZooKeeper high-risk operations

Operation	Risk	Severity	Workaround	Check Item
Delete or change ZooKeeper data directories.	This operation may cause service information loss.	▲ ▲ ▲	Follow the capacity expansion guide to change the ZooKeeper data directories.	Check whether services and associated components are started properly.
Modify the ZooKeeper instance start parameter **GC_OPTS**.	Services cannot start properly.	▲ ▲	Strictly follow the prompt information when modifying related configuration items. Ensure that new values are valid.	Check whether services can be started properly.

Operation	Risk	Severity	Workaround	Check Item
Modify the znode ACL information in ZooKeeper.	If znode permission is modified in ZooKeeper, other users may have no permission to access the znode and some system functions are abnormal.	▲ ▲ ▲ ▲	During the modification, strictly follow the ZooKeeper Configuration Guide and ensure that other components can use ZooKeeper properly after ACL information modification.	Check that other components that depend on ZooKeeper can properly start and provide services.

13 Interconnecting Jupyter Notebook with MRS Using Custom Python

13.1 Overview

Configuring Jupyter Notebook in MRS to use Pyspark improves the efficiency of machine learning, data exploration, and ETL application development.

This section describes how to configure Jupyter Notebook in MRS to use Pyspark. The procedure is as follows:

1. **Installing a Client on a Node Outside the Cluster**
2. **Installing Python 3**
3. **Configuring the MRS Client**
4. **Installing Jupyter Notebook**
5. **Verifying that Jupyter Notebook Can Access MRS**

NOTE

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

13.2 Installing a Client on a Node Outside the Cluster

Step 1 Prepare a Linux ECS outside the cluster. For details about the requirements, see **Installing a Client on a Node Outside a Cluster**.

Step 2 Install the client to a directory, for example, **/opt/client**, on the node outside the cluster by referring to **Installing a Client on a Node Outside a Cluster**.

Step 3 Check whether Kerberos authentication is enabled for the cluster.

- If yes, go to **Step 4**.
- If no, go to **Installing Python 3**.

Step 4 Log in to FusionInsight Manager by referring to **Accessing FusionInsight Manager (MRS 3.x or Later)**.

Step 5 Create a user, for example, **mrs-test**. Set **User Group** to **hadoop**, **Primary Group** to **hadoop**, and **Role** to **Manager_operator**.

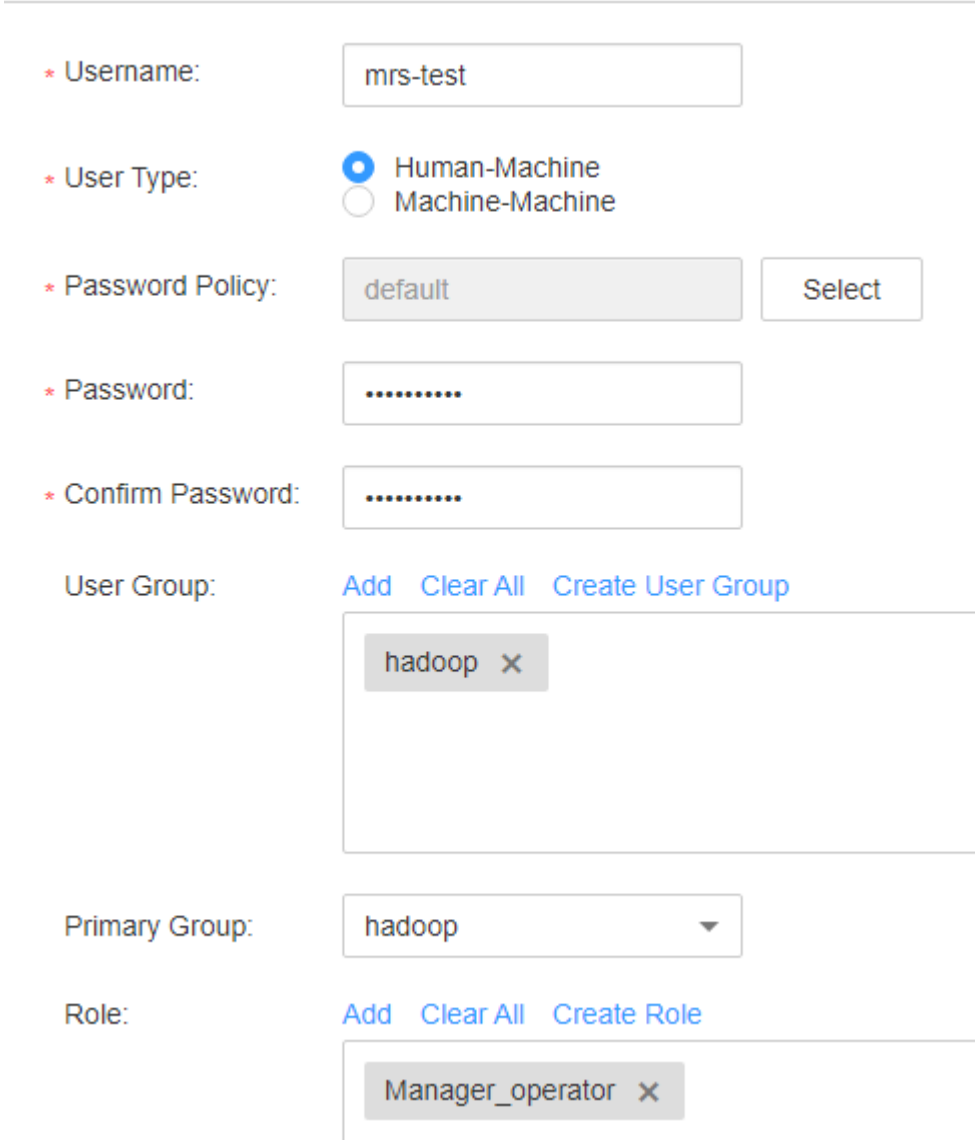

Step 6 Log in to the client node as user **root** and run the following commands to configure environment variables for security authentication:

source /opt/client/bigdata_env

kinit *mrs-test*

NOTE

Change the password upon the first authentication.

----End

13.3 Installing Python 3

Step 1 Log in to the client node outside the cluster as user **root** and run the following command to check whether Python 3 is installed:

python3 --version

```
[root@ecs-notebook FusionInsight_Cluster_1_Services_ClientConfig]# python3 --version
-bash: python3: command not found
```

- If yes, go to **Configuring the MRS Client**.
- If no, go to **Step 2**.

Step 2 Install Python. Python 3.6.6 is used as an example.

