

# MapReduce Service

## API Reference

**Issue** 10  
**Date** 2024-08-29



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

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

## **Trademarks and Permissions**



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

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

## **Notice**

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

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

---

# Contents

---

<b>1 Before You Start.....</b>	<b>1</b>
<b>2 API Overview.....</b>	<b>4</b>
<b>3 Selecting an API Type.....</b>	<b>9</b>
<b>4 Calling APIs.....</b>	<b>10</b>
4.1 Making an API Request.....	10
4.2 Authentication.....	13
4.3 Response.....	14
<b>5 Application Cases.....</b>	<b>16</b>
5.1 Creating an MRS Cluster.....	16
5.2 Scaling Out a Cluster.....	18
5.3 Scaling in a Cluster.....	19
5.4 Creating a Job.....	20
5.5 Terminating a Job.....	21
5.6 Terminating a Cluster.....	22
<b>6 API V2.....</b>	<b>24</b>
6.1 Cluster Management APIs.....	24
6.1.1 Creating a Cluster.....	24
6.1.2 Changing a Cluster Name.....	63
6.1.3 Creating a Cluster and Submitting a Job.....	65
6.1.4 Scaling Out a Cluster.....	94
6.1.5 Scaling In a Cluster.....	96
6.1.6 Adding Components to a Cluster.....	98
6.1.7 Querying the Cluster Node List.....	102
6.2 Job Management APIs.....	108
6.2.1 Adding and Executing a Job.....	108
6.2.2 Querying Information About a Job.....	114
6.2.3 Querying a List of Jobs.....	119
6.2.4 Terminating a Job.....	125
6.2.5 Obtaining SQL Results.....	127
6.2.6 Deleting Jobs in Batches.....	129
6.3 Auto Scaling APIs.....	130

6.3.1 Viewing Auto Scaling Policies.....	130
6.3.2 Updating an Auto Scaling Policy.....	136
6.3.3 Deleting an AS policy.....	142
6.3.4 Creating an AS policy.....	144
6.4 Cluster HDFS File API.....	150
6.4.1 Obtaining the List of Files from a Specified Directory.....	150
6.5 SQL APIs.....	154
6.5.1 Submitting a SQL Statement.....	154
6.5.2 Querying SQL Results.....	157
6.5.3 Canceling a SQL Execution Task.....	160
6.6 Agency Management.....	162
6.6.1 Querying the Mapping Between a User (Group) and an IAM Agency.....	162
6.6.2 Updating the Mapping Between a User (Group) and an IAM Agency.....	165
6.7 Data Connection Management.....	168
6.7.1 Creating a Data Connection.....	168
6.7.2 Querying the Data Connection List.....	170
6.7.3 Updating a Data Connection.....	173
6.7.4 Deleting a Data Connection.....	175
6.8 Querying Version Metadata.....	175
6.8.1 Obtaining MRS Version List.....	176
6.8.2 Querying Available Specifications of an MRS Cluster Version.....	177
6.9 IAM Synchronization.....	180
6.9.1 Obtaining Synchronized IAM Users and User Groups.....	180
6.9.2 Synchronizing an IAM User and User Group.....	181
6.9.3 Cancelling Synchronization of Specified Users and User Groups.....	183
6.10 Tag Management APIs.....	185
6.10.1 Enabling or Disabling the Default Tag of a Cluster.....	185
6.10.2 Querying the Status of Default Cluster Tags.....	186
6.10.3 Querying Tag Quotas.....	187
<b>7 API V1.1.....</b>	<b>190</b>
7.1 Cluster Management APIs.....	190
7.1.1 Creating a Cluster and Executing a Job.....	190
7.1.2 Resizing a Cluster.....	228
7.1.3 Querying a Cluster List.....	235
7.1.4 Querying Cluster Details.....	251
7.1.5 Querying a Host List.....	265
7.1.6 Terminating a Cluster.....	270
7.2 Auto Scaling APIs.....	271
7.2.1 Configuring an Auto Scaling Rule.....	271
7.3 Tag Management APIs.....	280
7.3.1 Adding Tags to a Specified Cluster.....	280
7.3.2 Querying Tags of a Specified Cluster.....	281

7.3.3 Deleting Tags from a Specified Cluster.....	283
7.3.4 Adding Tags to a Cluster in Batches.....	285
7.3.5 Deleting Tags from a Cluster in Batches.....	287
7.3.6 Querying All Tags.....	289
7.3.7 Querying a List of Clusters with Specified Tags.....	291
7.4 Availability Zones.....	297
7.4.1 Querying AZ Information.....	297
7.5 Version Metadata.....	300
7.5.1 Querying the Metadata of a Cluster Version.....	300
<b>8 Out-of-Date APIs.....</b>	<b>311</b>
8.1 Job API Management (Deprecated).....	311
8.1.1 Adding and Executing a Job (Deprecated).....	311
8.1.2 Querying the exe Object List of Jobs (Deprecated).....	322
8.1.3 Querying exe Object Details (Deprecated).....	327
8.1.4 Deleting a Job Execution Object (Deprecated).....	330
<b>9 Permissions Policies and Supported Actions.....</b>	<b>332</b>
9.1 Introduction.....	332
<b>10 Appendix.....</b>	<b>336</b>
10.1 ECS Specifications Used by MRS.....	336
10.2 BMS Specifications Used by MRS.....	338
10.3 Status Codes.....	339
10.4 Error Codes.....	343
10.5 Obtaining a Project ID.....	382
10.6 Obtaining Account ID.....	383
10.7 Obtaining the MRS Cluster Information.....	384
10.8 Roles and components supported by MRS.....	386

# 1 Before You Start

---

Welcome to *MapReduce Service (MRS) API Reference*. MRS provides enterprise-level big data clusters on the cloud. You can control your clusters and easily run big data components such as Hadoop, Spark, HBase, Kafka, and Storm in the clusters.

This document describes how to use application programming interfaces (APIs) to perform operations on MRS, such as creating or deleting clusters, adjusting nodes, as well as creating and executing jobs. For details about all supported operations, see [API Overview](#).

If you plan to access MRS through an API, ensure that you are familiar with MRS concepts. For details, see [Service Overview](#).

MRS 3.x does not support V1.1 APIs. You need to use V2 APIs. If you still need to use V1.1 APIs, obtain them from [Out-of-Date APIs](#).

## Endpoints

An endpoint is the **request address** for calling an API. Endpoints vary depending on services and regions. For the endpoints of services, see [Regions and Endpoints](#).

## Concepts

- Account

An account is created upon successful registration with HUAWEI CLOUD. The account has full access permissions for all of its cloud services and resources. It can be used to reset user passwords and grant user permissions. The account is a payment entity. For security purposes, do not directly use the account to perform routine management but create IAM users and grant them permissions for routine management.

- User

An IAM user is created using an account to use cloud services. Each IAM user has its own identity credentials (password and access keys).

An IAM user can view the account ID and user ID on the [My Credentials](#) page of the console. The account name, username, and password will be required for API authentication.

- **Region**


Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified as common regions and dedicated regions. A common region provides common cloud services that should be made available to all tenants. A dedicated region provides services of a specific type or only for specific tenants.

For details, see [Region and AZ](#).
- **AZ**

An availability zone (AZ) comprises one or more physical data centers equipped with independent ventilation, fire, water, and electricity facilities. Compute, network, storage, and other resources in an AZ are logically divided into multiple clusters. AZs within a region are connected using high-speed optical fibers to support cross-AZ high-availability systems.
- **Project**

Projects group and isolate resources (including compute, storage, and network resources) across physical regions. A default project is provided for each HUAWEI CLOUD region, and subprojects can be created under each default project. Users can be granted permissions to access all resources in a specific project. If you need more refined access control, you can create subprojects under a default project and purchase resources in subprojects. Then you can assign required permissions for users to access only resources in specific subprojects.

**Figure 1-1** Project isolating model



To view a project ID, go to the [My Credentials](#) page.

- Enterprise Project

Enterprise projects group and logically isolate resources. An enterprise project can contain resources from different regions, and resources can be transferred between enterprise projects.

For details about how to obtain enterprise project IDs and features, see the [Enterprise Management Service User Guide](#).



# 2 API Overview

MRS provides APIs that meet RESTful API design standards, as listed in [Table 2-1](#) and [Table 2-2](#).

Some APIs are gradually discarded during MRS version evolution. For details, see [Table 2-3](#).

**Table 2-1** V2 APIs

API	Function	API URI
Cluster management APIs	<a href="#">Creating a Cluster</a>	POST /v2/{project_id}/clusters
	<a href="#">Changing a Cluster Name</a>	PUT /v2/{project_id}/clusters/{cluster_id}/cluster-name
	<a href="#">Creating a Cluster and Submitting a Job</a>	POST /v2/{project_id}/run-job-flow
	<a href="#">Scaling Out a Cluster</a>	POST /v2/{project_id}/clusters/{cluster_id}/expand
	<a href="#">Scaling In a Cluster</a>	POST /v2/{project_id}/clusters/{cluster_id}/shrink
	<a href="#">Adding Components to a Cluster</a>	POST /v2/{project_id}/clusters/{cluster_id}/components
	<a href="#">Querying the Cluster Node List</a>	GET /v2/{project_id}/clusters/{cluster_id}/nodes
Job object APIs	<a href="#">Adding and Executing a Job</a>	POST /v2/{project_id}/clusters/{cluster_id}/job-executions
	<a href="#">Querying Information About a Job</a>	GET /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}
	<a href="#">Querying a List of Jobs</a>	GET /v2/{project_id}/clusters/{cluster_id}/job-executions

API	Function	API URI
	<b>Terminating a Job</b>	POST /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/kill
	<b>Obtaining SQL Results</b>	GET /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/sql-result
	<b>Deleting Jobs in Batches</b>	POST /v2/{project_id}/clusters/{cluster_id}/job-executions/batch-delete
Auto scaling APIs	<b>Viewing Auto Scaling Policies</b>	GET /v2/{project_id}/autoscaling-policy/{cluster_id}
	<b>Updating an Auto Scaling Policy</b>	PUT /v2/{project_id}/autoscaling-policy/{cluster_id}
	<b>Deleting an AS policy</b>	DELETE /v2/{project_id}/autoscaling-policy/{cluster_id}
	<b>Creating an AS policy</b>	POST /v2/{project_id}/autoscaling-policy/{cluster_id}
Cluster HDFS file APIs (V2)	<b>Obtaining the List of Files from a Specified Directory</b>	GET /v2/{project_id}/clusters/{cluster_id}/files?path={directory}&offset={offset}&limit={limit}&sort_key={sort_key}&order={order}
SQL APIs	<b>Submitting a SQL Statement</b>	POST /v2/{project_id}/clusters/{cluster_id}/sql-execution
	<b>Querying SQL Results</b>	GET /v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}
	<b>Canceling a SQL Execution Task</b>	POST /v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}/cancel
Agency management APIs (V2)	<b>Querying the Mapping Between a User (Group) and an IAM Agency</b>	GET /v2/{project_id}/clusters/{cluster_id}/agency-mapping
	<b>Updating the Mapping Between a User (Group) and an IAM Agency</b>	PUT /v2/{project_id}/clusters/{cluster_id}/agency-mapping
Data connection management APIs	<b>Creating a Data Connection</b>	POST /v2/{project_id}/data-connectors
	<b>Querying the Data Connection List</b>	GET /v2/{project_id}/data-connectors

API	Function	API URI
	<b>Updating a Data Connection</b>	PUT /v2/{project_id}/data-connectors/{connector_id}
	<b>Deleting a Data Connection</b>	DELETE /v2/{project_id}/data-connectors/{connector_id}
Version metadata query APIs	<b>Obtaining MRS Version List</b>	GET /v2/{project_id}/metadata/versions
	<b>Querying Available Specifications of an MRS Cluster Version</b>	GET /v2/{project_id}/metadata/version/{version_name}/available-flavor
IAM synchronization APIs	<b>Obtaining Synchronized IAM Users and User Groups</b>	GET /v2/{project_id}/clusters/{cluster_id}/iam-sync-user
	<b>Synchronizing an IAM User and User Group</b>	POST /v2/{project_id}/clusters/{cluster_id}/iam-sync-user
	<b>Cancelling Synchronization of Specified Users and User Groups</b>	DELETE /v2/{project_id}/clusters/{cluster_id}/iam-sync-user
Tag Management APIs	<b>Enabling or Disabling the Default Tag of a Cluster</b>	POST /v2/{project_id}/clusters/{cluster_id}/tags/switch
	<b>Querying the Status of Default Cluster Tags</b>	GET /v2/{project_id}/clusters/{cluster_id}/tags/status
	<b>Querying Tag Quotas</b>	GET /v2/{project_id}/clusters/{cluster_id}/tags/quota

**Table 2-2** V1.1 APIs

API	Function	API URI
Cluster management APIs	<b>Creating a Cluster and Executing a Job</b>	POST /v1.1/{project_id}/run-job-flow
	<b>Resizing a Cluster</b>	PUT /v1.1/{project_id}/cluster_infos/{cluster_id}
	<b>Querying a Cluster List</b>	GET /v1.1/{project_id}/cluster_infos

API	Function	API URI
	<b>Querying Cluster Details</b>	GET /v1.1/{project_id}/cluster_infos/{cluster_id}
	<b>Querying a Host List</b>	GET /v1.1/{project_id}/clusters/{cluster_id}/hosts
	<b>Terminating a Cluster</b>	DELETE /v1.1/{project_id}/clusters/{cluster_id}
Job object APIs	<b>Adding and Executing a Job (Deprecated)</b>	POST /v1.1/{project_id}/jobs/submit-job
	<b>Querying the exe Object List of Jobs (Deprecated)</b>	GET /v1.1/{project_id}/job-exes
	<b>Querying exe Object Details (Deprecated)</b>	GET /v1.1/{project_id}/job-exes/{job_exe_id}
Job execution object APIs	<b>Deleting a Job Execution Object (Deprecated)</b>	DELETE /v1.1/{project_id}/job-executions/{job_execution_id}
Auto scaling APIs	<b>Configuring an Auto Scaling Rule</b>	POST /v1.1/{project_id}/autoscaling-policy/{cluster_id}
Tag management APIs	<b>Adding Tags to a Specified Cluster</b>	POST /v1.1/{project_id}/clusters/{cluster_id}/tags
	<b>Deleting Tags from a Specified Cluster</b>	DELETE /v1.1/{project_id}/clusters/{cluster_id}/tags/{key}
	<b>Querying Tags of a Specified Cluster</b>	GET /v1.1/{project_id}/clusters/{cluster_id}/tags
	<b>Adding Tags to a Cluster in Batches</b>	POST /v1.1/{project_id}/clusters/{cluster_id}/tags/action
	<b>Querying All Tags</b>	GET /v1.1/{project_id}/clusters/tags
	<b>Querying a List of Clusters with Specified Tags</b>	POST /v1.1/{project_id}/clusters/resource_instances/action
AZ APIs	<b>Querying AZ Information</b>	GET /v1.1/{region_id}/available-zones
Version metadata	<b>Querying the Metadata of a Cluster Version</b>	GET /v1.1/{project_id}/metadata/versions/{version_name}

**Table 2-3** Historical APIs

API	Function	API URI
V1.1 Job Management API	<b>Adding and Executing a Job</b>	POST /v1.1/{project_id}/jobs/submit-job
	<b>Querying the exe Object List</b>	GET /v1.1/{project_id}/job-exes
	<b>Querying exe Object Details</b>	GET /v1.1/{project_id}/job-exes/{job_exe_id}
	<b>Deleting a Job Execution Object</b>	DELETE /v1.1/{project_id}/job-executions/{job_execution_id}

# 3 Selecting an API Type

---

Currently, MRS provides two types (V1.1 and V2) of APIs for cloud services with customized specifications. Only part of V2 APIs is available and mainly used for submitting jobs and SQL statements. If the API functions are the same, you are advised to use the V2 APIs first.

- All MRS versions support the V1.1 APIs.
- For MRS 1.X clusters, MRS 1.8.7 or later supports the V2 APIs.
- For MRS 2.X clusters, MRS 2.0.3 or later supports the V2 APIs.

Based on the V1.1 APIs, the V2 APIs have the following enhancements:

- Jobs can be submitted in a security cluster.
- The HiveSQL, Spark Python, and Flink jobs are supported.
- The SparkSQL and SparkScript results can be queried.

For details about the APIs and their functions, see [API Overview](#).

# 4 Calling APIs

## 4.1 Making an API Request

This section describes the structure of a REST API request, and describes how to call an API using [Creating an IAM User](#) as an example. This API obtains a user token, which can be used for authentication when other APIs are called.

### Request URI

A request URI is in the following format:


**{URI-scheme}://{Endpoint}/{resource-path}?{query-string}**

Although a request URI is included in the request header, most programming languages or frameworks require the request URI to be passed separately.

- **URI-scheme:** Protocol used to transmit requests. All APIs use **HTTPS**.
- **Endpoint:** Domain name or IP address of the server bearing the REST service. It can be obtained from [Endpoints](#). For example, The endpoint of IAM in the CN-Hong Kong region is **iam.ap-southeast-1.myhuaweicloud.com**.
- **resource-path:** API access path. Obtain the value from the URI of an API. For example, the **resource-path** of the API used by the administrator to create an IAM user is **/v3.0/OS-USER/users**.
- **query-string:** Query parameter, which is optional for APIs. Ensure that a question mark (?) is included before each query parameter that is in the format of "*Parameter name=Parameter value*". For example, **?limit=10** indicates that a maximum of 10 data records are allowed.

For example, if you want to create an IAM user, use the endpoint of any region because IAM is a global service. The endpoint of the CN-Hong Kong region is **iam.cn-north-4.myhuaweicloud.com**. **iam.ap-southeast-1.myhuaweicloud.com**. Find **resource-path** (**/v3.0/OS-USER/users**) in the URI of the [Creating an IAM User](#) API.

```
https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
```

**Figure 4-1** Example URI**NOTE**

To simplify the URI display in this document, each API is provided only with a **resource-path** and a request method. The **URI-scheme** of all APIs is **HTTPS**, and the endpoints of all APIs in the same region are identical.

## Request Methods

The HTTP protocol defines the following request methods that can be used to send a request to the server:

- **GET**: requests a server to return specified resources.
- **PUT**: requests a server to update specified resources.
- **POST**: requests a server to add a resource or perform special operations.
- **DELETE**: requests a server to delete specified resources, for example, to delete an object.
- **HEAD**: requests a server resource header.
- **PATCH**: requests a server to update the partial content of a specified resource. If the resource does not exist, a new resource will be created.