1. Run the following commands to install dependencies:

   yum install zlib zlib-devel zip -y

   yum install gcc-c++

   yum install openssl-devel

   yum install sqlite-devel -y

   If the pandas library requires the following dependencies:

   yum install -y xz-devel

   yum install bzip2-devel

2. Run the **wget https://www.python.org/ftp/python/3.6.6/Python-3.6.6.tgz** command to download the source code of Python.
3. Run the following command to decompress the Python source code package, for example, to the **opt** directory:

   cd /opt

   tar –xvf Python-3.6.6.tgz

4. Create a Python installation directory, for example, **/opt/python36**:

   mkdir /opt/python36

5. Compile Python.

   cd /opt/python-3.6.6

   ./configure --prefix=/opt/python36

   The following information is displayed if the commands are executed successfully:

   ```
   configure: creating ./config.status
   config.status: creating Makefile.pre
   config.status: creating Modules/Setup.config
   config.status: creating Misc/python.pc
   config.status: creating Misc/python-config.sh
   config.status: creating Modules/ld_so_aix
   config.status: creating pyconfig.h
   creating Modules/Setup
   creating Modules/Setup.local
   creating Makefile


   If you want a release build with all stable optimizations active (PGO, etc),
   please run ./configure --enable-optimizations
   ```

   Run the **make -j8** command. The following information is displayed if the command is executed successfully:


```
creating build/scripts-3.6
copying and adjusting /tmp/python366/Python-3.6.6/Tools/scripts/pydoc3 -> build/scripts-3.6
copying and adjusting /tmp/python366/Python-3.6.6/Tools/scripts/idle3 -> build/scripts-3.6
copying and adjusting /tmp/python366/Python-3.6.6/Tools/scripts/2to3 -> build/scripts-3.6
copying and adjusting /tmp/python366/Python-3.6.6/Tools/scripts/pyvenv -> build/scripts-3.6
changing mode of build/scripts-3.6/pydoc3 from 644 to 755
changing mode of build/scripts-3.6/idle3 from 644 to 755
changing mode of build/scripts-3.6/2to3 from 644 to 755
changing mode of build/scripts-3.6/pyvenv from 644 to 755
renaming build/scripts-3.6/pydoc3 to build/scripts-3.6/pydoc3.6
renaming build/scripts-3.6/idle3 to build/scripts-3.6/idle3.6
renaming build/scripts-3.6/2to3 to build/scripts-3.6/2to3-3.6
renaming build/scripts-3.6/pyvenv to build/scripts-3.6/pyvenv-3.6
```

Run the **make install** command. The following information is displayed if the command is executed successfully:

```
rm -f /opt/python36/share/man/man1/python3.1
(cd /opt/python36/share/man/man1; ln -s python3.6.1 python3.1)
if test "xupgrade" != "xno"  ; then \
        case upgrade in \
                upgrade) ensurepip="--upgrade" ;; \
                install|*) ensurepip="" ;; \
        esac; \
         ./python -E -m ensurepip \
                $ensurepip --root=/ ; \
fi
Looking in links: /tmp/tmp6ldv525m
Collecting setuptools
Collecting pip
Installing collected packages: setuptools, pip
Successfully installed pip-10.0.1 setuptools-39.0.1
```

6. Run the following commands to configure the Python environment:

   export PYTHON_HOME=/opt/python36

   export PATH=$PYTHON_HOME/bin:$PATH

7. Run the **python3 --version** command. Python has been installed if the following information is displayed:

   ```
   [root@ecs-notebook Python-3.6.6]# python3 --version
   Python 3.6.6
   ```

Step 3 Verify Python 3.

pip3 install helloword

python3

import helloworld

helloworld.say_hello("test")

```
[root@ecs-notebook Python-3.6.6]# pip3 install helloword
Collecting helloword
  Downloading https://files.pythonhosted.org/packages/1b/bf/f0f69f122158e0e98b5d95987a7ef5add3f8a340c6eb78d5871f855ca04e/helloword-0.0.1-py3-none-any.whl
Installing collected packages: helloword
Successfully installed helloword-0.0.1
You are using pip version 10.0.1, however version 21.3.1 is available.
You should consider upgrading via the 'pip install --upgrade pip' command.
[root@ecs-notebook Python-3.6.6]# python3
Python 3.6.6 (default, Dec 15 2021, 06:12:40)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-44)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import helloworld
helloworld.say_hello("test")Hello, Sara!
>>>
'Hello, test!'
>>>
```

Step 4 Install third-party Python libraries, such as pandas and sklearn.

pip3 install pandas


```
[root@ecs-mrs-test Python-3.6.6]# pip3 install pandas
Collecting pandas
  Downloading https://files.pythonhosted.org/packages/c3/e2/00cacecafbab071c787019f00ad84ca3185952f6bb9bca9550ed83870d4d/pandas-1.1.5-cp36-cp36m-man
x86_64.whl (9.5MB)
    100% |                                | 9.5MB 6.5MB/s
Collecting pytz>=2017.2 (from pandas)
  Downloading https://files.pythonhosted.org/packages/60/2e/dec1cc18c51b8df33c7c4d0a321b084cf38e1733b98f9d15018880fb4970/pytz-2022.1-py2.py3-none-an
503kB)
    100% |                                | 512kB 47.2MB/s
Collecting python-dateutil>=2.7.3 (from pandas)
  Downloading https://files.pythonhosted.org/packages/36/7a/87837f39d0296e723bb9b62bbb257d0355c7f6128853c78955f57342a56d/python_dateutil-2.8.2-py2.p
-any.whl (247kB)
    100% |                                | 256kB 54.5MB/s
Collecting numpy>=1.15.4 (from pandas)
  Downloading https://files.pythonhosted.org/packages/45/b2/6c7545bb7a38754d63048c7696804a0d947328125d81bf12beaa692c3ae3/numpy-1.19.5-cp36-cp36m-man
x86_64.whl (13.4MB)
    100% |                                | 13.4MB 4.2MB/s
Collecting six>=1.5 (from python-dateutil>=2.7.3->pandas)
  Downloading https://files.pythonhosted.org/packages/d9/5a/e7c31adbe875f2abbb91bd84cf2dc52d792b5a01506781dbcf25c91daf11/six-1.16.0-py2.py3-none-any
Installing collected packages: pytz, six, python-dateutil, numpy, pandas
Successfully installed numpy-1.19.5 pandas-1.1.5 python-dateutil-2.8.2 pytz-2022.1 six-1.16.0
You are using pip version 10.0.1, however version 21.3.1 is available.
You should consider upgrading via the 'pip install --upgrade pip' command.
```

pip3 install backports.lzma

```
[root@ecs-mrs-test Python-3.6.6]# pip3 install backports.lzma
Collecting backports.lzma
  Using cached https://files.pythonhosted.org/packages/21/0f/1a9990233076d48aa2084100ba289ca162975e73a688f3a56c0ee2bb441a/backports.lzma-0.0.14.tar.gz
Installing collected packages: backports.lzma
  Running setup.py install for backports.lzma ... done
Successfully installed backports.lzma-0.0.14
You are using pip version 10.0.1, however version 21.3.1 is available.
You should consider upgrading via the 'pip install --upgrade pip' command.
```

pip3 install sklearn

```
[root@ecs-mrs-test Python-3.6.6]# pip3 install sklearn
Collecting sklearn
  Downloading https://files.pythonhosted.org/packages/1e/7a/dbb3be0ce9bd5c8b7e3d87328e79063f8b263b2b1bfa4774cb1147bfcd3f/sklearn-0.0.tar.
Collecting scikit-learn (from sklearn)
  Downloading https://files.pythonhosted.org/packages/f5/ef/bcd79e8d59250d6e8478eb1290dc6e05be42b3be8a86e3954146adbc171a/scikit_learn-0.2
ylinux1_x86_64.whl (20.0MB)
    100% |                                | 20.0MB 3.4MB/s
Collecting joblib>=0.11 (from scikit-learn->sklearn)
  Downloading https://files.pythonhosted.org/packages/3e/d5/0163eb0cfa0b673aa4fe1cd3ea9d8a81ea0f32e50807b0c295871e4aab2e/joblib-1.1.0-py2
(306kB)
    100% |                                | 307kB 46.5MB/s
Requirement already satisfied: scipy>=0.19.1 in /root/.local/lib/python3.6/site-packages (from scikit-learn->sklearn) (1.5.4)
Collecting threadpoolctl>=2.0.0 (from scikit-learn->sklearn)
  Downloading https://files.pythonhosted.org/packages/61/cf/6e354304bcb9c6413c4e02a747b600061c21d38ba51e7e544ac7bc66aecc/threadpoolctl-3.
nl
Requirement already satisfied: numpy>=1.13.3 in /opt/python36/lib/python3.6/site-packages (from scikit-learn->sklearn) (1.19.5)
Installing collected packages: joblib, threadpoolctl, scikit-learn, sklearn
  Running setup.py install for sklearn ... done
Successfully installed joblib-1.1.0 scikit-learn-0.24.2 sklearn-0.0 threadpoolctl-3.1.0
You are using pip version 10.0.1, however version 21.3.1 is available.
You should consider upgrading via the 'pip install --upgrade pip' command.
```

Step 5 Run the **python3 -m pip list** command to check the installation result.

```
[root@ecs-mrs-test Python-3.6.6]# python3 -m pip list
Package         Version
--------------- -------
cycler          0.11.0
joblib          1.1.0
kiwisolver      1.3.1
numpy           1.19.5
pandas          1.1.5
pip             10.0.1
pyparsing       3.0.7
python-dateutil 2.8.2
pytz            2022.1
scikit-learn    0.24.2
scipy           1.5.4
setuptools      39.0.1
six             1.16.0
sklearn         0.0
threadpoolctl   3.1.0
```

Step 6 Pack them into **Python.zip**.

cd /opt/python36/

zip -r python36.zip ./*

Step 7 Create an HDFS directory and upload the package to the directory for future use.

hdfs dfs -mkdir /user/python

hdfs dfs -put python36.zip /user/python

----End

13.4 Configuring the MRS Client

Go to **/opt/client/Spark2x/spark/conf** (Spark client installation directory) and configure the following parameters in the **spark-defaults.conf** file:

```
spark.pyspark.driver.python=/usr/bin/python3
spark.yarn.dist.archives=hdfs://hacluster/user/python/python36.zip#Python
```

13.5 Installing Jupyter Notebook

Step 1 Log in to the client node as user **root** and run the following command to install Jupyter Notebook:

pip3 install jupyter notebook

The installation is successful if the following command output is displayed:

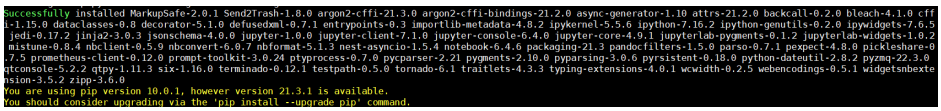

Step 2 To ensure security, you need to generate a ciphertext password for logging in to Jupyter and place it in the configuration file of Jupyter Notebook.

Run the following command and enter the password twice (exit at Out[3]):

ipython

```
[root@ecs-notebook python36]# ipython
Python 3.6.6 (default, Dec 20 2021, 09:32:25)
Type 'copyright', 'credits' or 'license' for more information
IPython 7.16.2 -- An enhanced Interactive Python. Type '?' for help.
In [1]: from notebook.auth import passwd
In [2]: passwd()
Enter password:
Verify password:
Out[2]: 'argon2:$argon2id$v=19$m=10240,t=10,p=8$g14BqLddl927n/unsyPlLQ
$YmoKJzbUfNG7LcxyUzm90bgbKWUIiHy6ZV+ObTzdcA
```

Step 3 Run the following command to generate the Jupyter configuration file:

jupyter notebook --generate-config

Step 4 Modify the configuration file:

vi ~/.jupyter/jupyter_notebook_config.py

Add the following configurations:

```
# -*- coding: utf-8 -*-
c.NotebookApp.ip='*' #Enter the internal IP address of the ECS.
c.NotebookApp.password = u'argon2:$argon2id$v=19$m=10240,t=10,p=8$NmoAVwd8F6vFP2rX5ZbV7w
$SyueJoC0a5TbCuHYzqfSx1vQcFvOTTryR+0uk2MNNZA' # Enter the ciphertext generated at Out[2] in step 2.
c.NotebookApp.open_browser = False # Disable automatic browser opening.
c.NotebookApp.port = 9999 # Specified port number
c.NotebookApp.allow_remote_access = True
```

----End

13.6 Verifying that Jupyter Notebook Can Access MRS

Step 1 Run the following command on the client node to start Jupyter Notebook:

PYSPARK_PYTHON=./Python/bin/python3 PYSPARK_DRIVER_PYTHON=jupyter-notebook PYSPARK_DRIVER_PYTHON_OPTS="--allow-root" pyspark --master yarn --executor-memory 2G --driver-memory 1G

Step 2 Use *EIP*:**9999** to log in to the Jupyter web UI (ensure that the ECS security group allows the local public IP address and port 9999). The login password is the password configured in **Step 2**.

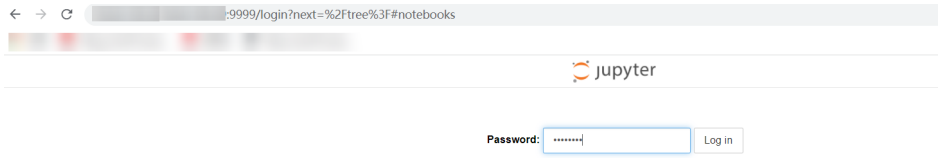

Step 3 Create code.

Create a Python 3 task and use Spark to read files.

The result is as follows:

Log in to FusionInsight Manager and view the submitted PySpark application on the YARN web UI.

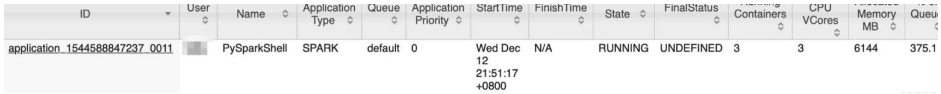

Step 4 Verify that the pandas library can be called.

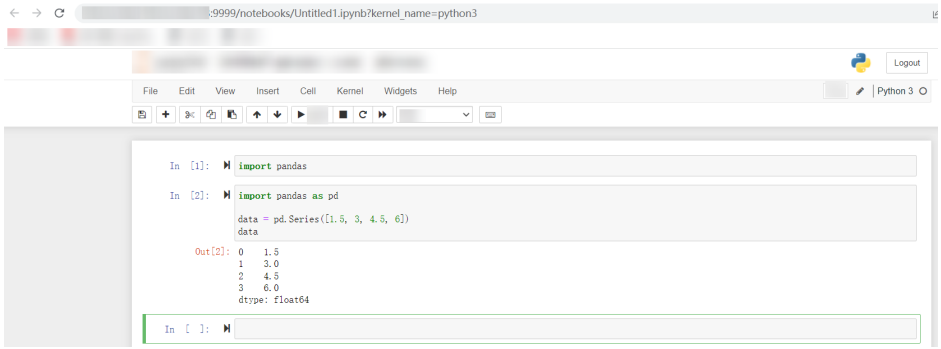

----End

13.7 FAQs

Question

When I import pandas from a local path, the following alarm is generated:

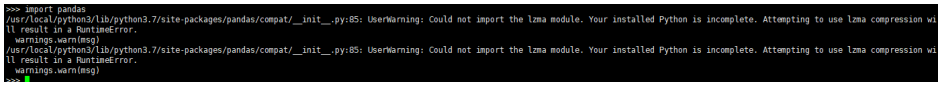

Procedure

Step 1 Run the **python -m pip install backports.lzma** command to install the LZMA module.

```
[root@master ~]# python -m pip install backports.lzma
Looking in indexes: http://mirrors.aliyun.com/pypi/simple/
Requirement already satisfied: backports.lzma in /usr/local/python3/lib/python3.7/site-packages (0.0.14)
You are using pip version 10.0.1, however version 19.3.1 is available.
You should consider upgrading via the 'pip install --upgrade pip' command.
```

Step 2 Go to the **/usr/local/python3/lib/python3.6** directory and edit the **lzma.py** file. The directory varies depending on hosts. You can run the **which** command to query the directory used by Python.