Check the URI of [Creating an IAM User](#). The request method is POST, and the request is as follows:

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
```

## Request Header

You can also add additional header fields to a request, such as the fields required by a specified URI or HTTP method. For example, to request for the authentication information, add **Content-Type**, which specifies the request body type.

Common request header fields are as follows:

- **Content-Type**: specifies the request body type or format. This field is mandatory and its default value is **application/json**. Other values of this field are provided for specific APIs, if any.
- **Authorization**: specifies signature authentication information. This field is optional. When AK/SK authentication is enabled, this field is automatically specified when SDK is used to sign the request. For more information, see "AK/SK-based Authentication" in [Authentication](#).
- **X-Sdk-Date**: specifies the time when a request is sent. This field is optional. When AK/SK authentication is enabled, this field is automatically specified when SDK is used to sign the request. For more information, see "AK/SK-based Authentication" in [Authentication](#).
- **X-Auth-Token**: specifies a user token, which is optional. This field is mandatory when token authentication is used. It is the response to the API



used to **obtain a user token**. This API is the only one that does not require authentication.

- **X-Project-ID**: specifies subproject ID. This field is optional and can be used in multi-project scenarios. The **X-Project-ID** field is mandatory in the request header for accessing resources in a sub-project through AK/SK-based authentication.
- **X-Domain-ID**: specifies account ID, which is optional. When you call APIs of global services using AK/SK-based authentication, **X-Domain-ID** needs to be configured in the request header.

For the **Creating an IAM User** API, if AK/SK authentication is used, the request with the header is as follows:

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
Content-Type: application/json
X-Sdk-Date: 20240416T095341Z
Authorization: SDK-HMAC-SHA256 Access=*****, SignedHeaders=content-type;host;x-sdk-date,
Signature=*****
```

## Request Body

A request body is generally sent in a structured format. It corresponds to **Content-Type** in the request header and transfers data except for the request header. If the request body needs to support Chinese characters, set **Content-type** to **UTF-8**. For example, **Content-Type: application/json; charset=utf-8**.

Request bodies vary with APIs. Some APIs do not require a request body, such as the APIs requested using GET and DELETE methods.

You can obtain the request parameters and parameter description from the API request. For details, see **Creating an IAM User**. The following provides an example request with a body included. Replace the bold fields with the actual values.

- **accountid** indicates the ID of the account to which the IAM user belongs.
- **username** indicates the IAM username to be created.
- **email** indicates the email address of the IAM user.
- **\*\*\*\*\*** indicates the login password of the IAM user.

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
Content-Type: application/json
X-Sdk-Date: 20240416T095341Z
Authorization: SDK-HMAC-SHA256 Access=*****, SignedHeaders=content-type;host;x-sdk-date,
Signature=*****
```

```
{
  "user": {
    "domain_id": "accountid",
    "name": "username",
    "password": "*****",
    "email": "email",
    "description": "IAM User Description"
  }
}
```

If all data required for the API request is available, you can send the request to call the API through curl, Postman, or coding.

## 4.2 Authentication

Requests for calling an API can be authenticated using either of the following methods:

- AK/SK-based authentication: Requests are authenticated by encrypting the request body using an AK/SK pair.
- Token-based authentication: Requests are authenticated using a token.

### AK/SK-based Authentication

#### NOTE

- AK/SK-based authentication supports API requests with a body not larger than 12 MB. For API requests with a larger body, token-based authentication is recommended.
- Both AK/SK in a permanent access key or in a temporary access are supported. The **X-Security-Token** field must be configured when the AK/SK in the temporary access key is used, and the field value is the **security\_token** of the temporary access key.

In AK/SK-based authentication, AK/SK is used to sign requests and the signature is then added to the requests for authentication.

- AK: access key ID, which is a unique identifier used in conjunction with a secret access key to sign requests cryptographically.
- SK: secret access key used in conjunction with an AK to sign requests cryptographically. It identifies a request sender and prevents the request from being modified.

In AK/SK-based authentication, you can use an AK/SK to sign requests based on the signature algorithm or use the signing SDK to sign requests. For details about how to sign requests and use the signing SDK, see [AK/SK Signature Guide](#).

---

#### NOTICE

The signing SDK is only used for signing requests and is different from the SDKs provided by services.

---

### Token-based Authentication

#### NOTE

- The validity period of a token is 24 hours. When using a token for authentication, cache it to prevent frequently API calling.
- Ensure that the token is valid when you use it. Using a token that will soon expire may cause API calling failures.

A token specifies temporary permissions in a computer system. During API authentication using a token, the token is added to requests to get permissions for calling the API.

The token can be obtained by calling the API in [Obtaining a User Token](#). A project-level token is required for calling this service API, that is, when calling the

API for [obtaining a user token](#), set the value of **auth.scope** in the request body to **project**.

```
{
  "auth": {
    "identity": {
      "methods": [
        "password"
      ],
      "password": {
        "user": {
          "name": "username",
          "password": "*****#",
          "domain": {
            "name": "domainname"
          }
        }
      }
    },
    "scope": {
      "project": {
        "name": "xxxxxxxx"
      }
    }
  }
}
```

After a token is obtained, the **X-Auth-Token** header field must be added to requests to specify the token when calling other APIs. For example, if the token is **ABCDEFJ....**, add **X-Auth-Token: ABCDEFJ....** in a request as follows:

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
Content-Type: application/json
X-Auth-Token: ABCDEFG....
```

## 4.3 Response

### Status Code

After sending a request, you will receive a response, including a status code, response header, and response body.

A status code is a group of digits, ranging from 1xx to 5xx. It indicates the status of a request. For more information, see [Status Codes](#).

For example, if status code 201 is returned after calling the [Creating an IAM User](#) API, the request is successful.

### Response Header

Similar to a request, a response also has a header, for example, **Content-Type**.

Check the response header of the [Creating an IAM User](#) API. The **x-subject-token** field is the user token. This token can then be used to authenticate the calling of other APIs.

**Figure 4-2** Response header

```
"X-Frame-Options": "SAMEORIGIN",
"X-IAM-ETag-id": "2562365939-d8f6f12921974cb097338ac11fceac8a",
"Transfer-Encoding": "chunked",
"Strict-Transport-Security": "max-age=31536000; includeSubdomains;",
"Server": "api-gateway",
"X-Request-Id": "af2953f2bcc67a42325a69a19e6c32a2",
"X-Content-Type-Options": "nosniff",
"Connection": "keep-alive",
"X-Download-Options": "noopen",
"X-XSS-Protection": "1; mode=block;",
"X-IAM-Trace-Id": "token_██████████_null_af2953f2bcc67a42325a69a19e6c32a2",
"Date": "Tue, 21 May 2024 09:03:40 GMT",
"Content-Type": "application/json; charset=utf8"
```

## Response Body

This part is optional. The body of a response is often returned in structured format (for example, JSON or XML) as specified in the **Content-Type** header field. The response body transfers content except the response header.

The following example shows a part of the response body of the [Creating an IAM User](#) API.

```
{
  "user": {
    "id": "c131886aec...",
    "name": "IAMUser",
    "description": "IAM User Description",
    "areacode": "",
    "phone": "",
    "email": "****@***.com",
    "status": null,
    "enabled": true,
    "pwd_status": false,
    "access_mode": "default",
    "is_domain_owner": false,
    "xuser_id": "",
    "xuser_type": "",
    "password_expires_at": null,
    "create_time": "2024-05-21T09:03:41.000000",
    "domain_id": "d78cbac1.....",
    "xdomain_id": "30086000.....",
    "xdomain_type": "",
    "default_project_id": null
  }
}
```

If an error occurs during API calling, an error code and a message will be displayed. The following shows an error response body.

```
{
  "error_msg": "Request body is invalid.",
  "error_code": "IAM.0011"
}
```

In the response body, **error\_code** is an error code, and **error\_msg** provides information about the error.

# 5 Application Cases

---

## 5.1 Creating an MRS Cluster

### Scenario

This section describes how to create an MRS analysis cluster using APIs. For details on how to call APIs, see [Making an API Request](#).

### Constraints

- A VPC and subnet have been created using the VPC service in the region where the cluster is to be created. For details about how to create a VPC, see [Querying VPCs](#) and [Creating a VPC](#). For details about how to create a subnet, see [Querying Subnets](#) and [Creating a Subnet](#).
- You have obtained the region and AZ information of the cluster to be created. For details, see [Endpoints](#).
- You have obtained the project ID of the region where the cluster is to be created. For details, see [Obtaining a Project ID](#).
- You have determined the version of the cluster to be created and the components supported by the version. For details, see [List of MRS Component Versions](#).
- In this example, an analysis cluster is created based on the pay-per-use billing mode.

### Procedure

- API  
URI format: POST /v2/{project\_id}/clusters  
For details, see [Creating a Cluster](#).
- Example request  
POST: https://{endpoint}/v2/{project\_id}/clusters
  - For details about **{endpoint}**, see [Endpoints](#).
  - For details about **{project\_id}**, see [Obtaining a Project ID](#).

- Obtain the value of **node\_size** from the cluster creation page on the MRS console.

## Request body:

```
{
  "cluster_version": "MRS 3.2.0-LTS.1",
  "cluster_name": "mrs_Demo",
  "cluster_type": "ANALYSIS",
  "charge_info": {
    "charge_mode": "postPaid"
  },
  "region": "",
  "availability_zone": "",
  "vpc_name": "vpc-37cd",
  "subnet_name": "subnet-ed99",
  "components": "Hadoop,Spark2x,HBase,Hive,Hue,Flink,Oozie,Ranger,Tez",
  "safe_mode": "KERBEROS",
  "manager_admin_password": "xxx",
  "login_mode": "PASSWORD",
  "node_root_password": "xxx",
  "log_collection": 1,
  "mrs_ecs_default_agency": "MRS_ECS_DEFAULT_AGENCY",
  "tags": [
    {
      "key": "tag1",
      "value": "111"
    },
    {
      "key": "tag2",
      "value": "222"
    }
  ],
  "node_groups": [
    {
      "group_name": "master_node_default_group",
      "node_num": 2,
      "node_size": "rc3.4xlarge.4.linux.bigdata",
      "root_volume": {
        "type": "SAS",
        "size": 480
      },
      "data_volume": {
        "type": "SAS",
        "size": 600
      },
      "data_volume_count": 1
    },
    {
      "group_name": "core_node_analysis_group",
      "node_num": 3,
      "node_size": "rc3.4xlarge.4.linux.bigdata",
      "root_volume": {
        "type": "SAS",
        "size": 480
      },
      "data_volume": {
        "type": "SAS",
        "size": 600
      },
      "data_volume_count": 1
    },
    {
      "group_name": "task_node_analysis_group",
      "node_num": 3,
      "node_size": "rc3.4xlarge.4.linux.bigdata",
      "root_volume": {
        "type": "SAS",
        "size": 480
      }
    }
  ]
}
```

```
"data_volume": {
  "type": "SAS",
  "size": 600
},
"data_volume_count": 1,
"auto_scaling_policy": {
  "auto_scaling_enable": true,
  "min_capacity": 0,
  "max_capacity": 1,
  "resources_plans": [],
  "exec_scripts": [],
  "rules": [
    {
      "name": "default-expand-1",
      "description": "",
      "adjustment_type": "scale_out",
      "cool_down_minutes": 5,
      "scaling_adjustment": "1",
      "trigger": {
        "metric_id": 2003,
        "metric_name": "StormSlotAvailablePercentage",
        "metric_value": 100,
        "comparison_operator_id": 2003,
        "comparison_operator": "LTOE",
        "evaluation_periods": "1"
      }
    }
  ]
}
]
```

For details about the parameters, see [Creating a Cluster](#).

- Example response

```
{
  "cluster_id": "da1592c2-bb7e-468d-9ac9-83246e95447a"
}
```

## 5.2 Scaling Out a Cluster

### Scenario

After a cluster is created, add Core or Task nodes to the cluster. After an MRS cluster is created, the number of Master nodes cannot be adjusted. That is, Master nodes cannot be scaled in or out. For details on how to call APIs, see [Making an API Request](#).

### Constraints

- A cluster has been created and is in the **Running** state.
- You have obtained the project ID of the region where the cluster is to be created. For details, see [Obtaining a Project ID](#).
- You have obtained the cluster ID, that is, the value of **cluster\_id** in the command output returned after the cluster is successfully created. For details about how to obtain the cluster ID, see [Obtaining a Cluster ID](#).
- This section uses the Core node as an example.

## Procedure

- API  
URI format: PUT /v1.1/{project\_id}/cluster\_infos/{cluster\_id}  
For details, see [Resizing a Cluster](#).
- Example request  
PUT: https://{endpoint}/v1.1/{project\_id}/cluster\_infos/{cluster\_id}
  - For details about **{endpoint}**, see [Endpoints](#).
  - For details about **{project\_id}**, see [Obtaining a Project ID](#).
  - **{cluster\_id}** indicates the value of **cluster\_id** in the command output returned after the cluster is successfully created. You can also obtain the value of **cluster\_id** by referring to [Obtaining a Cluster ID](#).

Request body:

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_out",
    "node_id": "node_orderadd",
    "node_group": "core_node_default_group",
    "instances": "1",
    "skip_bootstrap_scripts": false,
    "scale_without_start": false
  },
  "previous_values": {
    "plan_id": ""
  }
}
```

For details about the parameters, see [Resizing a Cluster](#).

- Example response

```
{
  "result": "succeeded"
}
```

## 5.3 Scaling in a Cluster

### Scenario

This section describes how to scale in a Core or Task node in the cluster after it is created. After an MRS cluster is created, the number of Master nodes cannot be adjusted. That is, Master nodes cannot be scaled in or out. For details on how to call APIs, see [Making an API Request](#).

### Constraints

- A cluster has been created and is in the **Running** state.
- You have obtained the project ID of the region where the cluster is to be created. For details, see [Obtaining a Project ID](#).
- You have obtained the cluster ID, that is, the value of **cluster\_id** in the command output returned after the cluster is successfully created. For details about how to obtain the cluster ID, see [Obtaining a Cluster ID](#).
- The Core node is used as an example.



## Procedure

- API  
URI format: PUT /v1.1/{project\_id}/cluster\_infos/{cluster\_id}  
For details, see [Resizing a Cluster](#).
- Example request  
PUT: https://{endpoint}/v1.1/{project\_id}/cluster\_infos/{cluster\_id}
  - For details about **{endpoint}**, see [Endpoints](#).
  - For details about **{project\_id}**, see [Obtaining a Project ID](#).
  - **{cluster\_id}** indicates the value of **cluster\_id** in the command output returned after the cluster is successfully created. You can also obtain the value of **cluster\_id** by referring to [Obtaining a Cluster ID](#).

Request body:

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_in",
    "node_id": "node_orderadd",
    "node_group": "core_node_default_group",
    "instances": "1"
  },
  "previous_values": {
    "plan_id": ""
  }
}
```

For details about the parameters, see [Resizing a Cluster](#).

- Example response  

```
{
  "result": "succeeded"
}
```

## 5.4 Creating a Job

### Scenario

This API is used to add and submit a job in an MRS cluster. For details on how to call APIs, see [Making an API Request](#).

### Constraints

- A cluster has been created and is in the **Running** state.
- You have obtained the project ID of the region where the cluster is to be created. For details, see [Obtaining a Project ID](#).
- You have obtained the cluster ID, that is, the value of **cluster\_id** in the command output returned after the cluster is successfully created. For details about how to obtain the cluster ID, see [Obtaining a Cluster ID](#).
- IAM users have been synchronized. On the Overview tab page of the cluster details page, click **Click to synchronize** on the right of **IAM User Sync** to synchronize IAM users.

- The job-related programs and input files have been stored in OBS.
- In this example, a MapReduce job is added.

## Procedure

- API  
URI format: POST /v2/{project\_id}/clusters/{cluster\_id}/job-executions  
For details, see [Adding and Executing a Job](#).
- Example request  
POST: `https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions`
  - For details about **{endpoint}**, see [Endpoints](#).
  - For details about **{project\_id}**, see [Obtaining a Project ID](#).
  - **{cluster\_id}** indicates the value of **cluster\_id** in the command output returned after the cluster is successfully created. You can also obtain the value of **cluster\_id** by referring to [Obtaining a Cluster ID](#).

Request body:

```
{
  "job_name": "MapReduceTest",
  "job_type": "MapReduce",
  "arguments": [
    "obs://obs-test/program/hadoop-mapreduce-examples-x.x.x.jar",
    "wordcount",
    "obs://obs-test/input/",
    "obs://obs-test/job/mapreduce/output"
  ],
  "properties": {
    "fs.obs.endpoint": "obs endpoint",
    "fs.obs.access.key": "xxx",
    "fs.obs.secret.key": "yyy"
  }
}
```

For details about the parameters, see [Adding and Executing a Job](#).

- Example response  

```
{
  "job_submit_result": {
    "job_id": "44b37a20-ffe8-42b1-b42b-78a5978d7e40",
    "state": "COMPLETE"
  }
}
```

## 5.5 Terminating a Job

### Scenario

This API is used to manually terminate the job if a job is not completed after being submitted. For details on how to call APIs, see [Making an API Request](#).

### Constraints

- A cluster has been created and is in the **Running** state.
- You have obtained the project ID of the region where the cluster is to be created. For details, see [Obtaining a Project ID](#).

- You have obtained the cluster ID, that is, the value of **cluster\_id** in the command output returned after the cluster is successfully created. For details about how to obtain the cluster ID, see [Obtaining a Cluster ID](#).
- You have obtained the job ID, that is, the value of **job\_id** in the returned result after the job is successfully submitted. For details about how to obtain the job ID, see [Obtaining a Job ID](#).
- IAM users have been synchronized. On the Overview tab page of the cluster details page, click **Click to synchronize** on the right of **IAM User Sync** to synchronize IAM users.
- The job-related programs and input files have been stored in OBS.
- In this example, a MapReduce job is added.

## Procedure

- API  
URI format: POST /v2/{project\_id}/clusters/{cluster\_id}/job-executions/{job\_execution\_id}/kill  
For details, see [Terminating a Job](#).
- Example request  
POST: `https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/kill`
  - For details about **{endpoint}**, see [Endpoints](#).
  - For details about **{project\_id}**, see [Obtaining a Project ID](#).
  - **{cluster\_id}** indicates the value of **cluster\_id** in the command output returned after the cluster is successfully created. You can also obtain the value of **cluster\_id** by referring to [Obtaining a Cluster ID](#).
  - **{job\_execution\_id}** indicates the job ID returned after the job is successfully submitted. You can also obtain the job ID by referring to [Obtaining a Job ID](#).Request body: None
- Example response  
None

## 5.6 Terminating a Cluster

### Scenario

This API is used to delete a cluster after data processing and analysis are completed or the cluster is abnormal.