Change

```
from _lzma import *
from _lzma import _encode_filter_properties, _decode_filter_properties
```

To

```
try:
    from _lzma import *
    from _lzma import _encode_filter_properties, _decode_filter_properties
except ImportError:
    from backports.lzma import *
    from backports.lzma import _encode_filter_properties, _decode_filter_properties
```

Before modification


```
"""Interface to the liblzma compression library.

This module provides a class for reading and writing compressed files,
classes for incremental (de)compression, and convenience functions for
one-shot (de)compression.

These classes and functions support both the XZ and legacy LZMA
container formats, as well as raw compressed data streams.
"""

__all__ = [
    "CHECK_NONE", "CHECK_CRC32", "CHECK_CRC64", "CHECK_SHA256",
    "CHECK_ID_MAX", "CHECK_UNKNOWN",
    "FILTER_LZMA1", "FILTER_LZMA2", "FILTER_DELTA", "FILTER_X86", "FILTER_IA64",
    "FILTER_ARM", "FILTER_ARMTHUMB", "FILTER_POWERPC", "FILTER_SPARC",
    "FORMAT_AUTO", "FORMAT_XZ", "FORMAT_ALONE", "FORMAT_RAW",
    "MF_HC3", "MF_HC4", "MF_BT2", "MF_BT3", "MF_BT4",
    "MODE_FAST", "MODE_NORMAL", "PRESET_DEFAULT", "PRESET_EXTREME",

    "LZMACompressor", "LZMADecompressor", "LZMAFile", "LZMAError",
    "open", "compress", "decompress", "is_check_supported",
]

import builtins
import io
import os
from _lzma import *
from _lzma import _encode_filter_properties, _decode_filter_properties
import _compression
```

After modification

```
These classes and functions support both the XZ and legacy LZMA
container formats, as well as raw compressed data streams.
"""

__all__ = [
    "CHECK_NONE", "CHECK_CRC32", "CHECK_CRC64", "CHECK_SHA256",
    "CHECK_ID_MAX", "CHECK_UNKNOWN",
    "FILTER_LZMA1", "FILTER_LZMA2", "FILTER_DELTA", "FILTER_X86", "FILTER_IA64",
    "FILTER_ARM", "FILTER_ARMTHUMB", "FILTER_POWERPC", "FILTER_SPARC",
    "FORMAT_AUTO", "FORMAT_XZ", "FORMAT_ALONE", "FORMAT_RAW",
    "MF_HC3", "MF_HC4", "MF_BT2", "MF_BT3", "MF_BT4",
    "MODE_FAST", "MODE_NORMAL", "PRESET_DEFAULT", "PRESET_EXTREME",

    "LZMACompressor", "LZMADecompressor", "LZMAFile", "LZMAError",
    "open", "compress", "decompress", "is_check_supported",
]

import builtins
import io
import os
#from _lzma import *
#from _lzma import _encode_filter_properties, _decode_filter_properties
try:
    from _lzma import *
    from _lzma import _encode_filter_properties, _decode_filter_properties
except ImportError:
    from backports.lzma import *
    from backports.lzma import _encode_filter_properties, _decode_filter_properties
import _compression
```

Step 3 Save the file and exit. Then, import pandas again.

```
[root@master python3.7]# python
Python 3.7.0 (default, Oct 26 2019, 01:19:22)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-36)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import pandas
>>>
```

----End

14 Appendix

14.1 ECS Specifications Used by MRS

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

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

ECS Flavor Naming Rules

AB.C.D

Example: m2.8xlarge.8

In the preceding flavor:

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

Specifications

Table 14-1 General computing-plus (C) ECS specifications

Type	vCPU	Memory (GB)	Flavor	Virtualization Type
C3	32	64	c3.8xlarge.2	KVM
C3	16	64	c3.4xlarge.4	KVM
C3	32	128	c3.8xlarge.4	KVM
C3	60	256	c3ne.15xlarge.4	KVM
C3ne	32	64	c3ne.8xlarge.2	KVM
C3ne	16	64	c3ne.4xlarge.4	KVM
C3ne	32	128	c3ne.8xlarge.4	KVM
C3ne	60	256	c3ne.15xlarge.4	KVM
C6	32	64	c6.8xlarge.2	KVM
C6	64	128	c6.16xlarge.2	KVM
C6	16	64	c6.4xlarge.4	KVM
C6	32	128	c6.8xlarge.4	KVM
C6	64	256	c6.16xlarge.4	KVM
C6s	32	64	c6s.8xlarge.2	KVM
C6s	64	128	c6s.16xlarge.2	KVM

Table 14-2 Memory-optimized ECS specifications

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

Type	vCPU	Memory (GB)	Flavor	Virtualization Type
M6	64	512	m6.16xlarge.8	KVM

Table 14-3 Ultra-high I/O ECS specifications

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

14.2 BMS Specifications Used by MRS

MRS uses the following BMS in different application scenarios.

Kunpeng V1 BMS

Specifications

Table 14-4 Kunpeng V1 BMS specifications

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

14.3 A Defect Exists After Core Nodes in the MRS Cluster Are Added

Symptom

Cores nodes are added, but some instances on the nodes may fail to be started. The symptoms are as follows:

1. A core node has been added and is displayed on the **Nodes** page.

2. Some tasks for adding nodes fail or are partially successful.

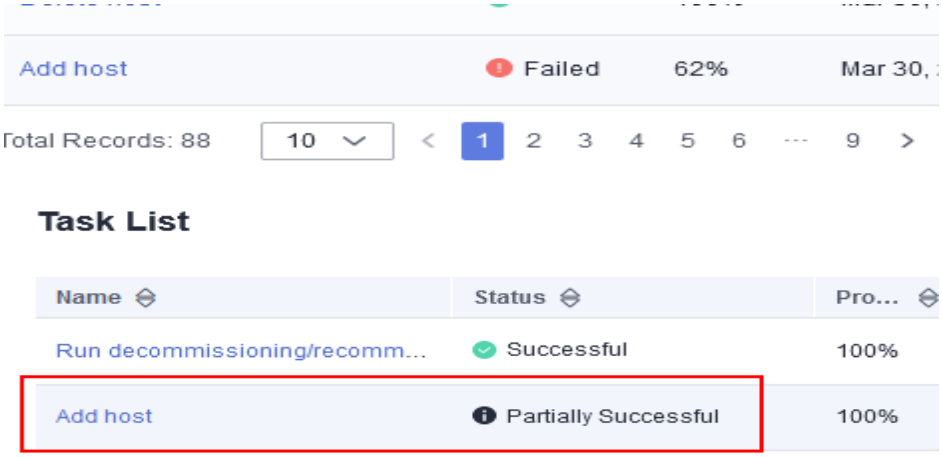

3. If IAM users have been synchronized, you can view unstarted roles on the **Components** page.
4. If they are not synchronized, you can view unstarted roles on the Manager page of this cluster.

Procedure

Scenario 1: The task for adding nodes fails before component installation.

Step 1 Perform the following steps if the MRS cluster is a pay-per-use cluster:

1. Log in to the MRS console.
2. Choose **Active Clusters** and click the cluster name to go to the cluster details page.

3. Click in the upper part of the page. In the **Task List** column, click the task for adding core nodes.
4. Records all nodes in the verification request parameter.
5. Click the **Nodes** tab, select the nodes recorded in **Step 1.4**, click **Stop** in the upper right corner, and stop the nodes as prompted.
6. Reduce nodes by referring to **Scaling In a Cluster**.

Step 2 If the MRS cluster is billed on a yearly/monthly basis, unsubscribe from the abnormal nodes by referring to **Unsubscribing from a Specified Node in a Yearly/Monthly Cluster**.

----End

Scenario 2: The task for adding nodes fails after component installation.

Step 1 Log in to the MRS console.

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

Step 3 On the **Dashboard** tab, click **Synchronize** next to **IAM User Sync** to synchronize IAM users.

Step 4 Click **Components** and check the role status of each service. If a role is not started, select the role, click **More**, and select **Start Instance** to start the instance.

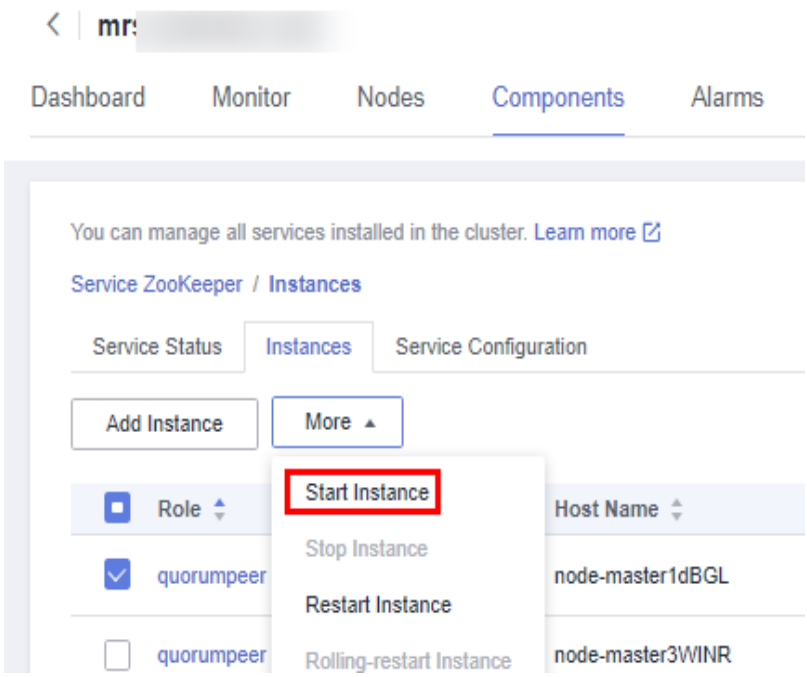

Step 5 If the startup fails, rectify the fault based on the error information in the task list and try again.

NOTE

- If there are many abnormal roles, click **Management Operations** in the upper right corner to start all components.
- For other exceptions that cannot be resolved, contact technical support.
- You can also start the instance on the Manager page of the cluster. For details, see **Overview**

----End

14.4 Data Migration Solution

14.4.1 Making Preparations

This section describes how to migrate HDFS, HBase, and Hive data to an MRS cluster in different scenarios. During data migration, data may be overwritten, lost, or damaged. This document is for reference only. Please cooperate with Huawei Cloud technical personnel to formulate and implement a specific data migration solution.

Make preparations on a source cluster before data migration to prevent the source cluster from generating new data during data migration, thereby avoiding data inconsistency between the source and destination clusters after data migration. Before data migration is complete, the destination cluster must be in the initial state and cannot run any other services except data migration jobs.

Stopping Cluster Services and the Related Services

- If the Kafka service is involved in your cluster, stop all jobs that generate data in Kafka. Wait until the Kafka consumption tasks have consumed the inventory data in Kafka, and then perform the next step.
- Stop all services and jobs related to HDFS, HBase, and Hive, and stop the HBase and Hive services.

Establishing a Data Transmission Channel

- If the source cluster and destination cluster are deployed in different VPCs in the same region, create a network connection between the two VPCs to establish a data transmission channel at the network layer. For details, see **VPC Peering Connection Overview**.
- If the source cluster and destination cluster are deployed in the same VPC but belong to different security groups, add security group rules to each security group on the VPC management console. In the security rules, **Protocol** is set to **ANY**, **Transfer Direction** is set to **Inbound**, and **Source** is set to **Security Group** (the security group of the peer cluster).
  - To add an inbound rule to the security group of the source cluster, select the security group of the destination cluster in **Source**.
  - To add an inbound rule to the security group of the destination cluster, select the security group of the source cluster in **Source**.
- If the source and destination clusters are deployed in the same security group of the same VPC and Kerberos authentication is enabled for both clusters, you need to configure mutual trust between the two clusters.

14.4.2 Exporting Metadata

To ensure that the data properties and permissions of the source cluster are consistent with those of the destination cluster, metadata of the source cluster needs to be exported to restore metadata after data migration. The metadata to be exported includes the owner, group, and permission information of the HDFS files and Hive table description.

Exporting HDFS Metadata

HDFS metadata information to be exported includes file and folder permissions and owner/group information. You can run the following command on the HDFS client to export the metadata:

```
$HADOOP_HOME/bin/hdfs dfs –ls –R <migrating_path> > /tmp/hdfs_meta.txt
```

The following provides description about the parameters in the preceding command.

- **$HADOOP_HOME**: installation directory of the Hadoop client in the source cluster
- **<migrating_path>**: HDFS data directory to be migrated
- **/tmp/hdfs_meta.txt**: local path for storing the exported metadata

NOTE

If the source cluster can communicate with the destination cluster and you run the **hadoop distcp** command as a super administrator to copy data, you can add the **-p** parameter to enable DistCp to restore the metadata of the corresponding file in the destination cluster while copying data. In this case, skip this step.

Exporting Hive Metadata

Hive table data is stored in HDFS. Table data and the metadata of the table data is centrally migrated in directories by HDFS in a unified manner. Metadata of Hive tables can be stored in different types of relational databases (such as MySQL, PostgreSQL, and Oracle) based on cluster configurations. The exported metadata of the Hive tables in this document is the Hive table description stored in the relational database.

The mainstream big data release editions in the industry support Sqoop installation. For on-premises big data clusters of the community version, you can download the Sqoop of the community version for installation. Use Sqoop to decouple the strong dependency between the metadata to be exported and the relational database and export Hive metadata to HDFS and migrate it together with the table data for restoration. The procedure is as follows:

Step 1 Download the Sqoop tool from the source cluster and install it. For details, see **http://sqoop.apache.org/**.

Step 2 Download the JDBC driver of the relational database to the **${Sqoop_Home}/lib** directory.

Step 3 Run the following command to export all Hive metadata tables: All exported data is stored in the **/user/**<*user_name*>/<*table_name*> directory on HDFS.

```
$Sqoop_Home/bin/sqoop import --connect jdbc:<driver_type>://<ip>:<port>/<database> --table
<table_name> --username <user> -password <passwd> -m 1
```


The following provides description about the parameters in the preceding command.