Clusters in any of the following states cannot be terminated:

- **Scaling-out**: The cluster is being scaled out.
- **Scaling-in**: The cluster is being scaled in.
- **Starting**: The cluster is being started.
- **Terminating**: The cluster is being deleted.
- **Terminated**: The cluster has been terminated.

- **Failed:** The cluster is failed.

For details on how to call APIs, see [Making an API Request](#).

## Constraints

- You have obtained the project ID of the region where the cluster is to be created. For details, see [Obtaining a Project ID](#).
- You have obtained the cluster ID, that is, the value of **cluster\_id** in the command output returned after the cluster is successfully created. For details about how to obtain the cluster ID, see [Obtaining a Cluster ID](#).

## Procedure

- API  
URI format: DELETE /V1.1/{project\_id}/clusters/{cluster\_id}  
For details, see [Terminating a Cluster](#).
- Example request  
DELETE: `https://{endpoint}/v1.1/{project_id}/clusters/{cluster_id}`
  - For details about **{endpoint}**, see [Endpoints](#).
  - For details about **{project\_id}**, see [Obtaining a Project ID](#).
  - **{cluster\_id}** indicates the value of **cluster\_id** in the command output returned after the cluster is successfully created. You can also obtain the value of **cluster\_id** by referring to [Obtaining a Cluster ID](#).

Request body: None

- Example response

```
{  
  "result": "succeeded"  
}
```

# 6 API V2

## 6.1 Cluster Management APIs

### 6.1.1 Creating a Cluster

#### Function

This API is used to create an MRS cluster.

Before using the API, you need to obtain the resources listed in [Table 6-1](#).

**Table 6-1** Obtaining resources

Resource	How to Obtain
VPC	See operation instructions in <a href="#">Querying VPCs</a> and <a href="#">Creating a VPC</a> in the <i>VPC API Reference</i> .
Subnet	See operation instructions in <a href="#">Querying Subnets</a> and <a href="#">Creating a Subnet</a> in the <i>VPC API Reference</i> .
Key Pair	See operation instructions in <a href="#">Querying SSH Key Pairs</a> and <a href="#">Creating and Importing an SSH Key Pair</a> in the <i>ECS API Reference</i> .
Zone	See <a href="#">Endpoints</a> for details about regions and AZs.
Version	Currently, MRS 1.9.2, 3.1.0, 3.1.5, 3.1.2-LTS.3, and 3.2.0-LTS.1 are supported.

Resource	How to Obtain
Component	<ul style="list-style-type: none"> <li>● MRS 3.2.0-LTS.1 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and Guardian.</li> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, Flume, and Guardian.</li> <li>- A custom cluster contains the following components: CDL, Hadoop, Spark2x, HBase, Hive, Hue, IoTDB, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse, and Guardian.</li> </ul> </li> <li>● MRS 3.1.5 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Sqoop, and Guardian</li> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Sqoop, Guardian, Kafka, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Kafka, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, Kudu, Sqoop, and Guardian.</li> </ul> </li> <li>● MRS 3.1.2-LTS.3 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, and Tez.</li> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse.</li> </ul> </li> <li>● MRS 3.1.0 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, and Kudu.</li> </ul> </li> </ul>

Resource	How to Obtain
	<ul style="list-style-type: none"> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Kafka, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Kafka, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, and Kudu.</li> <li>• MRS 3.0.5 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, and Alluxio.</li> <li>- A streaming cluster contains the following components: Kafka, Storm, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Alluxio, Kafka, Storm, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Kafka, Storm, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, Kudu, and Alluxio.</li> </ul> </li> <li>• MRS 2.1.0 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Presto, Hadoop, Spark, HBase, Hive, Hue, Loader, Tez, Impala, Kudu, and Flink.</li> <li>- A streaming cluster contains the following components: Kafka, Storm, and Flume.</li> </ul> </li> <li>• MRS 1.9.2 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Presto, Hadoop, Spark, HBase, OpenTSDB, Hive, Hue, Loader, Tez, Flink, Alluxio, and Ranger.</li> <li>- A streaming cluster contains the following components: Kafka, KafkaManager, Storm, and Flume.</li> </ul> </li> </ul>

## Constraints

None

## Debugging

You can debug this API through automatic authentication in API Explorer. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

POST /v2/{project\_id}/clusters

**Table 6-2** URI parameter

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .

## Request Parameters

**Table 6-3** Request body parameters

Parameter	Mandatory	Type	Description
cluster_version	Yes	String	Cluster version. Possible values: <ul style="list-style-type: none"> <li>• MRS 1.9.2</li> <li>• MRS 3.1.0</li> <li>• MRS 3.1.2-LTS.3</li> <li>• MRS 3.1.5</li> <li>• MRS 3.2.0-LTS.1</li> </ul>
cluster_name	Yes	String	Cluster name. It must be unique. A cluster name can contain only 1 to 64 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed.
cluster_type	Yes	String	Cluster type. Possible values: <ul style="list-style-type: none"> <li>• <b>ANALYSIS</b>: analysis cluster</li> <li>• <b>STREAMING</b>: streaming cluster</li> <li>• <b>MIXED</b>: hybrid cluster</li> <li>• <b>CUSTOM</b>: custom cluster, which is supported only by MRS 3.x.</li> </ul>
charge_info	No	object	The billing type. For details, see <a href="#">Table 6-5</a> .



Parameter	Mandatory	Type	Description
region	Yes	String	Information about the region where the cluster is located. For details, see <a href="#">Endpoints</a> .
is_dec_project	No	Boolean	Whether the cluster is specific for the DeC. The default value is <b>false</b> .
vpc_name	Yes	String	Name of the VPC where the subnet locates. Perform the following operations to obtain the VPC name from the VPC management console: <ol style="list-style-type: none"> <li>1. Log in to the management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>. On the <b>Virtual Private Cloud</b> page, obtain the VPC name from the list.</li> </ol>
subnet_id	No	String	Subnet ID, which can be obtained by performing the following operations on the VPC management console: <ol style="list-style-type: none"> <li>1. Log in to the VPC management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> <li>3. Locate the row that contains the target VPC and click the number in the <b>Subnets</b> column to view the subnet information.</li> <li>4. Click the subnet name to obtain the network ID. At least one of <b>subnet_id</b> and <b>subnet_name</b> must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. <b>subnet_id</b> is recommended.</li> </ol>

Parameter	Mandatory	Type	Description
subnet_name	Yes	String	Subnet name. Perform the following operations to obtain the subnet name from the VPC management console: <ol style="list-style-type: none"> <li>1. Log in to the management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> <li>3. Locate the row that contains the target VPC and click the number in the <b>Subnets</b> column to obtain the subnet name. At least one of <b>subnet_id</b> and <b>subnet_name</b> must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. If only <b>subnet_name</b> is configured and subnets with the same name exist in the VPC, the first subnet name in the VPC is used when a cluster is created. <b>subnet_id</b> is recommended.</li> </ol>
components	Yes	String	List of component names, which are separated by commas (.). For details about the component names, see the component list of each version in <a href="#">Table 6-1</a> .
external_data_sources	No	Array of ClusterDataConnect or Map objects	When deploying components such as Hive and Ranger, you can associate data connections and store metadata in associated databases. For details about the parameters, see <a href="#">Table 6-4</a> .
availability_zone	Yes	String	AZ name. Multi-AZ clusters are not supported. See <a href="#">Endpoints</a> for details about AZs.

Parameter	Mandatory	Type	Description
security_groups_id	No	String	<p>Security group ID of the cluster.</p> <ul style="list-style-type: none"> <li>If this parameter is left blank, MRS automatically creates a security group, whose name starts with <b>mrs_{cluster_name}</b>.</li> <li>If this parameter is not left blank, a fixed security group is used to create a cluster. The transferred ID must be the security group ID owned by the current tenant. The security group must include an inbound rule in which all protocols and all ports are allowed and the source is the IP address of the specified node on the management plane.</li> </ul>
auto_create_default_security_group	No	Boolean	<p>Whether to create the default security group for the MRS cluster. The default value is <b>false</b>. If this parameter is set to <b>true</b>, the default security group will be created for the cluster regardless of whether <b>security_groups_id</b> is specified.</p>
safe_mode	Yes	String	<p>Running mode of an MRS cluster.</p> <ul style="list-style-type: none"> <li><b>SIMPLE</b>: normal cluster. In a normal cluster, Kerberos authentication is disabled, and users can use all functions provided by the cluster.</li> <li><b>KERBEROS</b>: security cluster. In a security cluster, Kerberos authentication is enabled, and common users cannot use the file management and job management functions of an MRS cluster or view cluster resource usage and the job records of Hadoop and Spark. To use more functions, the users must obtain the relevant permissions from the Manager administrator.</li> </ul>

Parameter	Mandatory	Type	Description
manager_admin_password	Yes	String	<p>Password of the MRS Manager administrator. The password must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• Must contain 8 to 26 characters.</li> <li>• Must contain at least four of the following: uppercase letters, lowercase letters, digits, and special characters (!@\$%^_+= [{}]:./?), but must not contain spaces.</li> <li>• Cannot be the username or the username spelled backwards.</li> </ul>
login_mode	Yes	String	<p>Node login mode.</p> <ul style="list-style-type: none"> <li>• <b>PASSWORD</b>: password-based login. If this value is selected, <b>node_root_password</b> cannot be left blank.</li> <li>• <b>KEYPAIR</b>: specifies the key pair used for login. If this value is selected, <b>node_keypair_name</b> cannot be left blank.</li> </ul>
node_root_password	No	String	<p>Password of user <b>root</b> for logging in to a cluster node. A password must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• Must be 8 to 26 characters long.</li> <li>• Must contain at least four of the following: uppercase letters, lowercase letters, digits, and special characters (!@\$%^_+= [{}]:./?), but must not contain spaces.</li> <li>• Cannot be the username or the username spelled backwards.</li> </ul>
node_keypair_name	No	String	<p>Name of a key pair You can use a key pair to log in to the Master node in the cluster.</p>

Parameter	Mandatory	Type	Description
enterprise_project_id	No	String	<p>Enterprise project ID.</p> <p>When you create a cluster, associate the enterprise project ID with the cluster.</p> <p>The default value is <b>0</b>, indicating the <b>default</b> enterprise project.</p> <p>To obtain the enterprise project ID, see the <b>id</b> value in the <b>enterprise_project field data structure</b> table in section <b>Querying the Enterprise Project List</b> of the <i>Enterprise Management API Reference</i>.</p>
eip_address	No	String	<p>An EIP bound to an MRS cluster can be used to access MRS Manager. The EIP must have been created and must be in the same region as the cluster.</p>
eip_id	No	String	<p>ID of the bound EIP. This parameter is mandatory when <b>eip_address</b> is configured. To obtain the EIP ID, log in to the VPC console, choose <b>Network &gt; Elastic IP and Bandwidth &gt; Elastic IP</b>, click the EIP to be bound, and obtain the ID in the Basic Information area.</p>
mrs_ecs_default_agency	No	String	<p>Name of the agency bound to a cluster node by default. The value is fixed to <b>MRS_ECS_DEFAULT_AGENCY</b>.</p> <p>An agency allows ECS or BMS to manage MRS resources. You can configure an agency of the ECS type to automatically obtain the AK/SK to access OBS.</p> <p>The <b>MRS_ECS_DEFAULT_AGENCY</b> 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.</p>

Parameter	Mandatory	Type	Description
template_id	No	String	<p>Template used for node deployment when the cluster type is CUSTOM.</p> <ul style="list-style-type: none"> <li>• <b>mgmt_control_combined_v2:</b> template for jointly deploying the management and control nodes. The management and control roles are co-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.</li> <li>• <b>mgmt_control_separated_v2:</b> The management and control roles 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.</li> <li>• <b>mgmt_control_data_separated_v2:</b> 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.</li> </ul>
tags	No	Array of tag objects	<p>Cluster tag For more parameter description, see <a href="#">Table 6-6</a>.</p> <p>A maximum of 10 tags can be added to a cluster.</p>

Parameter	Mandatory	Type	Description
log_collection	No	Integer	<p>Whether to collect logs when cluster creation fails: The default value is <b>1</b>, indicating that OBS buckets will be created and only used to collect logs that record MRS cluster creation failures. Enumerated values:</p> <ul style="list-style-type: none"> <li>• <b>0</b>: Do not collect logs.</li> <li>• <b>1</b>: Collect logs.</li> </ul> <p>The default value is <b>1</b>, indicating that OBS buckets will be created and only used to collect logs that record MRS cluster creation failures.</p>
node_groups	Yes	Array of NodeGroupV2 objects	Information about the node groups in the cluster. For details about the parameters, see <a href="#">Table 6-7</a> .
bootstrap_scripts	No	Array of Bootstrap Script objects	<p>Bootstrap action script information. For more parameter description, see <a href="#">Table 6-9</a>.</p> <p>MRS 3.x does not support this parameter.</p>
add_jobs	No	Array of add_jobs objects	<p>You can submit a job when creating a cluster. Currently, only versions earlier than MRS 1.8.7 support this function. Currently, only one job can be submitted. You are advised to use the <b>steps</b> parameter in the Creating a Cluster and Submitting a Job API. For details about this parameter, see <a href="#">Table 6-10</a>.</p>

Parameter	Mandatory	Type	Description
log_uri	No	String	The OBS path to which cluster logs are dumped. After the log dump function is enabled, the read and write permissions on the OBS path are required to upload logs. Configure the default agency <b>MRS_ECS_DEFAULT_AGENCY</b> or customize an agency with the read and write permissions on the OBS path. For details, see <a href="#">Configuring a Storage-Compute Decoupled Cluster (Agency)</a> . This parameter is available only for cluster versions that support dumping cluster logs to OBS.
component_configs	No	Array of ComponentConfig objects	The custom configuration of cluster components. This parameter applies only to cluster versions that support the feature of creating a cluster by customizing component configurations. For details about this parameter, see <a href="#">ComponentConfig</a> .
smn_notify	No	SmnNotify object	SMN alarm notifications. For details about this parameter, see <a href="#">Table 6-18</a> .

**Table 6-4** ClusterDataConnectorMap

Parameter	Mandatory	Type	Description
map_id	No	Integer	Data connection association ID
connector_id	No	String	Data connection ID
component_name	No	String	Component name



Parameter	Mandatory	Type	Description
role_type	No	String	Component role type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>hive_metastore</b>: Hive Metastore role</li> <li>• <b>hive_data</b>: Hive role</li> <li>• <b>hbase_data</b>: HBase role</li> <li>• <b>ranger_data</b>: Ranger role</li> </ul>
source_type	No	String	Data connection type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>LOCAL_DB</b>: local metadata</li> <li>• <b>RDS_POSTGRES</b>: RDS PostgreSQL database</li> <li>• <b>RDS_MYSQL</b>: RDS MySQL database</li> <li>• <b>gaussdb-mysql</b>: GaussDB(for MySQL)</li> </ul>
cluster_id	No	String	ID of the associated cluster
status	No	Integer	Data connection status. The options are as follows: <ul style="list-style-type: none"> <li>• <b>0</b>: normal</li> <li>• <b>1</b>: in use</li> </ul>

**Table 6-5** ChargeInfo

Parameter	Mandatory	Type	Description
charge_mode	Yes	String	Billing mode. The options are as follows: <ul style="list-style-type: none"> <li>• <b>prePaid</b>: indicates the yearly/monthly billing mode. (This mode is now supported for the API for creating a cluster, but is not supported for the API for creating a cluster and submitting a job.)</li> <li>• <b>postPaid</b>: indicates the pay-per-use billing mode.</li> </ul>

Parameter	Mandatory	Type	Description
period_type	No	String	<p>Period type. The options are as follows:</p> <ul style="list-style-type: none"> <li>● <b>month</b>: The cluster is billed by month.</li> <li>● <b>year</b>: The cluster is billed by year.</li> <li>● <b>day</b>: The cluster is billed on a pay-per-use basis.</li> </ul>
period_num	No	Integer	<p>Number of periods. This parameter is valid and mandatory only when <b>charge_mode</b> is set to <b>prePaid</b>.</p> <ul style="list-style-type: none"> <li>● If <b>period_type</b> is set to <b>month</b>, the value ranges from 1 to 9.</li> <li>● If <b>period_type</b> is set to <b>year</b>, the value ranges from 1 to 3.</li> </ul>
is_auto_pay	No	Boolean	<p>Whether the order will be automatically paid. This parameter is available for yearly/monthly mode. By default, the automatic payment is disabled. The options are as follows:</p> <ul style="list-style-type: none"> <li>● <b>true</b>: The system automatically selects available discounts and coupons, and then pays for the order with the account balances. If the automatic payment fails, an order in <b>Pending payment</b> state is generated waiting for manual payment.</li> <li>● <b>false</b>: The user needs to pay for the bill after using available discounts and coupons.</li> </ul>

**Table 6-6** Tag parameters

Parameter	Mandatory	Type	Description
key	Yes	String	Tag key. <ul style="list-style-type: none"><li>• It contains a maximum of 36 Unicode characters and cannot be an empty string.</li><li>• The tag key cannot start or end with spaces or contain non-printable ASCII characters (0-31) and special characters (= * &lt; &gt; \,   /).</li><li>• The tag key of a resource must be unique.</li></ul>
value	Yes	String	Tag value. <ul style="list-style-type: none"><li>• The value can contain 0 to 43 unicode characters.</li><li>• The tag value cannot start or end with spaces or contain non-printable ASCII characters (0-31) and special characters (= * &lt; &gt; \,   /).</li></ul>

**Table 6-7** NodeGroup parameters

Parameter	Mandatory	Type	Description
group_name	Yes	String	<p>Node group name. The value can contain a maximum of 64 characters, including uppercase and lowercase letters, digits and underscores (_). The rules for configuring node groups are as follows:</p> <ul style="list-style-type: none"> <li>• <b>master_node_default_group</b>: master node group, which must be included in all cluster types.</li> <li>• <b>core_node_analysis_group</b>: analysis core node group, which must be included in both analysis and hybrid clusters.</li> <li>• <b>core_node_streaming_group</b>: streaming core node group, which must be included in both streaming and hybrid clusters.</li> <li>• <b>task_node_analysis_group</b>: analysis task node group, which can be selected for analysis clusters and hybrid clusters as needed.</li> <li>• <b>task_node_streaming_group</b>: streaming task node group, which can be selected for streaming clusters and hybrid clusters as needed.</li> <li>• <b>node_group{x}</b>: node group of a custom cluster. A maximum of nine such node groups can be added for a custom cluster.</li> </ul>
node_num	Yes	Integer	<p>Number of nodes. The value ranges from 0 to 500. The maximum number of Core and Task nodes is 500.</p>

Parameter	Mandatory	Type	Description
node_size	Yes	String	Instance specifications of a node. Example: <b>c3.4xlarge.2.linux.bigdata</b> The host specifications supported by MRS are determined by CPU, memory, and disk space. For details about instance specifications, see <a href="#">ECS Specifications Used by MRS</a> and <a href="#">BMS Specifications Used by MRS</a> . Obtain the instance specifications of the corresponding version in the corresponding region from the cluster creation page of the MRS management console.
root_volume	No	Volume object	System disk information of the node. This parameter is optional for some VMs or the system disk of the BMS. This parameter is mandatory in other cases. For details about the parameter description, see <a href="#">Table 6-8</a> .
data_volume	No	Volume object	Data disk information. This parameter is mandatory when <b>data_volume_count</b> is not 0. For details about this parameter, see <a href="#">Table 6-8</a> .
data_volume_count	No	Integer	Number of data disks of a node. Value range: 0 to 10
charge_info	No	ChargeInfo object	The billing type of a node group. The billing types of master and core node groups are the same as those of the cluster. The billing type of the task node group can be different. For details about the parameters, see <a href="#">Table 6-5</a> .
auto_scaling_policy	No	auto_scaling_policy object	Autoscaling rule corresponding to the node group. For details about the parameters, see <a href="#">Table 6-11</a> .