- *$Sqoop_Home*: Sqoop installation directory
- <*driver_type*>: Database type
- **<ip>**: IP address of the database in the source cluster
- **<port>**: Port number of the database in the source cluster
- *<table_name>*: Name of the table to be exported
- <*user*>: Username
- <*passwd*>: User password

NOTE

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

----End

14.4.3 Copying Data

Based on the regions of and network connectivity between the source cluster and destination cluster, data copy scenarios are classified as follows:

Same Region

If the source cluster and destination cluster are in the same region, follow instructions in **Establishing a Data Transmission Channel** to configure the network and set up a network transmission channel. Use the DistCp tool to run the following command to copy the HDFS, HBase, Hive data files and Hive metadata backup files from the source cluster to the destination cluster.

```
$HADOOP_HOME/bin/hadoop distcp <src> <dist> -p
```

The following provides description about the parameters in the preceding command.

- **$HADOOP_HOME**: installation directory of the Hadoop client in the destination cluster
- **<src>**: HDFS directory of the source cluster
- **<dist>**: HDFS directory of the destination cluster

Migrating Data from an Offline Cluster to a Cloud

You can use the following way to migrate data from an offline cluster to the cloud.

- Direct Connect

  Create a Direct Connect between the source cluster and target cluster, enable the network between the offline cluster egress gateway and the online VPC, and use DistCp to copy the data by referring to **Same Region**.

14.4.4 Restoring Data

HDFS File Property Restoration

Based on the exported permission information, run the HDFS commands in the background of the destination cluster to restore the file permission and owner and group information.

```
$HADOOP_HOME/bin/hdfs dfs –chmod <MODE> <path>
$HADOOP_HOME/bin/hdfs dfs –chown <OWNER> <path>
```

Hive Metadata Restoration

Install Sqoop and run the Sqoop command in the destination cluster to import the exported Hive metadata to DBService in the MRS cluster.

```
$Sqoop_Home/bin/sqoop export --connect jdbc:postgresql://<ip>:20051/hivemeta --table <table_name> --username hive -password <passwd> --export-dir <export_from>
```

The following provides description about the parameters in the preceding command.

- *$Sqoop_Home*: Sqoop installation directory in the destination cluster
- *<ip>*: IP address of the database in the destination cluster
- *<table_name>*: Name of the table to be restored
- *<passwd>*: Password of user **hive**
- *<export_from>*: HDFS address of the metadata in the destination cluster

NOTE

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

HBase Table Reconstruction

Restart the HBase service of the destination cluster to make data migration take effect. During the restart, HBase loads the data in the current HDFS and regenerates metadata. After the restart is complete, run the following command on the Master node client to load the HBase table data:

```
$HBase_Home/bin/hbase hbck -fixMeta -fixAssignments
```

After the command is executed, run the following command repeatedly to check the health status of the HBase cluster until the health status is normal:

```
hbase hbck
```

14.5 Precautions for MRS 3.x

Purpose

Custers of versions earlier than MRS 3.x use MRS Manager to manage and monitor MRS clusters. On the Cluster Management page of the MRS management console, you can view cluster details, manage nodes, components, alarms, patches,

files, jobs, tenants, and backup and restoration. In addition, you can configure Bootstrap actions and manage tags.

MRS 3.x uses FusionInsight Manager to manage and monitor clusters. On the Cluster Management page of the MRS management console, you can view cluster details, manage nodes, components, alarms, files, jobs, Bootstrap actions, and tags.

Some maintenance operations of the MRS 3.x cluster are different from those of earlier versions. For details, see **MRS Manager Operation Guide (Applicable to 2.x and Earlier Versions)** and **FusionInsight Manager Operation Guide (Applicable to 3.x)**.

Accessing MRS Manager

- For details about how to access MRS Manager of versions earlier than MRS 3.x, see **Accessing MRS Manager (MRS 2.x or Earlier)**.
- For details about how to access FusionInsight Manager of MRS 3.x, see **Accessing FusionInsight Manager (MRS 3.x or Later)**.

Modifying MRS Cluster Service Configuration Parameters

- For versions earlier than MRS 3.x, you can modify service configuration parameters on the cluster management page of the MRS management console.
  a. Log in to the MRS console. In the left navigation pane, choose **Clusters** > **Active Clusters**, and click a cluster name.
  b. Choose **Components** > *Name of the desired service* > **Service Configuration**.

     The **Basic Configurations** tab page is displayed by default. To modify more parameters, click the **All Configurations** tab. The navigation tree displays all configuration parameters of the service. The level-1 nodes in the navigation tree are service names or role names. The parameter category is displayed after the level-1 node is expanded.
  c. In the navigation tree, select the specified parameter category and change the parameter values on the right.

     If you are not sure about the location of a parameter, you can enter the parameter name in search box in the upper right corner. The system searches for the parameter in real time and displays the result.
  d. Click **Save Configuration**. In the displayed dialog box, click **OK**.
  e. Wait until the message "Operation succeeded" is displayed. Click **Finish**. The configuration is modified.

     Check whether there is any service whose configuration has expired in the cluster. If yes, restart the corresponding service or role instance for the configuration to take effect. You can also select **Restart the affected services or instances** when saving the configuration.
- In MRS 3.x, you need to log in to FusionInsight Manager to modify service configuration parameters.
  a. Log in to FusionInsight Manager.
  b. Choose **Cluster** > **Services**.

c. Click the specified service name on the service management page.

d. Click **Configurations**.

The **Basic Configurations** tab page is displayed by default. To modify more parameters, click the **All Configurations** tab. The navigation tree displays all configuration parameters of the service. The level-1 nodes in the navigation tree are service names or role names. The parameter category is displayed after the level-1 node is expanded.

e. In the navigation tree, select the specified parameter category and change the parameter values on the right.

If you are not sure about the location of a parameter, you can enter the parameter name in search box in the upper right corner. The Manager searches for the parameter in real time and displays the result.

f. Click **Save**. In the confirmation dialog box, click **OK**.

g. Wait until the message "Operation succeeded" is displayed. Click **Finish**. The configuration is modified.

Check whether there is any service whose configuration has expired in the cluster. If yes, restart the corresponding service or role instance for the configuration to take effect.

14.6 Installing the Flume Client

14.6.1 Installing the Flume Client on Clusters of Versions Earlier Than MRS 3.x

Scenario

To use Flume to collect logs, you must install the Flume client on a log host. You can create an ECS and install the Flume client on it.

This section applies to MRS 3.*x* or earlier clusters.

Prerequisites

- A streaming cluster with the Flume component has been created.
- The log host is in the same VPC and subnet with the MRS cluster.
- You have obtained the username and password for logging in to the log host.

Procedure

Step 1 Create an ECS that meets the requirements.

Step 2 Go to the cluster details page and choose **Components**.

NOTE

If the **Components** tab is unavailable, complete IAM user synchronization first. (On the **Dashboard** page, click **Synchronize** on the right side of **IAM User Sync** to synchronize IAM users.)

Step 3 Click **Download Client**.

1. In **Client Type**, select **All client files**.
2. In **Download to**, select **Remote host**.
3. Set **Host IP Address** to the IP address of the ECS, **Host Port** to **22**, and **Save Path** to **/tmp**.
   - If the default port **22** for logging in to an ECS through SSH has been changed, set **Host Port** to a new port.
   - The value of **Save Path** contains a maximum of 256 characters.
4. Set **Login User** to **root**.

   If another user is used, ensure that the user has permissions to read, write, and execute the save path.
5. Select **Password** or **SSH Private Key** for **Login Mode**.
   - **Password**: Enter the password of user **root** set during cluster creation.
   - **SSH Private Key**: Select and upload the key file used for creating the cluster.
6. Click **OK** to generate a client file.

   If the following information is displayed, the client package is saved.

   Client files downloaded to the remote host successfully.

   If the following information is displayed, check the username, password, and security group configurations of the remote host. Ensure that the username and password are correct and an inbound rule of the SSH (22) port has been added to the security group of the remote host. And then, go to **Step 3** to download the client again.

   Failed to connect to the server. Please check the network connection or parameter settings.

Figure 14-1 Downloading a client

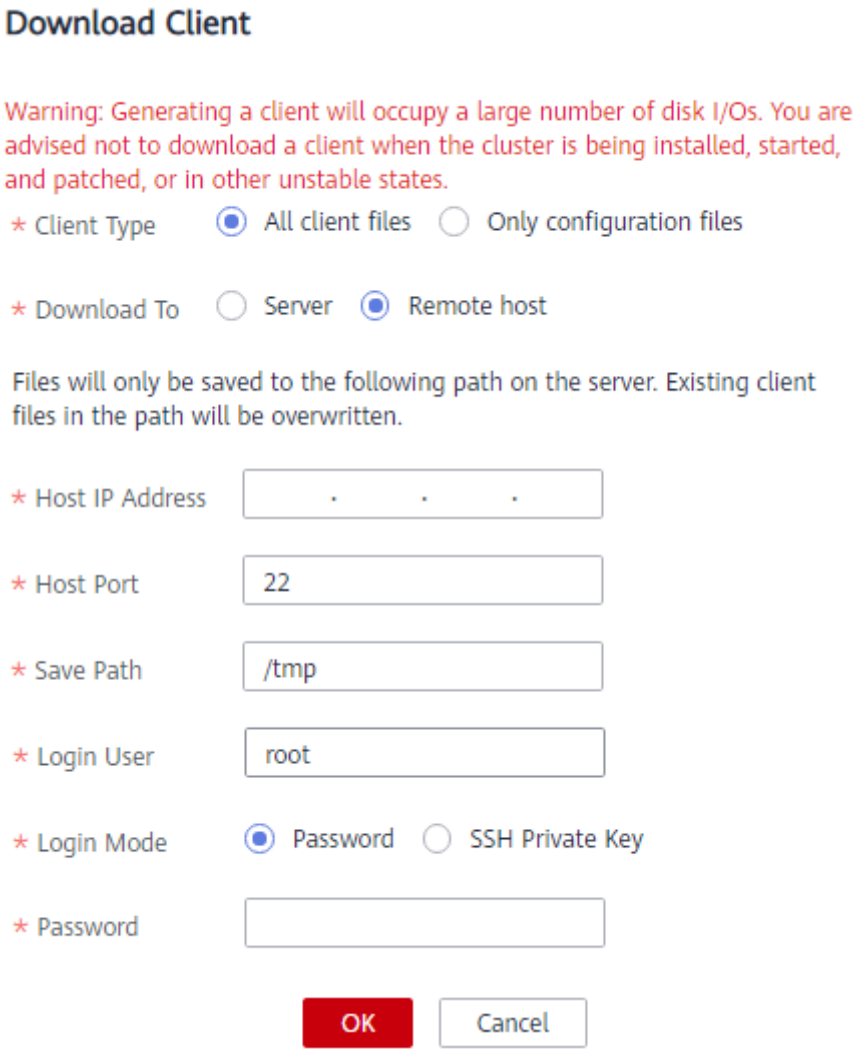

Step 4 Choose **Flume** > **Instance**. Query the **Business IP Address** of any Flume instance and any two MonitorServer instances.

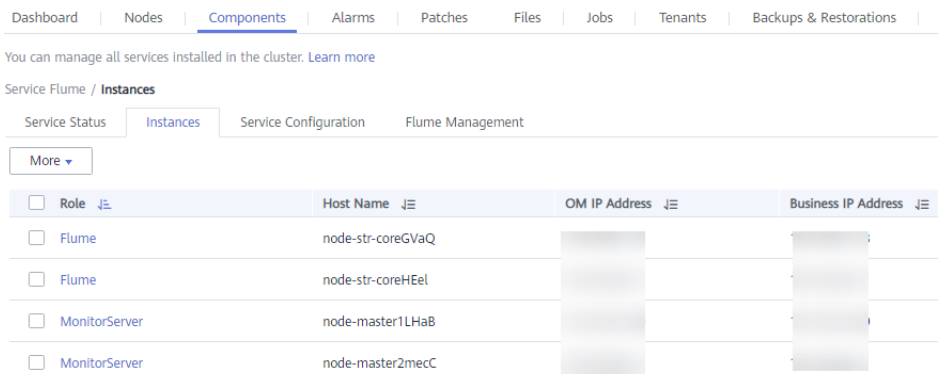

Step 5 Log in to the ECS using VNC. For details, see "**User Guide**" > "**Login Using VNC**" in *Elastic Cloud Server*. ("**Instances**" > "**Logging In to a Linux ECS**").

All images support Cloud-Init. The preset username for Cloud-Init is **root** and the password is the one you set during cluster creation. You are advised to change the password upon the first login.

Step 6 On the ECS, switch to user **root** and copy the installation package to the **/opt** directory.

sudo su - root

cp /tmp/MRS_Flume_Client.tar /opt

Step 7 Run the following command in the **/opt** directory to decompress the package and obtain the verification file and the configuration package of the client:

tar -xvf MRS_Flume_Client.tar

Step 8 Run the following command to verify the configuration package of the client:

sha256sum -c MRS_Flume_ClientConfig.tar.sha256

If the following information is displayed, the file package is successfully verified:

```
MRS_Flume_ClientConfig.tar: OK
```

Step 9 Run the following command to decompress **MRS_Flume_ClientConfig.tar**:

tar -xvf MRS_Flume_ClientConfig.tar

Step 10 Run the following command to install the client running environment to a new directory, for example, **/opt/Flumeenv**. A directory is automatically generated during the client installation.

sh /opt/MRS_Flume_ClientConfig/install.sh /opt/Flumeenv

If the following information is displayed, the client running environment is successfully installed:

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

Step 11 Run the following command to configure environment variables:

source /opt/Flumeenv/bigdata_env

Step 12 Run the following commands to decompress the Flume client package:

cd /opt/MRS_Flume_ClientConfig/Flume

tar -xvf FusionInsight-Flume-1.6.0.tar.gz

Step 13 Run the following command to check whether the password of the current user has expired:

chage -l root

If the value of **Password expires** is earlier than the current time, the password has expired. Run the **chage -M -1 root** command to validate the password.