Parameter	Mandatory	Type	Description
assigned_roles	No	Array of strings	<p>This parameter is mandatory when the cluster type is CUSTOM. You can specify the roles deployed in a node group. This parameter is a character string array. Each character string represents a role expression.</p> <p>Role expression definition:</p> <ul style="list-style-type: none"> <li>If the role is deployed on all nodes in the node group, set this parameter to <i>&lt;role name&gt;</i>, for example, <b>DataNode</b>.</li> <li>If the role is deployed on a specified subscript node in the node group: <i>&lt;role name&gt;:&lt;index1&gt;,&lt;index2&gt;...,&lt;indexN&gt;</i>, for example, <b>NameNode:1,2</b>. The subscript starts from 1.</li> <li>Some roles support multi-instance deployment (that is, multiple instances of the same role are deployed on a node): <i>&lt;role name&gt;[&lt;instance count&gt;]</i>, for example, <b>EsNode[9]</b>.</li> </ul> <p>For details about available roles, see <a href="#">Roles and components supported by MRS</a>.</p>

**Table 6-8** Volume

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Disk type.</p> <p>The following disk types are supported:</p> <ul style="list-style-type: none"> <li><b>SATA</b>: common I/O disk</li> <li><b>SAS</b>: high I/O disk</li> <li><b>SSD</b>: ultra-high I/O disk</li> </ul>
size	Yes	Integer	<p>Specifies the data disk size, in GB. The value range is <b>10</b> to <b>32768</b>.</p>

**Table 6-9** BootstrapScript

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Name of a bootstrap action script. It must be unique in a cluster.</p> <p>The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and must not start with a space.</p> <p>The value can contain 1 to 64 characters.</p>
uri	Yes	String	<p>Path of a bootstrap action script. Set this parameter to an OBS bucket path or a local VM path.</p> <ul style="list-style-type: none"> <li>• OBS bucket path: Enter a script path manually. For example, enter the path of the public sample script provided by MRS. Example: <b>s3a://bootstrap/presto/presto-install.sh</b>. If <b>dualroles</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>dualroles</b>. If <b>worker</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>worker</b>. Based on the Presto usage habit, you are advised to install <b>dualroles</b> on the active master nodes and <b>worker</b> on the Core nodes.</li> <li>• Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul>
parameters	No	String	Bootstrap action script parameters.
nodes	Yes	Array of strings	Name of the node group where the bootstrap action script is executed
active_master	No	Boolean	<p>Whether the bootstrap action script runs only on active master nodes.</p> <p>The default value is <b>false</b>, indicating that the bootstrap action script can run on all master nodes.</p>

Parameter	Mandatory	Type	Description
before_component_start	No	Boolean	Time when the bootstrap action script is executed. Currently, the following two options are available: <b>Before component start</b> and <b>After component start</b> . The default value is <b>false</b> , indicating that the bootstrap action script is executed after the component is started.
fail_action	Yes	String	Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to be executed. <ul style="list-style-type: none"> <li>• <b>continue</b>: Continue to execute subsequent scripts.</li> <li>• <b>errorout</b>: Stop the action.</li> </ul> The default value is <b>errorout</b> , indicating that the action is stopped. <p><b>NOTE</b></p> You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.
start_time	No	Long	The execution time of one bootstrap action script.
state	No	String	The running status of one bootstrap action script. <ul style="list-style-type: none"> <li>• PENDING</li> <li>• IN_PROGRESS</li> <li>• SUCCESS</li> <li>• FAILURE</li> </ul>



Parameter	Mandatory	Type	Description
action_stages	No	Array of strings	Select the time when the bootstrap action script is executed. <ul style="list-style-type: none"> <li>• <b>BEFORE_COMPONENT_FIRST_START</b>: before initial component starts</li> <li>• <b>AFTER_COMPONENT_FIRST_START</b>: after initial component starts</li> <li>• <b>BEFORE_SCALE_IN</b>: before scale-in</li> <li>• <b>AFTER_SCALE_IN</b>: after scale-in</li> <li>• <b>BEFORE_SCALE_OUT</b>: before scale-out</li> <li>• <b>AFTER_SCALE_OUT</b>: after scale-out</li> </ul>

**Table 6-10** add\_jobs parameters

Parameter	Mandatory	Type	Description
job_type	Yes	Integer	Job type code. <ul style="list-style-type: none"> <li>• 1: MapReduce</li> <li>• 2: Spark</li> <li>• 3: Hive Script</li> <li>• 4: HiveQL (not supported currently)</li> <li>• 5: DistCp, importing and exporting data (not supported currently)</li> <li>• 6: Spark Script</li> <li>• 7: Spark SQL, submitting Spark SQL statements (not supported currently).</li> </ul>
job_name	Yes	String	Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. <b>NOTE</b> Identical job names are allowed but not recommended.

Parameter	Mandatory	Type	Description
jar_path	No	String	<p>Path of the JAR or SQL file for program execution. The parameter must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• Contains a maximum of 1023 characters, excluding special characters such as ; &amp;&gt;&lt;'\$. The parameter value cannot be empty or full of spaces.</li> <li>• Files can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> <li>– OBS: The path must start with <b>s3a://</b>. Files or programs encrypted by KMS are not supported.</li> <li>– HDFS: The path starts with a slash (/).</li> </ul> </li> <li>• Spark Script must end with <b>.sql</b> while MapReduce and Spark Jar must end with <b>.jar.sql</b> and <b>jar</b> are case-insensitive.</li> </ul>
arguments	No	String	<p>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.</p> <p>The parameter contains a maximum of 2047 characters, excluding special characters such as ; &amp;&gt;&lt;'\$, and can be left blank.</p>

Parameter	Mandatory	Type	Description
input	No	String	<p>Address for inputting data.</p> <p>Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> <li>• OBS: The path must start with <b>s3a://</b>. Files or programs encrypted by KMS are not supported.</li> <li>• HDFS: The path starts with a slash (/).</li> </ul> <p>The parameter contains a maximum of 1023 characters, excluding special characters such as ; &amp;&gt;'&lt;\$, and can be left blank.</p>
output	No	String	<p>Address for outputting data.</p> <p>Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> <li>• OBS: The path must start with <b>s3a://</b>.</li> <li>• HDFS: The path starts with a slash (/).</li> </ul> <p>If the specified path does not exist, the system will automatically create it.</p> <p>The parameter contains a maximum of 1023 characters, excluding special characters such as ; &amp;&gt;'&lt;\$, and can be left blank.</p>
job_log	No	String	<p>Path for storing job logs that record job running status.</p> <p>Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> <li>• OBS: The path must start with <b>s3a://</b>.</li> <li>• HDFS: The path starts with a slash (/).</li> </ul> <p>The parameter contains a maximum of 1023 characters, excluding special characters such as ; &amp;&gt;'&lt;\$, and can be left blank.</p>

Parameter	Mandatory	Type	Description
shutdown_cluster	No	Boolean	Whether to delete the cluster after the job execution is complete. <ul style="list-style-type: none"> <li>• <b>true</b>: Yes</li> <li>• <b>false</b>: No</li> </ul>
file_action	No	String	Data import and export. <ul style="list-style-type: none"> <li>• import</li> <li>• export</li> </ul>
submit_job_once_cluster_run	Yes	Boolean	<ul style="list-style-type: none"> <li>• <b>true</b>: Submit a job during cluster creation.</li> <li>• <b>false</b>: Submit a job after the cluster is created.</li> </ul> Set this parameter to <b>true</b> in this example.
hql	No	String	HiveQL statement
hive_script_path	No	String	SQL program path. This parameter is needed by Spark Script and Hive Script jobs only, and must meet the following requirements: <ul style="list-style-type: none"> <li>• Contains a maximum of 1023 characters, excluding special characters such as ; &amp;&gt;&lt;'\$. The address cannot be empty or full of spaces.</li> <li>• Files can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> <li>– OBS: The path must start with <b>s3a://</b>. Files or programs encrypted by KMS are not supported.</li> <li>– HDFS: The path starts with a slash (/).</li> </ul> </li> <li>• Ends with <b>.sql</b>. <b>sql</b> is case-insensitive.</li> </ul>

**Table 6-11** auto\_scaling\_policy parameters

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	Whether to enable the auto scaling rule.

Parameter	Mandatory	Type	Description
min_capacity	Yes	Integer	Minimum number of nodes allowed in the node group. Value range: [0, 500]
max_capacity	Yes	Integer	Maximum number of nodes in the node group. Value range: [0, 500]
resources_plans	No	Array of resources_plan objects	Resource plan list. For details, see <a href="#">Table 6-12</a> . If this parameter is left blank, the resource plan is disabled.  When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured.
exec_scripts	No	Array of scale_script objects	List of custom scaling automation scripts. For details, see <a href="#">Table 6-13</a> . If this parameter is left blank, a hook script is disabled.
rules	No	Array of rules objects	List of auto scaling rules. For details, see <a href="#">Table 6-14</a> .  When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured.

**Table 6-12** ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	Cycle type of a resource plan. Currently, only the following cycle type is supported: <ul style="list-style-type: none"> <li>daily</li> </ul>
start_time	Yes	String	The start time of a resource plan. The value is in the format of <b>hour:minute</b> , indicating that the time ranges from 00:00 to 23:59.
end_time	Yes	String	End time of a resource plan. The value is in the same format as that of <b>start_time</b> . The interval between <b>end_time</b> and <b>start_time</b> must be greater than or equal to 30 minutes.

Parameter	Mandatory	Type	Description
min_capacity	Yes	Integer	Minimum number of the preserved nodes in a node group in a resource plan. Value range: [0, 500]
max_capacity	Yes	Integer	Maximum number of the preserved nodes in a node group in a resource plan. Value range: [0, 500]
effective_days	No	Array of strings	Effective date of a resource plan. If this parameter is left blank, it indicates that the resource plan takes effect every day. The options are as follows: <b>MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</b>

**Table 6-13** scale\_script parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Name of a custom automation script. It must be unique in a same cluster. 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.
uri	Yes	String	Path of a custom automation script. Set this parameter to an OBS bucket path or a local VM path. <ul style="list-style-type: none"> <li>• OBS bucket path: Enter a script path manually. for example, <b>s3a://XXX/scale.sh</b>.</li> <li>• Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul>

Parameter	Mandatory	Type	Description
parameters	No	String	<p>Parameters of a custom automation script.</p> <ul style="list-style-type: none"> <li>Multiple parameters are separated by space.</li> <li>The following predefined system parameters can be transferred: <ul style="list-style-type: none"> <li><i>\${mrs_scale_node_num}</i>: Number of the nodes to be added or removed</li> <li><i>\${mrs_scale_type}</i>: Scaling type. The value can be <b>scale_out</b> or <b>scale_in</b>.</li> <li><i>\${mrs_scale_node_hostnames}</i>: Host names of the nodes to be added or removed</li> <li><i>\${mrs_scale_node_ips}</i>: IP addresses of the nodes to be added or removed</li> <li><i>\${mrs_scale_rule_name}</i>: Name of the rule that triggers auto scaling</li> </ul> </li> <li>Other user-defined parameters are used in the same way as those of common shell scripts. Parameters are separated by space.</li> </ul>
nodes	Yes	List<String>	Type of a node where the custom automation script is executed. The node type can be Master, Core, or Task.
active_master	No	Boolean	<p>Whether the custom automation script runs only on the active master node.</p> <p>The default value is <b>false</b>, indicating that the custom automation script can run on all Master nodes.</p>

Parameter	Mandatory	Type	Description
action_stage	Yes	String	Time when a script is executed. The following four options are supported: <ul style="list-style-type: none"> <li>• <b>before_scale_out</b>: before scale-out</li> <li>• <b>before_scale_in</b>: before scale-in</li> <li>• <b>after_scale_out</b>: after scale-out</li> <li>• <b>after_scale_in</b>: after scale-in</li> </ul>
fail_action	Yes	String	Whether to continue to execute subsequent scripts and create a cluster after the custom automation script fails to be executed. <ul style="list-style-type: none"> <li>• <b>continue</b>: Continue to execute subsequent scripts.</li> <li>• <b>errorout</b>: Stop the action.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the custom automation script is executed successfully.</li> <li>• The scale-in operation cannot be undone. Therefore, <b>fail_action</b> must be set to <b>continue</b> for the scripts that are executed after scale-in.</li> </ul>

**Table 6-14** rules parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Name of an auto scaling rule. A cluster name can contain only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. Rule names must be unique in a node group.



Parameter	Mandatory	Type	Description
description	No	String	Description about an auto scaling rule. It contains a maximum of 1024 characters.
adjustment_type	Yes	String	Auto scaling rule adjustment type. Possible values: <ul style="list-style-type: none"> <li>● <b>scale_out</b>: cluster scale-out</li> <li>● <b>scale_in</b>: cluster scale-in</li> </ul>
cool_down_minutes	Yes	Integer	Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute. Value range: 0 to 10,080. One week is equal to 10,080 minutes.
scaling_adjustment	Yes	Integer	Number of nodes that can be adjusted once. Value range: [1, 100]
trigger	Yes	Trigger object	Condition for triggering a rule. For details, see <a href="#">Table 6-15</a> .

**Table 6-15** trigger parameters

Parameter	Mandatory	Type	Description
metric_name	Yes	String	Metric name. This triggering condition makes a judgment according to the value of the metric. A metric name contains a maximum of 64 characters.
metric_value	Yes	String	Metric threshold to trigger a rule. The value must be an integer or a number with two decimal places.
comparison_operator	No	String	Metric judgment logic operator. Possible values: <ul style="list-style-type: none"> <li>● <b>LT</b>: less than</li> <li>● <b>GT</b>: greater than</li> <li>● <b>LTOE</b>: less than or equal to</li> <li>● <b>GTOE</b>: greater than or equal to</li> </ul>

Parameter	Mandatory	Type	Description
evaluation_periods	Yes	Integer	Number of consecutive five-minute periods, during which a metric threshold is reached. Value range: 1 to 288

**Table 6-16** ComponentConfig

Parameter	Mandatory	Type	Description
component_name	Yes	String	The component name
configs	No	Array of Config objects	The component configuration item list. For details about this parameter, see <a href="#">Table 6-17</a> .

**Table 6-17** Config

Parameter	Mandatory	Type	Description
key	Yes	String	The configuration name. Only the configuration names displayed on the MRS component configuration page are supported.
value	Yes	String	The configuration value
config_file_name	Yes	String	The configuration file name. Only the file names displayed on the MRS component configuration page are supported.

**Table 6-18** SmnNotify

Parameter	Mandatory	Type	Description
topic_urn	No	String	SMN topic URN. This parameter is mandatory if alarm notifications are enabled.

Parameter	Mandatory	Type	Description
subscription_name	No	String	Name of a subscription rule. If this parameter is not set, the default value <b>default_alert_rule</b> will be used.

## Response Parameters

Status code: 200

Table 6-19 Response parameters

Parameter	Type	Description
cluster_id	String	Cluster ID, which is returned by the system after the cluster is created.

## Example Request

- Create an MRS 3.2.0-LTS.1 cluster for analysis. There are a Master node group with two nodes, a Core node group with three nodes, and a Task node group with three nodes. Autoscaling is enabled from 12:00 to 13:00 every Monday.  
POST /v2/{project\_id}/clusters

```
{
  "cluster_version": "MRS 3.2.0-LTS.1",
  "cluster_name": "mrs_DyJA_dm",
  "cluster_type": "ANALYSIS",
  "charge_info": {
    "charge_mode": "postPaid"
  },
  "region": "",
  "availability_zone": "",
  "vpc_name": "vpc-37cd",
  "subnet_id": "1f8c5ca6-1f66-4096-bb00-baf175954f6e",
  "subnet_name": "subnet",
  "components": "Hadoop,Spark2x,HBase,Hive,Hue,Flink,Oozie,Ranger,Tez",
  "safe_mode": "KERBEROS",
  "manager_admin_password": "your password",
  "login_mode": "PASSWORD",
  "node_root_password": "your password",
  "log_collection": 1,
  "mrs_ecs_default_agency": "MRS_ECS_DEFAULT_AGENCY",
  "tags": [ {
    "key": "tag1",
    "value": "111"
  }, {
    "key": "tag2",
    "value": "222"
  } ],
  "node_groups": [ {
    "group_name": "master_node_default_group",
    "node_num": 2,
    "node_size": "rc3.4xlarge.4.linux.bigdata",
    "root_volume": {
      "type": "SAS",
```

```
    "size" : 480
  },
  "data_volume" : {
    "type" : "SAS",
    "size" : 600
  },
  "data_volume_count" : 1
}, {
  "group_name" : "core_node_analysis_group",
  "node_num" : 3,
  "node_size" : "rc3.4xlarge.4.linux.bigdata",
  "root_volume" : {
    "type" : "SAS",
    "size" : 480
  },
  "data_volume" : {
    "type" : "SAS",
    "size" : 600
  },
  "data_volume_count" : 1
}, {
  "group_name" : "task_node_analysis_group",
  "node_num" : 3,
  "node_size" : "rc3.4xlarge.4.linux.bigdata",
  "root_volume" : {
    "type" : "SAS",
    "size" : 480
  },
  "data_volume" : {
    "type" : "SAS",
    "size" : 600
  },
  "data_volume_count" : 1,
  "auto_scaling_policy" : {
    "auto_scaling_enable" : true,
    "min_capacity" : 0,
    "max_capacity" : 1,
    "resources_plans" : [ {
      "period_type" : "daily",
      "start_time" : "12:00",
      "end_time" : "13:00",
      "min_capacity" : 2,
      "max_capacity" : 3,
      "effective_days" : [ "MONDAY" ]
    } ],
    "exec_scripts" : [ {
      "name" : "test",
      "uri" : "s3a://obs-mrstest/bootstrap/basic_success.sh",
      "parameters" : "",
      "nodes" : [ "master_node_default_group", "core_node_analysis_group",
"task_node_analysis_group" ],
      "active_master" : false,
      "action_stage" : "before_scale_out",
      "fail_action" : "continue"
    } ],
    "rules" : [ {
      "name" : "default-expand-1",
      "description" : "",
      "adjustment_type" : "scale_out",
      "cool_down_minutes" : 5,
      "scaling_adjustment" : "1",
      "trigger" : {
        "metric_name" : "YARNAppRunning",
        "metric_value" : 100,
        "comparison_operator" : "GTOE",
        "evaluation_periods" : "1"
      }
    } ]
  } ]
}
```

- ```
    }  
  }  
}
```
- Create an MRS 3.2.0-LTS.1 cluster for stream analysis. There are a Master node group with two nodes, a Core node group with three nodes, and a Task node group with no node. Autoscaling is enabled from 12:00 to 13:00 every Monday.