Step 14 Run the following command to install the Flume client to a new directory, for example, **/opt/FlumeClient**. A directory is automatically generated during the client installation.

sh /opt/MRS_Flume_ClientConfig/Flume/install.sh -d /opt/FlumeClient -f *service IP address of the MonitorServer instance* **-c** *path of the Flume*

configuration file **-l /var/log/ -e** *service IP address of Flume* **-n** *name of the Flume client*

The parameters are described as follows:

- **-d**: indicates the installation path of the Flume client.
- (Optional) **-f**: indicates the service IP addresses of the two MonitorServer instances, separated by a comma (,). If the IP addresses are not configured, the Flume client will not send alarm information to MonitorServer, and the client information will not be displayed on MRS Manager.
- (Optional) **-c**: indicates the **properties.properties** configuration file that the Flume client loads after installation. If this parameter is not specified, the **fusioninsight-flume-1.6.0/conf/properties.properties** file in the client installation directory is used by default. The configuration file of the client is empty. You can modify the configuration file as required and the Flume client will load it automatically.
- (Optional) **-l**: indicates the log directory. The default value is **/var/log/Bigdata**.
- (Optional) **-e**: indicates the service IP address of the Flume instance. It is used to receive the monitoring indicators reported by the client.
- (Optional) **-n**: indicates the name of the Flume client.
- IBM JDK does not support **-Xloggc**. You must change **-Xloggc** to **-Xverbosegclog** in **flume/conf/flume-env.sh**. For 32-bit JDK, the value of **-Xmx** must not exceed 3.25 GB.
- In **flume/conf/flume-env.sh**, the default value of **-Xmx** is 4 GB. If the client memory is too small, you can change it to 512 MB or even 1 GB.

For example, run **sh install.sh -d /opt/FlumeClient**.

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

```
install flume client successfully.
```

----End

14.6.2 Installing the Flume Client on MRS 3.*x* or Later Clusters

Scenario

To use Flume to collect logs, you must install the Flume client on a log host. You can create an ECS and install the Flume client on it.

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

Prerequisites

- A cluster with the Flume component has been created.
- The log host is in the same VPC and subnet with the MRS cluster.
- You have obtained the username and password for logging in to the log host.
- The installation directory is automatically created if it does not exist. If it exists, the directory must be left blank. The directory path cannot contain any space.

Procedure

Step 1 Obtain the software package.

Log in to the FusionInsight Manager. Choose **Cluster** > *Name of the target cluster* > **Services** > **Flume**. On the Flume service page that is displayed, choose **More** > **Download Client** in the upper right corner and set **Select Client Type** to **Complete Client** to download the Flume service client file.

The file name of the client is **FusionInsight_Cluster_**<*Cluster ID*>**_Flume_Client.tar**. This section takes the client file **FusionInsight_Cluster_1_Flume_Client.tar** as an example.

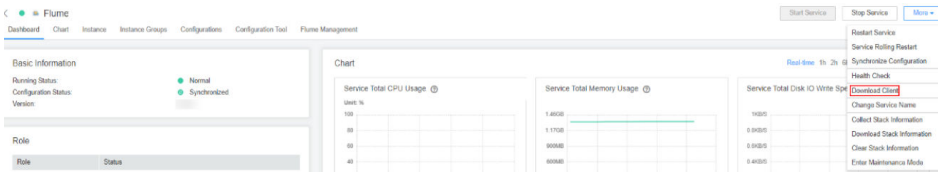

Step 2 Upload the software package.

Upload the software package to a directory, for example, **/opt/client**, on the node where the Flume client is to be installed as user **user**.

NOTE

user is the user who installs and runs the Flume client.

Step 3 Decompress the software package.

Log in to the node where the Flume service client is to be installed as user **user**. Go to the directory where the installation package is installed, for example, **/opt/client**, and run the following command to decompress the installation package to the current directory:

cd /opt/client

tar -xvf FusionInsight_Cluster_1_Flume_Client.tar

Step 4 Verify the software package.

Run the **sha256sum -c** command to verify the decompressed file. If **OK** is returned, the verification is successful. Example:

sha256sum -c FusionInsight_Cluster_1_Flume_ClientConfig.tar.sha256

```
FusionInsight_Cluster_1_Flume_ClientConfig.tar: OK
```

Step 5 Decompress the package.

tar -xvf FusionInsight_Cluster_1_Flume_ClientConfig.tar

Step 6 To install the Flume client on a node outside the cluster, perform the following steps to configure the installation environment. Skip this step if you install Flume client on a node in the cluster.

1. Run the following command to install the client running environment to a new directory, for example, **/opt/Flumeenv**. A directory is automatically generated during the client installation.

sh /opt/client/FusionInsight_Cluster_1_Flume_ClientConfig/install.sh /opt/Flumeenv

If the following information is displayed, the client running environment is successfully installed:

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

2. Run the following command to set environment variables:

   source /opt/Flumeenv/bigdata_env

Step 7 Run the following command in the Flume client installation directory to install the client to a specified directory (for example, **opt/FlumeClient**): After the client is installed successfully, the installation is complete.

cd /opt/client/FusionInsight_Cluster_1_Flume_ClientConfig/Flume/FlumeClient

./install.sh -d */opt/FlumeClient* **-f** *MonitorServerService IP address or host name of the role* **-c** *User service configuration filePath for storing properties.properties* **-s** *CPU threshold* **-l /var/log/Bigdata -e** *FlumeServer service IP address or host name* **-n** *Flume*

NOTE

- **-d**: Flume client installation path
- (Optional) **-f**: IP addresses or host names of two MonitorServer roles. The IP addresses or host names are separated by commas (,). If this parameter is not configured, the Flume client does not send alarm information to MonitorServer and information about the client cannot be viewed on the FusionInsight Manager GUI.
- (Optional) **-c**: Service configuration file, which needs to be generated on the configuration tool page of the Flume server based on your service requirements. Upload the file to any directory on the node where the client is to be installed. If this parameter is not specified during the installation, you can upload the generated service configuration file **properties.properties** to the **/opt/FlumeClient/fusioninsight-flume-1.9.0/conf** directory after the installation.
- (Optional) **-s**: cgroup threshold. The value is an integer ranging from 1 to 100 x *N*. *N* indicates the number of CPU cores. The default threshold is **-1**, indicating that the processes added to the cgroup are not restricted by the CPU usage.
- (Optional) **-l**: Log path. The default value is **/var/log/Bigdata**. The user **user** must have the write permission on the directory. When the client is installed for the first time, a subdirectory named **flume-client** is generated. After the installation, subdirectories named **flume-client-***n* will be generated in sequence. The letter *n* indicates a sequence number, which starts from 1 in ascending order. In the **/conf/** directory of the Flume client installation directory, modify the **ENV_VARS** file and search for the **FLUME_LOG_DIR** attribute to view the client log path.
- (Optional) **-e**: Service IP address or host name of FlumeServer, which is used to receive statistics for the monitoring indicator reported by the client.
- (Optional) **-n**: Name of the Flume client. You can choose **Cluster** > *Name of the desired cluster* > **Service** > **Flume** > **Flume Management** on FusionInsight Manager to view the client name on the corresponding node.
- If the following error message is displayed, run the **export JAVA_HOME=*JDK path*** command. You can run the **echo $JAVA_HOME** command to query the JDK path.
  JAVA_HOME is null in current user,please install the JDK and set the JAVA_HOME
- IBM JDK does not support **-Xloggc**. You must change **-Xloggc** to **-Xverbosegclog** in **flume/conf/flume-env.sh**. For 32-bit JDK, the value of **-Xmx** must not exceed 3.25 GB.
- When installing a cross-platform client in a cluster, go to the **/opt/client/FusionInsight_Cluster_1_Flume_ClientConfig/Flume/FusionInsight-Flume-1.9.0.tar.gz** directory to install the Flume client.

----End