POST /v2/{project\_id}/clusters

```
{  
  "cluster_version": "MRS 3.2.0-LTS.1",  
  "cluster_name": "mrs_Dokle_dm",  
  "cluster_type": "STREAMING",  
  "charge_info": {  
    "charge_mode": "postPaid"  
  },  
  "region": "",  
  "availability_zone": "",  
  "vpc_name": "vpc-37cd",  
  "subnet_id": "1f8c5ca6-1f66-4096-bb00-baf175954f6e",  
  "subnet_name": "subnet",  
  "components": "Storm,Kafka,Flume,Ranger",  
  "safe_mode": "KERBEROS",  
  "manager_admin_password": "your password",  
  "login_mode": "PASSWORD",  
  "node_root_password": "your password",  
  "log_collection": 1,  
  "mrs_ecs_default_agency": "MRS_ECS_DEFAULT_AGENCY",  
  "tags": [ {  
    "key": "tag1",  
    "value": "111"  
  }, {  
    "key": "tag2",  
    "value": "222"  
  } ],  
  "node_groups": [ {  
    "group_name": "master_node_default_group",  
    "node_num": 2,  
    "node_size": "rc3.4xlarge.4.linux.bigdata",  
    "root_volume": {  
      "type": "SAS",  
      "size": 480  
    },  
    "data_volume": {  
      "type": "SAS",  
      "size": 600  
    },  
    "data_volume_count": 1  
  }, {  
    "group_name": "core_node_streaming_group",  
    "node_num": 3,  
    "node_size": "rc3.4xlarge.4.linux.bigdata",  
    "root_volume": {  
      "type": "SAS",  
      "size": 480  
    },  
    "data_volume": {  
      "type": "SAS",  
      "size": 600  
    },  
    "data_volume_count": 1  
  }, {  
    "group_name": "task_node_streaming_group",  
    "node_num": 0,  
    "node_size": "rc3.4xlarge.4.linux.bigdata",  
    "root_volume": {  
      "type": "SAS",  
      "size": 480  
    },  
  } ],  
}
```

```
"data_volume" : {
  "type" : "SAS",
  "size" : 600
},
"data_volume_count" : 1,
"auto_scaling_policy" : {
  "auto_scaling_enable" : true,
  "min_capacity" : 0,
  "max_capacity" : 1,
  "resources_plans" : [ {
    "period_type" : "daily",
    "start_time" : "12:00",
    "end_time" : "13:00",
    "min_capacity" : 2,
    "max_capacity" : 3,
    "effective_days" : [ "MONDAY" ]
  } ],
  "rules" : [ {
    "name" : "default-expand-1",
    "description" : "",
    "adjustment_type" : "scale_out",
    "cool_down_minutes" : 5,
    "scaling_adjustment" : "1",
    "trigger" : {
      "metric_name" : "StormSlotAvailablePercentage",
      "metric_value" : 100,
      "comparison_operator" : "LTOE",
      "evaluation_periods" : "1"
    }
  } ]
} ]
}
```

- Create an MRS 3.2.0-LTS.1 cluster for hybrid analysis. There are a Master node group with two nodes, two Core node groups with three nodes in each, and two Task node groups with three nodes in one group and one node in the other group.

POST /v2/{project\_id}/clusters

```
{
  "cluster_version" : "MRS 3.2.0-LTS.1",
  "cluster_name" : "mrs_onmm_dm",
  "cluster_type" : "MIXED",
  "charge_info" : {
    "charge_mode" : "postPaid"
  },
  "region" : "",
  "availability_zone" : "",
  "vpc_name" : "vpc-37cd",
  "subnet_id" : "1f8c5ca6-1f66-4096-bb00-baf175954f6e",
  "subnet_name" : "subnet",
  "components" : "Hadoop,Spark2x,HBase,Hive,Hue,Loader,Kafka,Storm,Flume,Flink,Oozie,Ranger,Tez",
  "safe_mode" : "KERBEROS",
  "manager_admin_password" : "your password",
  "login_mode" : "PASSWORD",
  "node_root_password" : "your password",
  "log_collection" : 1,
  "mrs_ecs_default_agency" : "MRS_ECS_DEFAULT_AGENCY",
  "tags" : [ {
    "key" : "tag1",
    "value" : "111"
  }, {
    "key" : "tag2",
    "value" : "222"
  } ],
  "node_groups" : [ {
    "group_name" : "master_node_default_group",
    "node_num" : 2,
```

```

"node_size" : "Sit3.4xlarge.4.linux.bigdata",
"root_volume" : {
  "type" : "SAS",
  "size" : 480
},
"data_volume" : {
  "type" : "SAS",
  "size" : 600
},
"data_volume_count" : 1
}, {
"group_name" : "core_node_streaming_group",
"node_num" : 3,
"node_size" : "Sit3.4xlarge.4.linux.bigdata",
"root_volume" : {
  "type" : "SAS",
  "size" : 480
},
"data_volume" : {
  "type" : "SAS",
  "size" : 600
},
"data_volume_count" : 1
}, {
"group_name" : "core_node_analysis_group",
"node_num" : 3,
"node_size" : "Sit3.4xlarge.4.linux.bigdata",
"root_volume" : {
  "type" : "SAS",
  "size" : 480
},
"data_volume" : {
  "type" : "SAS",
  "size" : 600
},
"data_volume_count" : 1
}, {
"group_name" : "task_node_analysis_group",
"node_num" : 1,
"node_size" : "Sit3.4xlarge.4.linux.bigdata",
"root_volume" : {
  "type" : "SAS",
  "size" : 480
},
"data_volume" : {
  "type" : "SAS",
  "size" : 600
},
"data_volume_count" : 1
}, {
"group_name" : "task_node_streaming_group",
"node_num" : 0,
"node_size" : "Sit3.4xlarge.4.linux.bigdata",
"root_volume" : {
  "type" : "SAS",
  "size" : 480
},
"data_volume" : {
  "type" : "SAS",
  "size" : 600
},
"data_volume_count" : 1
}
}]]
}

```

- Create a cluster where custom management nodes and control nodes are the same nodes. The cluster version is MRS 3.2.0-LTS.1. There is a Master node group with three nodes, two Core node groups with three nodes in one group and one node in the other group.

```
POST /v2/{project_id}/clusters

{
  "cluster_version": "MRS 3.2.0-LTS.1",
  "cluster_name": "mrs_heshe_dm",
  "cluster_type": "CUSTOM",
  "charge_info": {
    "charge_mode": "postPaid"
  },
  "region": "",
  "availability_zone": "",
  "vpc_name": "vpc-37cd",
  "subnet_id": "1f8c5ca6-1f66-4096-bb00-baf175954f6e",
  "subnet_name": "subnet",
  "components":
  "Hadoop,Spark2x,HBase,Hive,Hue,Kafka,Flume,Flink,Oozie,HetuEngine,Ranger,Tez,ZooKeeper,ClickHous
e",
  "safe_mode": "KERBEROS",
  "manager_admin_password": "your password",
  "login_mode": "PASSWORD",
  "node_root_password": "your password",
  "mrs_ecs_default_agency": "MRS_ECS_DEFAULT_AGENCY",
  "template_id": "mgmt_control_combined_v2",
  "log_collection": 1,
  "tags": [ {
    "key": "tag1",
    "value": "111"
  }, {
    "key": "tag2",
    "value": "222"
  } ],
  "node_groups": [ {
    "group_name": "master_node_default_group",
    "node_num": 3,
    "node_size": "Sit3.4xlarge.4.linux.bigdata",
    "root_volume": {
      "type": "SAS",
      "size": 480
    },
    "data_volume": {
      "type": "SAS",
      "size": 600
    },
    "data_volume_count": 1,
    "assigned_roles": [ "OMSServer:1,2", "SlapdServer:1,2", "KerberosServer:1,2", "KerberosAdmin:1,2",
"quorumpeer:1,2,3", "NameNode:2,3", "Zkfc:2,3", "JournalNode:1,2,3", "ResourceManager:2,3",
"JobHistoryServer:2,3", "DBServer:1,3", "Hue:1,3", "LoaderServer:1,3", "MetaStore:1,2,3",
"WebHCat:1,2,3", "HiveServer:1,2,3", "HMaster:2,3", "MonitorServer:1,2", "Nimbus:1,2", "UI:1,2",
"JDBCServer2x:1,2,3", "JobHistory2x:2,3", "SparkResource2x:1,2,3", "oozie:2,3", "LoadBalancer:2,3",
"TezUI:1,3", "TimelineServer:3", "RangerAdmin:1,2", "UserSync:2", "TagSync:2", "KerberosClient",
"SlapdClient", "meta", "HSConsole:2,3", "FlinkResource:1,2,3", "DataNode:1,2,3",
"NodeManager:1,2,3", "IndexServer2x:1,2", "ThriftServer:1,2,3", "RegionServer:1,2,3",
"ThriftServer1:1,2,3", "RESTServer:1,2,3", "Broker:1,2,3", "Supervisor:1,2,3", "Logviewer:1,2,3",
"Flume:1,2,3", "HSBroker:1,2,3" ]
  }, {
    "group_name": "node_group_1",
    "node_num": 3,
    "node_size": "Sit3.4xlarge.4.linux.bigdata",
    "root_volume": {
      "type": "SAS",
      "size": 480
    },
    "data_volume": {
      "type": "SAS",
      "size": 600
    },
    "data_volume_count": 1,
    "assigned_roles": [ "DataNode", "NodeManager", "RegionServer", "Flume:1", "Broker",
"Supervisor", "Logviewer", "HBaseIndexer", "KerberosClient", "SlapdClient", "meta", "HSBroker:1,2",
```



```
"ThriftServer", "ThriftServer1", "RETServer", "FlinkResource" ]
}, {
  "group_name": "node_group_2",
  "node_num": 1,
  "node_size": "Sit3.4xlarge.4.linux.bigdata",
  "root_volume": {
    "type": "SAS",
    "size": 480
  },
  "data_volume": {
    "type": "SAS",
    "size": 600
  },
  "data_volume_count": 1,
  "assigned_roles": [ "NodeManager", "KerberosClient", "SlapdClient", "meta", "FlinkResource" ]
}]
}
```

- Create a cluster where custom management nodes and control nodes are independent nodes. The cluster version is MRS 3.2.0-LTS.1. There a Master node group with five nodes and a Core node group with three nodes.

POST /v2/{project\_id}/clusters

```
{
  "cluster_version": "MRS 3.2.0-LTS.1",
  "cluster_name": "mrs_jdRU_dm01",
  "cluster_type": "CUSTOM",
  "charge_info": {
    "charge_mode": "postPaid"
  },
  "region": "",
  "availability_zone": "",
  "vpc_name": "vpc-37cd",
  "subnet_id": "1f8c5ca6-1f66-4096-bb00-baf175954f6e",
  "subnet_name": "subnet",
  "components":
  "Hadoop,Spark2x,HBase,Hive,Hue,Kafka,Flume,Flink,Oozie,HetuEngine,Ranger,Tez,Ranger,Tez,ZooKeeper,ClickHouse",
  "safe_mode": "KERBEROS",
  "manager_admin_password": "your password",
  "login_mode": "PASSWORD",
  "node_root_password": "your password",
  "mrs_ecs_default_agency": "MRS_ECS_DEFAULT_AGENCY",
  "log_collection": 1,
  "template_id": "mgmt_control_separated_v2",
  "tags": [ {
    "key": "aaa",
    "value": "111"
  }, {
    "key": "bbb",
    "value": "222"
  } ],
  "node_groups": [ {
    "group_name": "master_node_default_group",
    "node_num": 5,
    "node_size": "rc3.4xlarge.4.linux.bigdata",
    "root_volume": {
      "type": "SAS",
      "size": 480
    },
    "data_volume": {
      "type": "SAS",
      "size": 600
    },
    "data_volume_count": 1,
    "assigned_roles": [ "OMSServer:1,2", "SlapdServer:3,4", "KerberosServer:3,4", "KerberosAdmin:3,4", "quorumpeer:3,4,5", "NameNode:4,5", "Zkfc:4,5", "JournalNode:1,2,3,4,5", "ResourceManager:4,5", "JobHistoryServer:4,5", "DBServer:3,5", "Hue:1,2", "LoaderServer:1,2", "MetaStore:1,2,3,4,5", "WebHCat:1,2,3,4,5", "HiveServer:1,2,3,4,5", "HMaster:4,5", "MonitorServer:1,2", "Nimbus:1,2",
```

```
"UI:1,2", "JDBCServer2x:1,2,3,4,5", "JobHistory2x:4,5", "SparkResource2x:1,2,3,4,5", "oozie:1,2",
"LoadBalancer:1,2", "TezUI:1,2", "TimelineServer:5", "RangerAdmin:1,2", "KerberosClient",
"SlapdClient", "meta", "HSConsole:1,2", "FlinkResource:1,2,3,4,5", "DataNode:1,2,3,4,5",
"NodeManager:1,2,3,4,5", "IndexServer2x:1,2", "ThriftServer:1,2,3,4,5", "RegionServer:1,2,3,4,5",
"ThriftServer1:1,2,3,4,5", "RETSer:1,2,3,4,5", "Broker:1,2,3,4,5", "Supervisor:1,2,3,4,5",
"Logviewer:1,2,3,4,5", "Flume:1,2,3,4,5", "HBaseIndexer:1,2,3,4,5", "TagSync:1", "UserSync:1" ]
}, {
  "group_name": "node_group_1",
  "node_num": 3,
  "node_size": "rc3.4xlarge.4.linux.bigdata",
  "root_volume": {
    "type": "SAS",
    "size": 480
  },
  "data_volume": {
    "type": "SAS",
    "size": 600
  },
  "data_volume_count": 1,
  "assigned_roles": [ "DataNode", "NodeManager", "RegionServer", "Flume:1", "Broker",
"Supervisor", "Logviewer", "HBaseIndexer", "KerberosClient", "SlapdClient", "meta", "HSBroker:1,2",
"ThriftServer", "ThriftServer1", "RETSer", "FlinkResource" ]
}]
}
```

- Create a cluster where data nodes are deployed independently from other nodes. The cluster version is MRS 3.2.0-LTS.1. There are a Master node group with nine nodes, four Core node groups with three nodes in each group.

POST /v2/{project\_id}/clusters

```
{
  "cluster_version": "MRS 3.2.0-LTS.1",
  "cluster_name": "mrs_jdRU_dm02",
  "cluster_type": "CUSTOM",
  "charge_info": {
    "charge_mode": "postPaid"
  },
  "region": "",
  "availability_zone": "",
  "vpc_name": "vpc-37cd",
  "subnet_id": "1f8c5ca6-1f66-4096-bb00-baf175954f6e",
  "subnet_name": "subnet",
  "components":
"Hadoop,Spark2x,HBase,Hive,Hue,Kafka,Flume,Flink,Oozie,Ranger,Tez,Ranger,Tez,ZooKeeper,ClickHous
e",
  "safe_mode": "KERBEROS",
  "manager_admin_password": "your password",
  "login_mode": "PASSWORD",
  "node_root_password": "your password",
  "mrs_ecs_default_agency": "MRS_ECS_DEFAULT_AGENCY",
  "template_id": "mgmt_control_data_separated_v2",
  "log_collection": 1,
  "tags": [ {
    "key": "aaa",
    "value": "111"
  }, {
    "key": "bbb",
    "value": "222"
  } ],
  "node_groups": [ {
    "group_name": "master_node_default_group",
    "node_num": 9,
    "node_size": "rc3.4xlarge.4.linux.bigdata",
    "root_volume": {
      "type": "SAS",
      "size": 480
    },
    "data_volume": {
      "type": "SAS",
```

```
"size" : 600
},
"data_volume_count" : 1,
"assigned_roles" : [ "OMSServer:1,2", "SlapdServer:5,6", "KerberosServer:5,6", "KerberosAdmin:5,6",
"quorumpeer:5,6,7,8,9", "NameNode:3,4", "Zkfc:3,4", "JournalNode:5,6,7", "ResourceManager:8,9",
"JobHistoryServer:8", "DBServer:8,9", "Hue:8,9", "FlinkResource:3,4", "LoaderServer:3,5",
"MetaStore:8,9", "WebHCat:5", "HiveServer:8,9", "HMaster:8,9", "FTP-Server:3,4", "MonitorServer:3,4",
"Nimbus:8,9", "UI:8,9", "JDBCServer2x:8,9", "JobHistory2x:8,9", "SparkResource2x:5,6,7", "oozie:4,5",
"EsMaster:7,8,9", "LoadBalancer:8,9", "TezUI:5,6", "TimelineServer:5", "RangerAdmin:4,5",
"UserSync:5", "TagSync:5", "KerberosClient", "SlapdClient", "meta", "HSBroker:5", "HSConsole:3,4",
"FlinkResource:3,4" ]
}, {
"group_name" : "node_group_1",
"node_num" : 3,
"node_size" : "rc3.4xlarge.4.linux.bigdata",
"root_volume" : {
"type" : "SAS",
"size" : 480
},
"data_volume" : {
"type" : "SAS",
"size" : 600
},
"data_volume_count" : 1,
"assigned_roles" : [ "DataNode", "NodeManager", "RegionServer", "Flume:1", "GraphServer",
"KerberosClient", "SlapdClient", "meta", "HSBroker:1,2" ]
}, {
"group_name" : "node_group_2",
"node_num" : 3,
"node_size" : "rc3.4xlarge.4.linux.bigdata",
"root_volume" : {
"type" : "SAS",
"size" : 480
},
"data_volume" : {
"type" : "SAS",
"size" : 600
},
"data_volume_count" : 1,
"assigned_roles" : [ "HBaseIndexer", "SolrServer[3]", "EsNode[2]", "KerberosClient", "SlapdClient",
"meta", "SolrServerAdmin:1,2" ]
}, {
"group_name" : "node_group_3",
"node_num" : 3,
"node_size" : "rc3.4xlarge.4.linux.bigdata",
"root_volume" : {
"type" : "SAS",
"size" : 480
},
"data_volume" : {
"type" : "SAS",
"size" : 600
},
"data_volume_count" : 1,
"assigned_roles" : [ "Redis[2]", "KerberosClient", "SlapdClient", "meta" ]
}, {
"group_name" : "node_group_4",
"node_num" : 3,
"node_size" : "rc3.4xlarge.4.linux.bigdata",
"root_volume" : {
"type" : "SAS",
"size" : 480
},
"data_volume" : {
"type" : "SAS",
"size" : 600
},
"data_volume_count" : 1,
"assigned_roles" : [ "Broker", "Supervisor", "Logviewer", "KerberosClient", "SlapdClient", "meta" ]
```

```
}]  
}
```

## Example Response

- Example of a successful response

```
{  
  "cluster_id": "da1592c2-bb7e-468d-9ac9-83246e95447a"  
}
```

- Example of a failed response

```
{  
  "error_code": "MRS.0002",  
  "error_msg": "The parameter is invalid."  
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.1.2 Changing a Cluster Name

### Function

This API is used to change a cluster name.

### Constraints

None

### Debugging

You can debug this API through automatic authentication in API Explorer. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

PUT /v2/{project\_id}/clusters/{cluster\_id}/cluster-name

**Table 6-20** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                  |
|------------|-----------|--------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

| Parameter  | Mandatory | Type   | Description                                                                                                  |
|------------|-----------|--------|--------------------------------------------------------------------------------------------------------------|
| cluster_id | Yes       | String | The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> . |

## Request Parameters

**Table 6-21** Request body parameter

| Parameter    | Mandatory | Type   | Description           |
|--------------|-----------|--------|-----------------------|
| cluster_name | Yes       | String | The new cluster name. |

## Response Parameters

**Status code: 200**

**Table 6-22** Response body parameter

| Parameter | Type   | Description                                                                                                                                                                                |
|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| result    | String | The operation result of the mapping update request. Value <b>**succeeded**</b> indicates that the operation is successful, and value <b>**failed**</b> indicates that the operation fails. |

**Status code: 400**

**Table 6-23** Response body parameters

| Parameter  | Type   | Description        |
|------------|--------|--------------------|
| error_code | String | The error code.    |
| error_msg  | String | The error message. |

## Example Request

Change the MRS cluster name to **mrs\_jdRU\_dm01**.

```
{
  "cluster_name": "mrs_jdRU_dm01"
}
```

## Example Response

**Status code: 200**

The cluster name is changed.

```
{  
  "result" : "succeeded"  
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.1.3 Creating a Cluster and Submitting a Job

### Function

This API is used to create an MRS cluster, submit a job, and terminate the cluster after the job is complete. This API is supported in MRS 1.8.9 or later. Before using this API, you need to obtain the following resource information:

- Create or query a VPC and subnet.
- Create or query a key pair using an ECS.
- Obtain the region information by referring to [Endpoints](#).
- Obtain the MRS version and the components supported by the MRS version by referring to [Obtaining the MRS Cluster Information](#).

### Constraints

None

### Debugging

You can debug this API through automatic authentication in API Explorer. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

POST /v2/{project\_id}/run-job-flow

**Table 6-24** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                  |
|------------|-----------|--------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

## Request Parameters

**Table 6-25** Request body parameters

| Parameter       | Mandatory | Type              | Description                                                                                                                                                                                                                                                                                              |
|-----------------|-----------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| is_dec_project  | No        | Boolean           | Whether the cluster is specific for the DeC. The default value is <b>false</b> .                                                                                                                                                                                                                         |
| cluster_version | Yes       | String            | The cluster version. Example: MRS 3.1.0.                                                                                                                                                                                                                                                                 |
| cluster_name    | Yes       | String            | The cluster name, which must be unique.<br>A cluster name can contain only 1 to 64 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed.                                                                                                                                      |
| cluster_type    | Yes       | String            | The cluster type. The options are as follows: <ul style="list-style-type: none"> <li>● <b>ANALYSIS</b>: analysis cluster</li> <li>● <b>STREAMING</b>: streaming cluster</li> <li>● <b>MIXED</b>: hybrid cluster</li> <li>● <b>CUSTOM</b>: custom cluster, which is supported only by MRS 3.x.</li> </ul> |
| charge_info     | No        | ChargeInfo object | The billing type. For details, see <a href="#">Table 6-30</a> .                                                                                                                                                                                                                                          |
| region          | Yes       | String            | Information about the region where the cluster is located. For details, see <a href="#">Endpoints</a> .                                                                                                                                                                                                  |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vpc_name  | Yes       | String | <p>The name of the VPC where the subnet is located. Obtain the VPC name by performing the following operations on the VPC management console:</p> <ol style="list-style-type: none"> <li>1. Log in to the VPC management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>. On the <b>Virtual Private Cloud</b> page, obtain the VPC name from the list.</li> </ol>                                                                                                                                                                                                                                                                                                       |
| subnet_id | No        | String | <p>The subnet ID. Obtain the subnet ID by performing the following operations on the VPC management console:</p> <ol style="list-style-type: none"> <li>1. Log in to the VPC management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> <li>3. Locate the row containing the target VPC and click the number in the <b>Subnets</b> column to view the subnet information.</li> <li>4. Click the subnet name to obtain the network ID. At least one of <b>subnet_id</b> and <b>subnet_name</b> must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. <b>subnet_id</b> is recommended.</li> </ol> |



| Parameter             | Mandatory | Type                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------|-----------|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| subnet_name           | Yes       | String                                   | The subnet name. Obtain the subnet name by performing the following operations on the VPC management console: <ol style="list-style-type: none"> <li>1. Log in to the management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> <li>3. Locate the row containing the target VPC and click the number in the <b>Subnets</b> column to obtain the subnet name. At least one of <b>subnet_id</b> and <b>subnet_name</b> must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. If only <b>subnet_name</b> is configured and subnets with the same name exist in the VPC, the first subnet name in the VPC is used when a cluster is created. <b>subnet_id</b> is recommended.</li> </ol> |
| components            | Yes       | String                                   | List of component names, which are separated by commas (.). For details about the components that are supported, see "Components Supported by MRS" in <a href="#">Obtaining the MRS Cluster Information</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| external_data sources | No        | Array of ClusterDataConnectorMap objects | When deploying components such as Hive and Ranger, you can associate data connections and store metadata in associated databases. For details about the parameters, see <a href="#">Table 6-26</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| availability_zone     | Yes       | String                                   | The AZ name. Multi-AZ clusters are not supported. For details about AZs, see <a href="#">Endpoints</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

| Parameter                          | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| security_groups_id                 | No        | String  | <p>The ID of the security group configured for the cluster.</p> <ul style="list-style-type: none"> <li>If this parameter is left blank, MRS automatically creates a security group whose name starts with <b>mrs_{cluster_name}</b>.</li> <li>If this parameter is configured, a fixed security group is used to create a cluster. The transferred ID must be the security group ID owned by the current tenant. The security group must include an inbound rule in which all protocols and all ports are allowed and the source is the IP address of the specified node on the management plane.</li> <li>Multiple security group IDs are supported and separated by commas (,).</li> </ul> |
| auto_create_default_security_group | No        | Boolean | <p>Whether to create the default security group for the MRS cluster. The default value is <b>false</b>. If this parameter is set to <b>true</b>, the default security group will be created for the cluster regardless of whether <b>security_groups_id</b> is specified.</p>                                                                                                                                                                                                                                                                                                                                                                                                                |

| Parameter              | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| safe_mode              | Yes       | String | <p>The running mode of an MRS cluster. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>SIMPLE</b>: normal cluster. In a normal cluster, Kerberos authentication is disabled, and users can use all functions provided by the cluster.</li> <li>• <b>KERBEROS</b>: security cluster. In a security cluster, Kerberos authentication is enabled, and common users cannot use the file management and job management functions of an MRS cluster or view cluster resource usage and the job records of Hadoop and Spark. To use more functions, the users must obtain the relevant permissions from the Manager administrator.</li> </ul> |
| manager_admin_password | Yes       | String | <p>Password of the MRS Manager administrator. The password must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• Must contain 8 to 26 characters.</li> <li>• Must contain at least four of the following: uppercase letters, lowercase letters, numbers, and special characters (!@#\$%^&amp;*_+[]{};:./?), but must not contain spaces.</li> <li>• Cannot be the username or the username spelled backwards.</li> </ul>                                                                                                                                                                                                              |

| Parameter             | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| login_mode            | Yes       | String | Node login mode. <ul style="list-style-type: none"> <li>• <b>PASSWORD</b>: password-based login. If this value is selected, <b>node_root_password</b> cannot be left blank.</li> <li>• <b>KEYPAIR</b>: key pair used for login. If this value is selected, <b>node_keypair_name</b> cannot be left blank.</li> </ul>                                                                                                                                                           |
| node_root_password    | No        | String | The password of user <b>root</b> for logging in to a cluster node. A password must meet the following complexity requirements: <ul style="list-style-type: none"> <li>• Must be 8 to 26 characters long.</li> <li>• Must contain at least four of the following: uppercase letters, lowercase letters, numbers, and special characters (!@\$%^&amp;*_+[]{};:./?), but must not contain spaces.</li> <li>• Cannot be the username or the username spelled backwards.</li> </ul> |
| node_keypair_name     | No        | String | The name of a key pair. You can use a key pair to log in to a cluster node.                                                                                                                                                                                                                                                                                                                                                                                                    |
| enterprise_project_id | No        | String | The enterprise project ID. When you create a cluster, associate the enterprise project ID with the cluster. The default value is <b>0</b> , indicating the <b>default</b> enterprise project. To obtain the enterprise project ID, see the <b>id</b> value in the <b>enterprise_project field data structure</b> table in "Querying the Enterprise Project List" in <i>Enterprise Management API Reference</i> .                                                               |

| Parameter              | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| eip_address            | No        | String | EIP bound to an MRS cluster, which can be used to access MRS Manager. The EIP must have been created and must be in the same region as the cluster.                                                                                                                                                                                                                                                                                                                                                                                                          |
| eip_id                 | No        | String | ID of the bound EIP. This parameter is mandatory when <b>eip_address</b> is configured. To obtain the EIP ID, log in to the VPC console, choose <b>Network &gt; Elastic IP and Bandwidth &gt; Elastic IP</b> , click the EIP to be bound, and obtain the ID in the <b>Basic Information</b> area.                                                                                                                                                                                                                                                            |
| mrs_ecs_default_agency | No        | String | Name of the agency bound to a cluster node by default. The value is fixed to <b>MRS_ECS_DEFAULT_AGENCY</b> . 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. The <b>MRS_ECS_DEFAULT_AGENCY</b> agency has the <b>OBS OperateAccess</b> permission of OBS and the <b>CES FullAccess</b> (for users who have enabled fine-grained policies), <b>CES Administrator</b> , and <b>KMS Administrator</b> permissions in the region where the cluster is located. |

| Parameter   | Mandatory | Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-----------|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| template_id | No        | String               | <p>The template used for node deployment when the cluster type is <b>CUSTOM</b>.</p> <ul style="list-style-type: none"> <li>• <b>mgmt_control_combined_v2</b>: template for jointly deploying management and controller nodes. The management and controller roles are co-deployed on the master node, and data instances are deployed in the same node group. This deployment model applies to scenarios where there are less than 100 nodes, reducing costs.</li> <li>• <b>mgmt_control_separated_v2</b>: The management and control roles are deployed on different master nodes, and data instances are deployed in the same node group. This deployment model applies to a cluster with 100 to 500 nodes and delivers better performance in high-concurrency load scenarios.</li> <li>• <b>mgmt_control_data_separated_v2</b>: The management and control roles are deployed on different master nodes, and data instances are deployed in different node groups. This deployment model applies to a cluster with more than 500 nodes. Components can be deployed separately, which can be used for a larger cluster scale.</li> </ul> |
| tags        | No        | Array of Tag objects | <p>The cluster tags. One cluster can have a maximum of 10 tags. The tag name (key) must be unique. For details about the parameters, see <a href="#">Table 6-27</a>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Parameter         | Mandatory | Type                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------|-----------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| log_collection    | No        | Integer                          | Whether to collect logs when cluster creation fails. The default value is <b>1</b> , indicating that OBS buckets will be created and only used to collect logs that record MRS cluster creation failures. Enumerated values: <ul style="list-style-type: none"> <li>• <b>0</b>: Do not collect logs.</li> <li>• <b>1</b>: Collect logs.</li> </ul>                                                                                                                                                    |
| node_groups       | Yes       | Array of NodeGroupV2 objects     | Information about the node groups that form the cluster. For details about the parameters, see <a href="#">Table 6-28</a> .                                                                                                                                                                                                                                                                                                                                                                           |
| bootstrap_scripts | No        | Array of BootstrapScript objects | The bootstrap action script. For details about the parameters, see <a href="#">Table 6-36</a> .                                                                                                                                                                                                                                                                                                                                                                                                       |
| log_uri           | No        | String                           | The OBS path to which cluster logs are dumped. After the log dump function is enabled, the read and write permissions on the OBS path are required to upload logs. Configure the default agency <b>MRS_ECS_DEFAULT_AGENCY</b> or customize an agency with the read and write permissions on the OBS path. For details, see <a href="#">Configuring a Storage-Compute Decoupled Cluster (Agency)</a> . This parameter is available only for cluster versions that support dumping cluster logs to OBS. |
| component_configs | No        | Array of ComponentConfig objects | The custom configuration of cluster components. This parameter applies only to cluster versions that support the feature of creating a cluster by customizing component configurations. For details about this parameter, see <a href="#">Table 6-37</a> .                                                                                                                                                                                                                                            |

| Parameter            | Mandatory | Type                        | Description                                                                                                   |
|----------------------|-----------|-----------------------------|---------------------------------------------------------------------------------------------------------------|
| delete_when_no_steps | No        | Boolean                     | Whether to automatically terminate the cluster after the job is complete. The default value is <b>false</b> . |
| steps                | Yes       | Array of StepConfig objects | The job list. For details about this parameter, see <a href="#">Table 6-39</a> .                              |

**Table 6-26** ClusterDataConnectorMap

| Parameter      | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                           |
|----------------|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| map_id         | No        | Integer | The data connection association ID                                                                                                                                                                                                                                                                    |
| connector_id   | No        | String  | The data connection ID                                                                                                                                                                                                                                                                                |
| component_name | No        | String  | The component name                                                                                                                                                                                                                                                                                    |
| role_type      | No        | String  | The component role type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>hive_metastore</b>: Hive Metastore role</li> <li>• <b>hive_data</b>: Hive role</li> <li>• <b>hbase_data</b>: HBase role</li> <li>• <b>ranger_data</b>: Ranger role</li> </ul>                        |
| source_type    | No        | String  | The data connection type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>LOCAL_DB</b>: local metadata</li> <li>• <b>RDS_POSTGRES</b>: RDS PostgreSQL database</li> <li>• <b>RDS_MYSQL</b>: RDS MySQL database</li> <li>• <b>gaussdb-mysql</b>: GaussDB(for MySQL)</li> </ul> |
| cluster_id     | No        | String  | The ID of the associated cluster                                                                                                                                                                                                                                                                      |
| status         | No        | Integer | The data connection status. The options are as follows: <ul style="list-style-type: none"> <li>• <b>0</b>: normal</li> <li>• <b>1</b>: in use</li> </ul>                                                                                                                                              |



**Table 6-27** Tag

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                          |
|-----------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| key       | Yes       | String | <p>The tag key</p> <ul style="list-style-type: none"> <li>It contains a maximum of 36 Unicode characters and cannot be an empty string.</li> <li>The tag key cannot start or end with spaces or contain non-printable ASCII characters (0-31) and special characters (=*&lt;&gt;\, /).</li> <li>The tag key of a resource must be unique.</li> </ul> |
| value     | Yes       | String | <p>The value</p> <ul style="list-style-type: none"> <li>The value contains a maximum of 43 Unicode characters and can be an empty string.</li> <li>The tag value cannot start or end with spaces or contain non-printable ASCII characters (0-31) and special characters (=*&lt;&gt;\, /).</li> </ul>                                                |

**Table 6-28** NodeGroupV2

| Parameter  | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| group_name | Yes       | String  | <p>The node group name. The value can contain a maximum of 64 characters, including letters, numbers, and underscores (_). The rules for configuring node groups are as follows:</p> <ul style="list-style-type: none"> <li>• <b>master_node_default_group</b>: master node group, which must be included in all cluster types.</li> <li>• <b>core_node_analysis_group</b>: analysis core node group, which must be included in both analysis and hybrid clusters.</li> <li>• <b>core_node_streaming_group</b>: streaming core node group, which must be included in both streaming and hybrid clusters.</li> <li>• <b>task_node_analysis_group</b>: analysis task node group, which can be selected for analysis clusters and hybrid clusters as needed.</li> <li>• <b>task_node_streaming_group</b>: streaming task node group, which can be selected for streaming clusters and hybrid clusters as needed.</li> <li>• <b>node_group{x}</b>: node group of a custom cluster. A maximum of nine such node groups can be added for a custom cluster.</li> </ul> |
| node_num   | Yes       | Integer | <p>The number of nodes. The value ranges from 0 to 500. The maximum number of core and task nodes is 500.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

| Parameter           | Mandatory | Type                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------|-----------|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_size           | Yes       | String                                   | The instance specifications of a node, for example, <b>c3.4xlarge.2.linux.bigdata</b> . For details about instance specifications, see <a href="#">ECS Specifications Used by MRS</a> and <a href="#">BMS Specifications Used by MRS</a> . You are advised to obtain the specifications supported by the corresponding version in the corresponding region from the cluster creation page on the MRS console. |
| root_volume         | No        | Volume object                            | The system disk information of the node. This parameter is optional for some VMs or the system disk of the BMS and mandatory in other cases. For details about this parameter, see <a href="#">Table 6-29</a> .                                                                                                                                                                                               |
| data_volume         | No        | Volume object                            | The data disk information. This parameter is mandatory when <b>data_volume_count</b> is not <b>0</b> . For details about this parameter, see <a href="#">Table 6-29</a> .                                                                                                                                                                                                                                     |
| data_volume_count   | No        | Integer                                  | The number of data disks on a node. The value ranges from 0 to 20.                                                                                                                                                                                                                                                                                                                                            |
| charge_info         | No        | ChargeInfo object                        | The billing type of a node group. The billing types of master and core node groups are the same as those of the cluster. The billing type of the task node group can be different. For details about this parameter, see <a href="#">Table 6-30</a> .                                                                                                                                                         |
| auto_scaling_policy | No        | <a href="#">AutoScalingPolicy</a> object | The auto scaling rule information. For details about this parameter, see <a href="#">Table 6-31</a> .                                                                                                                                                                                                                                                                                                         |

| Parameter      | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------|-----------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| assigned_roles | No        | Array of strings | <p>This parameter is mandatory when the cluster type is <b>CUSTOM</b>. You can specify the roles deployed in a node group. This parameter is a string array. Each string represents a role expression. Role expression definition:</p> <ul style="list-style-type: none"> <li>• If the role is deployed on all nodes in a node group, set this parameter to <i>{role name}</i>, for example, <b>DataNode</b>.</li> <li>• If the role is deployed on a specified subscript node in the node group: <i>{role name}:{index1},{index2}..., {indexN}</i>, for example, <b>NameNode:1,2</b>. The subscript starts from 1.</li> <li>• Some roles support multi-instance deployment (that is, multiple instances of the same role are deployed on a node): <i>{role name} [{instance count}]</i>, for example, <b>EsNode[9]</b>. For details about the available roles, see <a href="#">Roles and components supported by MRS</a>.</li> </ul> |

Table 6-29 Volume

| Parameter | Mandatory | Type    | Description                                                                                                                                                                                                                                                          |
|-----------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| type      | Yes       | String  | <p>The disk type. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O disk</li> <li>• <b>SAS</b>: high I/O disk</li> <li>• <b>SSD</b>: ultra-high I/O disk</li> <li>• <b>GPSSD</b>: general-purpose SSD disk</li> </ul> |
| size      | Yes       | Integer | The data disk size, in GB. The value ranges from 10 to 32768.                                                                                                                                                                                                        |

**Table 6-30** ChargeInfo

| Parameter   | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                      |
|-------------|-----------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| charge_mode | Yes       | String  | <p>The billing mode. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>prePaid</b>: the yearly/monthly billing mode. This mode is now supported for the API used to create a cluster, but is not supported for the API used to create a cluster and submit a job.</li> <li>• <b>postPaid</b>: the pay-per-use billing mode.</li> </ul> |
| period_type | No        | String  | <p>Subscription period. The value can be:</p> <ul style="list-style-type: none"> <li>• <b>month</b>: The cluster is billed by month.</li> <li>• <b>year</b>: The cluster is billed by year.</li> <li>• <b>day</b>: The cluster is billed on a pay-per-use basis.</li> </ul>                                                                                      |
| period_num  | No        | Integer | <p>Number of periods. This parameter is valid and mandatory only when <b>charge_mode</b> is set to <b>prePaid</b>.</p> <ul style="list-style-type: none"> <li>• If <b>period_type</b> is set to <b>month</b>, the value ranges from 1 to 9.</li> <li>• If <b>period_type</b> is set to <b>year</b>, the value ranges from 1 to 3.</li> </ul>                     |

| Parameter   | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| is_auto_pay | No        | Boolean | <p>Whether the order will be automatically paid. This parameter is available for yearly/monthly mode. By default, the automatic payment is disabled. Available values are as follows:</p> <ul style="list-style-type: none"> <li>• <b>true</b>: The system automatically selects available discounts and coupons, and then pays for the order with the account balances. If the automatic payment fails, an order in <b>Pending payment</b> state is generated waiting for manual payment.</li> <li>• <b>false</b>: The user needs to pay for the bill after using available discounts and coupons.</li> </ul> |

**Table 6-31** AutoScalingPolicy

| Parameter           | Mandatory | Type                           | Description                                                                                                                                                                                                                                                                  |
|---------------------|-----------|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| auto_scaling_enable | Yes       | Boolean                        | Whether to enable the auto scaling policy.                                                                                                                                                                                                                                   |
| min_capacity        | Yes       | Integer                        | The minimum number of nodes reserved in the node group. Value range: [0, 500]                                                                                                                                                                                                |
| max_capacity        | Yes       | Integer                        | The maximum number of nodes in the node group. Value range: [0, 500]                                                                                                                                                                                                         |
| resources_plans     | No        | Array of ResourcesPlan objects | The resource plan list. If this parameter is left blank, the resource plan is disabled. When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>rules</b> must be configured. For details about this parameter, see <a href="#">Table 6-32</a> . |

| Parameter    | Mandatory | Type                         | Description                                                                                                                                                                                                                    |
|--------------|-----------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rules        | No        | Array of Rule objects        | The list of auto scaling rules. When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>resources_plans</b> must be configured. For details about this parameter, see <a href="#">Table 6-33</a> . |
| exec_scripts | No        | Array of ScaleScript objects | The list of custom scaling automation scripts. If this parameter is left blank, the automation script is disabled. For details about this parameter, see <a href="#">Table 6-35</a> .                                          |

**Table 6-32** ResourcesPlan

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                             |
|--------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| period_type  | Yes       | String  | The cycle type of a resource plan. Currently, only the following cycle type is supported:<br><b>daily</b>                                                                                               |
| start_time   | Yes       | String  | The start time of a resource plan. The value is in the format of <b>hour:minute</b> , indicating that the time ranges from 00:00 to 23:59.                                                              |
| end_time     | Yes       | String  | The end time of a resource plan. The value is in the same format as that of <b>start_time</b> . The interval between <b>end_time</b> and <b>start_time</b> must be greater than or equal to 30 minutes. |
| min_capacity | Yes       | Integer | The minimum number of reserved nodes in a node group in a resource plan. Value range: [0, 500]                                                                                                          |
| max_capacity | Yes       | Integer | The maximum number of reserved nodes in a node group in a resource plan. Value range: [0, 500]                                                                                                          |

| Parameter      | Mandatory | Type             | Description                                                                                                                                                                                                                                  |
|----------------|-----------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| effective_days | No        | Array of strings | The effective date of a resource plan. If this parameter is left blank, it indicates that the resource plan takes effect every day. The options are as follows:<br><b>MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</b> |

**Table 6-33** Rule

| Parameter          | Mandatory | Type           | Description                                                                                                                                                                                                      |
|--------------------|-----------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name               | Yes       | String         | The name of an auto scaling rule. The name can contain only 1 to 64 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed. Rule names must be unique in a node group.                  |
| description        | No        | String         | The description about an auto scaling rule. It contains a maximum of 1,024 characters.                                                                                                                           |
| adjustment_type    | Yes       | String         | The adjustment type of an auto scaling rule. The options are as follows: <ul style="list-style-type: none"> <li>● <b>scale_out</b>: cluster scale-out</li> <li>● <b>scale_in</b>: cluster scale-in</li> </ul>    |
| cool_down_minutes  | Yes       | Integer        | The cluster cooling time after an auto scaling rule is triggered, in minutes, during which period no auto scaling operation is performed. The value ranges from 0 to 10080. One week is equal to 10,080 minutes. |
| scaling_adjustment | Yes       | Integer        | The number of cluster nodes that can be adjusted at a time. Value range: [1, 100]                                                                                                                                |
| trigger            | Yes       | Trigger object | The condition for triggering a rule. For details about this parameter, see <a href="#">Table 6-34</a> .                                                                                                          |



**Table 6-34** Trigger

| Parameter           | Mandatory | Type    | Description                                                                                                                                                                                                                                                                |
|---------------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| metric_name         | Yes       | String  | The metric name. This triggering condition makes a judgment according to the value of the metric. A metric name contains a maximum of 64 characters.                                                                                                                       |
| metric_value        | Yes       | String  | The metric threshold to trigger a rule.<br>The value must be an integer or a number with two decimal places.                                                                                                                                                               |
| comparison_operator | No        | String  | The metric judgment logic operator. The options are as follows: <ul style="list-style-type: none"> <li>• <b>LT</b>: less than</li> <li>• <b>GT</b>: greater than</li> <li>• <b>LTOE</b>: less than or equal to</li> <li>• <b>GTOE</b>: greater than or equal to</li> </ul> |
| evaluation_periods  | Yes       | Integer | The number of consecutive five-minute periods, during which a metric threshold is reached. The value ranges from 1 to 288.                                                                                                                                                 |

**Table 6-35** ScaleScript

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                               |
|-----------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | Yes       | String | The name of a custom automation script. Names must be unique in a cluster. The value can contain only numbers, letters, spaces, hyphens (-), and underscores (_) and cannot start with a space. The value can contain 1 to 64 characters. |

| Parameter  | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------|-----------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri        | Yes       | String           | <p>The path of a custom automation script. Set this parameter to an OBS bucket path or a local VM path.</p> <ul style="list-style-type: none"> <li>• OBS bucket path: Enter a script path, for example, <b>obs://XXX/scale.sh</b>.</li> <li>• Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| parameters | No        | String           | <p>Parameters of a custom automation script. Multiple parameters are separated by space. The following predefined system parameters can be transferred:</p> <ul style="list-style-type: none"> <li>• <i>`\${mrs_scale_node_num}`</i>: The number of nodes to be added or removed</li> <li>• <i>`\${mrs_scale_type}`</i>: The scaling type. The value can be <b>scale_out</b> or <b>scale_in</b>.</li> <li>• <i>`\${mrs_scale_node_hostname s}`</i>: Host names of the nodes to be added or removed</li> <li>• <i>`\${mrs_scale_node_ips}`</i>: IP addresses of the nodes to be added or removed</li> <li>• <i>`\${mrs_scale_rule_name}`</i>: Name of the rule that triggers auto scaling Other user-defined parameters are used in the same way as those of common shell scripts. Parameters are separated by space.</li> </ul> |
| nodes      | Yes       | Array of strings | <p>The name of the node group where the custom automation script is executed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| Parameter     | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|-----------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active_master | No        | Boolean | Whether the custom automation script runs only on the active master node. The default value is <b>false</b> , indicating that the custom automation script can run on all master nodes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| fail_action   | Yes       | String  | Whether to continue executing subsequent scripts and creating a cluster after the custom automation script fails to be executed. Notes: <ul style="list-style-type: none"> <li>You are advised to set this parameter to <b>continue</b> in the commissioning phase so the cluster can continue to be installed and started no matter whether the custom automation script is executed successfully.</li> <li>The scale-in operation cannot be undone. <b>fail_action</b> must be set to <b>continue</b> for the scripts that are executed after scale-in. Enumerated values: <ul style="list-style-type: none"> <li><b>continue</b>: Continue to execute subsequent scripts.</li> <li><b>errorout</b>: Stop the action.</li> </ul> </li> </ul> |
| action_stage  | Yes       | String  | The time when a script is executed. Enumerated values: <ul style="list-style-type: none"> <li><b>before_scale_out</b>: before scale-out</li> <li><b>before_scale_in</b>: before scale-in</li> <li><b>after_scale_out</b>: after scale-out</li> <li><b>after_scale_in</b>: after scale-in</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Table 6-36** BootstrapScript

| Parameter     | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------|-----------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name          | Yes       | String           | The name of a bootstrap action script, which must be unique in a cluster. The value can contain only numbers, letters, spaces, hyphens (-), and underscores (_) and cannot start with a space. The value can contain 1 to 64 characters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| uri           | Yes       | String           | The path of a bootstrap action script. Set this parameter to an OBS bucket path or a local VM path. OBS bucket path: Enter a script path, For example, enter the path of the public sample script provided by MRS. Example: <b>obs://bootstrap/presto/presto-install.sh</b> . If <b>dualroles</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>dualroles</b> . If <b>worker</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>worker</b> . Based on the Presto usage habit, you are advised to install <b>dualroles</b> on the active master nodes and <b>worker</b> on the core nodes. Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b> . |
| parameters    | No        | String           | The bootstrap action script parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| nodes         | Yes       | Array of strings | The name of the node group where the bootstrap action script is executed                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| active_master | No        | Boolean          | Whether the bootstrap action script runs only on active master nodes. The default value is <b>false</b> , indicating that the bootstrap action script can run on all master nodes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

| Parameter              | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fail_action            | Yes       | String  | Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to execute. The default value is <b>errorout</b> , indicating that the action is stopped. Note: You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful. Enumerated values: <ul style="list-style-type: none"> <li>• <b>continue</b>: Continue to execute subsequent scripts.</li> <li>• <b>errorout</b>: Stop the action.</li> </ul> |
| before_component_start | No        | Boolean | The time when the bootstrap action script is executed. Currently, the following two options are available: <b>Before component start</b> and <b>After component start</b> . The default value is <b>false</b> , indicating that the bootstrap action script is executed after the component is started.                                                                                                                                                                                                                                                                                             |
| start_time             | No        | Long    | The execution time of one bootstrap action script.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| state                  | No        | String  | The running status of one bootstrap action script. The options are as follows: <ul style="list-style-type: none"> <li>• <b>PENDING</b></li> <li>• <b>IN_PROGRESS</b></li> <li>• <b>SUCCESS</b></li> <li>• <b>FAILURE</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                 |

| Parameter     | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|-----------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| action_stages | No        | Array of strings | Select the time when the bootstrap action script is executed. <ul style="list-style-type: none"> <li>• <b>BEFORE_COMPONENT_FIRST_START</b>: before initial component starts</li> <li>• <b>AFTER_COMPONENT_FIRST_START</b>: after initial component starts</li> <li>• <b>BEFORE_SCALE_IN</b>: before scale-in</li> <li>• <b>AFTER_SCALE_IN</b>: after scale-in</li> <li>• <b>BEFORE_SCALE_OUT</b>: before scale-out</li> <li>• <b>AFTER_SCALE_OUT</b>: after scale-out</li> </ul> |

**Table 6-37** ComponentConfig

| Parameter      | Mandatory | Type                    | Description                                                                                               |
|----------------|-----------|-------------------------|-----------------------------------------------------------------------------------------------------------|
| component_name | Yes       | String                  | The component name                                                                                        |
| configs        | No        | Array of Config objects | The component configuration item list. For details about this parameter, see <a href="#">Table 6-38</a> . |

**Table 6-38** Config

| Parameter        | Mandatory | Type   | Description                                                                                                           |
|------------------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------|
| key              | Yes       | String | The configuration name. Only the configuration names displayed on the MRS component configuration page are supported. |
| value            | Yes       | String | The configuration value                                                                                               |
| config_file_name | Yes       | String | The configuration file name. Only the file names displayed on the MRS component configuration page are supported.     |

**Table 6-39** StepConfig

| Parameter     | Mandatory | Type                | Description                                                                           |
|---------------|-----------|---------------------|---------------------------------------------------------------------------------------|
| job_execution | Yes       | JobExecution object | The job parameter. For details about this parameter, see <a href="#">Table 6-40</a> . |

**Table 6-40** JobExecution

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_type  | Yes       | String | <p>The job type. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>MapReduce</b></li> <li>• <b>SparkSubmit</b></li> <li>• <b>SparkPython</b>: Example request of a SparkPython job (Jobs of this type will be converted to SparkSubmit jobs for submission. The job type is displayed as SparkSubmit on the MRS console. Select SparkSubmit when you call an API to query the job list.)</li> <li>• <b>HiveScript</b></li> <li>• <b>HiveSql</b></li> <li>• <b>DistCp</b>: imports and exports data.</li> <li>• <b>SparkScript</b></li> <li>• <b>SparkSql</b></li> <li>• <b>Flink</b></li> </ul> |
| job_name  | Yes       | String | The job name. The value can contain 1 to 64 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed. Identical job names are allowed but not recommended.                                                                                                                                                                                                                                                                                                                                                                                                                                         |

| Parameter  | Mandatory | Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------|-----------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| arguments  | No        | Array of strings   | <p>The 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. The value can contain a maximum of 150,000 characters. Special characters (; &amp;&gt;'&lt;\$!'") are not allowed. This parameter can be left blank. Notes:</p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• If you need to access files stored in OBS via the path starting with <b>obs://</b> when submitting a HiveScript or HiveSQL job, search for the <b>core.site.customized.configs</b> parameter on the Hive service configuration page, add the endpoint configuration item (<b>fs.obs.endpoint</b>) of OBS, and set the value to the endpoint corresponding to OBS. For details, see <a href="#">Endpoints</a>.</li> </ul> |
| properties | No        | Map<String,String> | <p>The program system parameter. The value can contain a maximum of 2,048 characters. Special characters (; &amp;&gt;'&lt;\$!\) are not allowed. This parameter can be left blank.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Response Parameters

Status code: 200



**Table 6-41** Response body parameter

| Parameter  | Type   | Description                                                                   |
|------------|--------|-------------------------------------------------------------------------------|
| cluster_id | String | The cluster ID, which is returned by the system after the cluster is created. |

## Example Request

Create an MRS 3.2.0-LTS.1 cluster where the custom management nodes and control nodes are the same nodes and submit a HiveScript job.

POST /v2/{project\_id}/run-job-flow

```
{
  "cluster_version": "MRS 3.2.0-LTS.1",
  "cluster_name": "mrs_heshe_dm",
  "cluster_type": "CUSTOM",
  "charge_info": {
    "charge_mode": "postPaid"
  },
  "region": "",
  "availability_zone": "",
  "vpc_name": "vpc-37cd",
  "subnet_id": "1f8c5ca6-1f66-4096-bb00-baf175954f6e",
  "subnet_name": "subnet",
  "components": "Hadoop,Spark2x,HBase,Hive,Hue,Loader,Kafka,Storm,Flume,Flink,Oozie,Ranger,Tez",
  "safe_mode": "KERBEROS",
  "manager_admin_password": "your password",
  "login_mode": "PASSWORD",
  "node_root_password": "your password",
  "mrs_ecs_default_agency": "MRS_ECS_DEFAULT_AGENCY",
  "template_id": "mgmt_control_combined_v2",
  "log_collection": 1,
  "tags": [ {
    "key": "tag1",
    "value": "111"
  }, {
    "key": "tag2",
    "value": "222"
  } ],
  "node_groups": [ {
    "group_name": "master_node_default_group",
    "node_num": 3,
    "node_size": "Sit3.4xlarge.4.linux.bigdata",
    "root_volume": {
      "type": "SAS",
      "size": 480
    },
    "data_volume": {
      "type": "SAS",
      "size": 600
    },
    "data_volume_count": 1,
    "assigned_roles": [ "OMSServer:1,2", "SlapdServer:1,2", "KerberosServer:1,2", "KerberosAdmin:1,2",
"quorumpeer:1,2,3", "NameNode:2,3", "Zkfc:2,3", "JournalNode:1,2,3", "ResourceManager:2,3",
"JobHistoryServer:2,3", "DBServer:1,3", "Hue:1,3", "LoaderServer:1,3", "MetaStore:1,2,3", "WebHCat:1,2,3",
"HiveServer:1,2,3", "HMaster:2,3", "MonitorServer:1,2", "Nimbus:1,2", "UI:1,2", "JDBCServer2x:1,2,3",
"JobHistory2x:2,3", "SparkResource2x:1,2,3", "oozie:2,3", "LoadBalancer:2,3", "TezUI:1,3", "TimelineServer:3",
"RangerAdmin:1,2", "UserSync:2", "TagSync:2", "KerberosClient", "SlapdClient", "meta", "HSConsole:2,3",
"FlinkResource:1,2,3", "DataNode:1,2,3", "NodeManager:1,2,3", "IndexServer2x:1,2", "ThriftServer:1,2,3",
"RegionServer:1,2,3", "ThriftServer1:1,2,3", "RESTServer:1,2,3", "Broker:1,2,3", "Supervisor:1,2,3",
"Logviewer:1,2,3", "Flume:1,2,3", "HSBroker:1,2,3" ]
  }, {
    "group_name": "node_group_1",
```

```
"node_num" : 3,
"node_size" : "Sit3.4xlarge.4.linux.bigdata",
"root_volume" : {
  "type" : "SAS",
  "size" : 480
},
"data_volume" : {
  "type" : "SAS",
  "size" : 600
},
"data_volume_count" : 1,
"assigned_roles" : [ "DataNode", "NodeManager", "RegionServer", "Flume:1", "Broker", "Supervisor",
"Logviewer", "HBaseIndexer", "KerberosClient", "SlapdClient", "meta", "HSBroker:1,2", "ThriftServer",
"ThriftServer1", "RESTServer", "FlinkResource" ]
}, {
  "group_name" : "node_group_2",
  "node_num" : 1,
  "node_size" : "Sit3.4xlarge.4.linux.bigdata",
  "root_volume" : {
    "type" : "SAS",
    "size" : 480
  },
  "data_volume" : {
    "type" : "SAS",
    "size" : 600
  },
  "data_volume_count" : 1,
  "assigned_roles" : [ "NodeManager", "KerberosClient", "SlapdClient", "meta", "FlinkResource" ]
}],
"log_uri" : "obs://bucketTest/logs",
"delete_when_no_steps" : true,
"steps" : [ {
  "job_execution" : {
    "job_name" : "import_file",
    "job_type" : "DistCp",
    "arguments" : [ "obs://test/test.sql", "/user/hive/input" ]
  }
}, {
  "job_execution" : {
    "job_name" : "hive_test",
    "job_type" : "HiveScript",
    "arguments" : [ "obs://test/hive/sql/HiveScript.sql" ]
  }
}
]
}
```

## Example Response

**Status code: 200**

Example successful response

```
{
  "cluster_id" : "da1592c2-bb7e-468d-9ac9-83246e95447a"
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.1.4 Scaling Out a Cluster

### Function

This API is used to scale out an MRS cluster.

### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/expand

**Table 6-42** URI parameters

| Parameter  | Mandatory | Type   | Description                                                   |
|------------|-----------|--------|---------------------------------------------------------------|
| project_id | Yes       | String | Project ID. The value can contain a maximum of 64 characters. |
| cluster_id | Yes       | String | Cluster ID. The value can contain a maximum of 64 characters. |

### Request Parameters

**Table 6-43** Request body parameters

| Parameter              | Mandatory | Type    | Description                                                                                                                                                                                             |
|------------------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_group_name        | Yes       | String  | Node group name. <ul style="list-style-type: none"> <li>• Minimum: 1 character</li> <li>• Maximum: 64 characters</li> </ul>                                                                             |
| count                  | Yes       | Integer | Number of nodes to added                                                                                                                                                                                |
| skip_bootstrap_scripts | No        | Boolean | Whether to skip the specified bootstrap action during cluster creation on the new node during scale-out. The default value is <b>true</b> , indicating that the bootstrap action is skipped by default. |

| Parameter           | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                |
|---------------------|-----------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scale_without_start | No        | Boolean | Whether the components remain stopped on the added nodes after cluster scale-out. The default value is <b>false</b> , indicating that the components will be started. <ul style="list-style-type: none"> <li>• <b>true</b>: Do not start components after scale-out.</li> <li>• <b>false</b>: Start components after scale-out.</li> </ul> |

## Response Parameters

Status code: 200

Table 6-44 Response body parameter

| Parameter | Type   | Description                                                                                                                                                                                                                                                                   |
|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| result    | String | Result of the request. Value <b>succeeded</b> indicates that the operation is successful, and value <b>failed</b> indicates that the operation fails. Value <b>success</b> is returned if the request for scaling out a non-yearly/monthly node group is successfully issued. |
| order_id  | String | Order ID. When you scale out a yearly/monthly node group, the order ID generated during the scale-out is returned. You need to pay for the order to trigger the scale-out.                                                                                                    |

## Example Request

- Add a node to the node group **node\_group\_1**. Skip the bootstrap action and start the components by default.

```
/v2/ff8080828997cb24018a1b2db3440b80/clusters/f7f45c04-4303-411c-9b71-d2cb730dd162/expand
```

```
{
  "node_group_name": "node_group_1",
  "count": "1"
}
```

- Add a node to the node group **node\_group\_1**. Do not skip the bootstrap action or start the components.

```
/v2/ff8080828997cb24018a1b2db3440b80/clusters/f7f45c04-4303-411c-9b71-d2cb730dd162/expand
```

```
{
  "node_group_name": "node_group_1",
```

```
"count" : "1",
"skip_bootstrap_scripts" : false,
"scale_without_start" : true
}
```

## Example Response

### Status code: 200

The response body varies depending on the billing mode of node groups. The order ID is returned for you to settle payment for scaling out a yearly/monthly node group.

- A pay-per-use node group is scaled out successfully.
- A yearly/monthly node group is scaled out successfully. The ID of the order you need to pay is CS231XXXXXXXXXX.

```
{
  "result" : "succeeded"
}
```

```
{
  "order_id" : "CS231XXXXXXXXXX"
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.1.5 Scaling In a Cluster

### Function

This API is used to scale in an MRS cluster.

### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/shrink

**Table 6-45** URI parameters

| Parameter  | Mandatory | Type   | Description                                                   |
|------------|-----------|--------|---------------------------------------------------------------|
| project_id | Yes       | String | Project ID. The value can contain a maximum of 64 characters. |
| cluster_id | Yes       | String | Cluster ID. The value can contain a maximum of 64 characters. |

## Request Parameters

**Table 6-46** Request body parameters

| Parameter       | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                 |
|-----------------|-----------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_group_name | Yes       | String           | Node group name <ul style="list-style-type: none"> <li>• Minimum length: 1 character</li> <li>• Maximum length: 64 characters</li> </ul>                                                                                                                                                                    |
| count           | No        | Integer          | Number of nodes to be removed. If a specified node is to be removed, this parameter can be left blank.                                                                                                                                                                                                      |
| resource_ids    | No        | Array of strings | Resource IDs of nodes to be removed during node scale-in. If this parameter is left blank, nodes are automatically removed based on system rules. Only abnormal ECS nodes can be removed. The specified nodes are forcibly removed. You can obtain the value of <b>resource_id</b> by calling the host API. |

## Response Parameters

Status code: 200

**Table 6-47** Response body parameter

| Parameter | Type   | Description                                                                                                                                      |
|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| result    | String | Operation result. Value <b>succeeded</b> indicates that the operation is successful, and value <b>failed</b> indicates that the operation fails. |

## Example Request

- Remove a node from node group **node\_group\_1**.  
/v2/ff8080828997cb24018a1b2db3440b80/clusters/f7f45c04-4303-411c-9b71-d2cb730dd162/shrink

```

{
  "node_group_name": "node_group_1",
  "count": 1
}

```

- Forcibly remove a node whose resource ID is **678050cd-ba1d-4550-942d-f2e396b1c6fb** from node group **node\_group\_1**.

/v2/ff8080828997cb24018a1b2db3440b80/clusters/f7f45c04-4303-411c-9b71-d2cb730dd162/shrink

```
{
  "node_group_name" : "node_group_1",
  "resource_ids" : [ "678050cd-ba1d-4550-942d-f2e396b1c6fb" ]
}
```

## Example Response

Status code: 200

The scale-in is successful.

```
{
  "result" : "succeeded"
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.1.6 Adding Components to a Cluster

### Function

This API is used to add components to a cluster.

### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/components

**Table 6-48** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                              |
|------------|-----------|--------|----------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |
| cluster_id | Yes       | String | Cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> . |

## Request Parameters

**Table 6-49** Request body parameters

| Parameter               | Mandatory | Type                                  | Description                                                                                        |
|-------------------------|-----------|---------------------------------------|----------------------------------------------------------------------------------------------------|
| components_install_mode | Yes       | Array of ComponentInstallMode objects | Component installation details. For details about the parameters, see <a href="#">Table 6-50</a> . |

**Table 6-50** ComponentInstallMode

| Parameter               | Mandatory | Type                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|-----------|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| component               | Yes       | String                             | Component name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| node_groups             | Yes       | Array of AssignedNodeGroup objects | Role deployment information of the component. For details, see <a href="#">Table 6-51</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| component_user_password | No        | String                             | User password of the component. The password is used for machine-machine account to connect to the ClickHouse component. The password: <ul style="list-style-type: none"> <li>• Must contain 8 to 26 characters.</li> <li>• Cannot be the username or the username spelled backwards.</li> <li>• Must contain every type of the following: <ul style="list-style-type: none"> <li>• Lowercase letters</li> <li>• Uppercase letters</li> <li>• Numbers</li> <li>• Special characters (!@\$%^_+=+[]{};./?)</li> </ul> </li> </ul> |



| Parameter                  | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|-----------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| component_default_password | No        | String | <p>Default user password of the component. The password is used for human-machine account to connect to the ClickHouse component. The password:</p> <ul style="list-style-type: none"> <li>• Must contain 8 to 26 characters.</li> <li>• Cannot be the username or the username spelled backwards.</li> <li>• Must contain every type of the following: <ul style="list-style-type: none"> <li>• Lowercase letters</li> <li>• Uppercase letters</li> <li>• Numbers</li> <li>• Special characters (!@\$%^-_=+[]{};./?)</li> </ul> </li> </ul> |

**Table 6-51** AssignedNodeGroup

| Parameter | Mandatory | Type   | Description     |
|-----------|-----------|--------|-----------------|
| name      | Yes       | String | Node group name |

| Parameter      | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|-----------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| assigned_roles | Yes       | Array of strings | <p>Role deployment information. You can specify the roles deployed in a node group. This parameter is a string array. Each string represents a role expression. Role expression definition:</p> <ul style="list-style-type: none"> <li>• If the role is deployed on all nodes in a node group, set this parameter to <i>{role name}</i>, for example, <b>DataNode</b>.</li> <li>• If the role is deployed on a specified subscript node in the node group: <i>{role name}:{index1},{index2}..., {indexN}</i>, for example, <b>NameNode:1,2</b>. The subscript starts from 1.</li> <li>• For some roles, multiple instances of the same role can be deployed on the same node with <i>{role name} [{instance count}]</i>, for example, <b>EsNode[9]</b>. You do not need to specify the role location. By default, multiple instances are deployed on all nodes in a node group. For details about the available roles, see the <a href="#">Roles and Components Supported by MRS</a>.</li> </ul> |

## Response Parameters

Status code: 200

Table 6-52 Response body parameter

| Parameter | Type   | Description                                                                                                                                      |
|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| result    | String | Operation result. Value <b>succeeded</b> indicates that the operation is successful, and value <b>failed</b> indicates that the operation fails. |

## Example Request

- Add the ClickHouse component to a cluster in normal mode.  
v2/f77c10d14a544393a24e5f0bf53202b6/clusters/ff879d3a-e5d5-4485-a9b6-c673b52673fa/components  

```
{
  "components_install_mode" : [ {
    "component" : "ClickHouse",
    "node_groups" : [ {
      "name" : "master_node_default_group",
      "assigned_roles" : [ "ClickHouseServer:1,2" ]
    }, {
      "name" : "node_group_1",
      "assigned_roles" : [ "ClickHouseServer", "ClickHouseBalancer" ]
    } ],
    "component_user_password" : "*****",
    "component_default_password" : "*****"
  } ]
}
```
- Add the HBase component.  
v2/f77c10d14a544393a24e5f0bf53202b6/clusters/ff879d3a-e5d5-4485-a9b6-c673b52673fa/components  

```
{
  "components_install_mode" : [ {
    "component" : "HBase",
    "node_groups" : [ {
      "name" : "master_node_default_group",
      "assigned_roles" : [ "RegionServer", "HMaster" ]
    } ]
  } ]
}
```

## Example Response

**Status code: 200**

Processing result of a request:

```
{
  "result" : "succeeded"
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.1.7 Querying the Cluster Node List

### Function

This API is used to query the cluster node list.

### URI

GET /v2/{project\_id}/clusters/{cluster\_id}/nodes

**Table 6-53** URI parameters

| Parameter  | Mandatory | Type   | Description |
|------------|-----------|--------|-------------|
| project_id | Yes       | String | Project ID  |
| cluster_id | Yes       | String | Cluster ID  |

**Table 6-54** Query parameters

| Parameter         | Mandatory | Type    | Description                                                                                                         |
|-------------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------|
| node_group        | No        | String  | Name of the node groups to be queried.                                                                              |
| limit             | No        | Integer | Number of records displayed on each page in the returned result. Default value: 10                                  |
| offset            | No        | Integer | Offset from which the job list starts to be queried. The default value is <b>1</b> .                                |
| node_name         | No        | String  | Node name. Fuzzy search is supported.                                                                               |
| sort_key          | No        | String  | Sorting key. You can set node names as the key.                                                                     |
| sort_dir          | No        | String  | Sorting order. <b>desc</b> indicates descending order, and <b>asc</b> indicates ascending order.                    |
| query_node_detail | No        | Boolean | Whether to query node details. If this parameter is set to <b>true</b> , the interface performance may be affected. |
| query_ecs_detail  | No        | Boolean | Whether to query ECS details. ECS APIs may be called.                                                               |
| internal_ip       | No        | String  | Internal IP address.                                                                                                |

## Request Parameters

None

## Response Parameters

Status code: 200

**Table 6-55** Response body parameter

| Parameter  | Type                         | Description                                              |
|------------|------------------------------|----------------------------------------------------------|
| nodes      | Array of ClusterNode objects | Node list. For details, see <a href="#">Table 6-56</a> . |
| node_total | Integer                      | Number of nodes.                                         |

**Table 6-56** ClusterNode

| Parameter       | Type                           | Description                                                                                                                           |
|-----------------|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| node_name       | String                         | Node name, which is the same as that on Manager.                                                                                      |
| resource_id     | String                         | Specifies the resource ID. Unique identity of a node. For yearly/monthly nodes, the ID can be used to query bills.                    |
| node_group_name | String                         | Node group name.                                                                                                                      |
| node_type       | String                         | Node type. Available options include Task, Core, and Master.                                                                          |
| billing_type    | String                         | <b>on-period:</b> yearly/monthly; <b>on-quantity:</b> pay-per-use.                                                                    |
| deployment_type | String                         | Deployment type. The Server hosts are supported.                                                                                      |
| server_info     | ServerInfo object              | If the deployment type is Server, this parameter cannot be empty. For details, see <a href="#">Table 6-57</a> .                       |
| tags            | Array of Tag objects           | Node tag. For details, see <a href="#">Table 6-59</a> .                                                                               |
| node_detail     | NodeDetail object              | Node monitoring information. This parameter is returned only after IAM synchronization. For details, see <a href="#">Table 6-60</a> . |
| node_status     | String                         | Node status. Status of operations on the corresponding page.                                                                          |
| component_infos | Array of ComponentInfo objects | Component instance information array. For details, see <a href="#">Table 6-61</a> .                                                   |

**Table 6-57** ServerInfo

| Parameter    | Type                        | Description                                                |
|--------------|-----------------------------|------------------------------------------------------------|
| server_id    | String                      | Server ID.                                                 |
| server_name  | String                      | Server name.                                               |
| server_type  | String                      | Server type. Available options are ECS or BMS.             |
| data_volumes | Array of VolumeInfo objects | Data disk. For details, see <a href="#">Table 6-58</a> .   |
| root_volume  | VolumeInfo object           | System disk. For details, see <a href="#">Table 6-58</a> . |
| cpu_type     | String                      | CPU type. Available options are x86 and Arm.               |
| cpu          | String                      | Number of vCPUs.                                           |
| mem          | String                      | Memory size.                                               |
| internal_ip  | String                      | Internal IP address.                                       |

**Table 6-58** VolumeInfo

| Parameter | Type    | Description                                                                                                                                                                                                                   |
|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| type      | String  | Disk type. <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O disk</li> <li>• <b>SAS</b>: high I/O disk</li> <li>• <b>SSD</b>: ultra-high I/O disk</li> <li>• <b>GPSSD</b>: general-purpose SSD disk</li> </ul> |
| size      | Integer | Storage. Unit: GB.                                                                                                                                                                                                            |
| count     | Integer | Disk quantity.                                                                                                                                                                                                                |

**Table 6-59** Tag

| Parameter | Type   | Description |
|-----------|--------|-------------|
| key       | String | Tag key     |
| value     | String | Tag value   |

**Table 6-60** NodeDetail

| Parameter                 | Type   | Description                                  |
|---------------------------|--------|----------------------------------------------|
| running_status            | String | Running status                               |
| cpu_usage                 | String | CPU utilization rate of the HBase cluster    |
| memory_usage              | String | Memory utilization rate of the HBase cluster |
| disk_usage                | String | Disk usage                                   |
| total_memory              | String | Total memory. Unit: MB.                      |
| available_memory          | String | Available memory. Unit: MB.                  |
| total_hard_disk_space     | String | Total hard disk space. Unit: GB              |
| available_hard_disk_space | String | Available hard disk space. Unit: GB          |
| network_read              | String | Network read speed. Unit: Byte/s             |
| network_write             | String | Network write speed. Unit: Byte/s            |

**Table 6-61** ComponentInfo

| Parameter           | Type   | Description                                      |
|---------------------|--------|--------------------------------------------------|
| id                  | String | Component ID                                     |
| name                | String | Component name                                   |
| instance_group_name | String | Name of the group to which the component belongs |
| running_status      | String | Running status                                   |
| ha_status           | String | HA status                                        |
| config_status       | String | Configuration status                             |
| role_name           | String | Role                                             |
| role_short_name     | String | Abbreviation of a role name                      |
| role_type           | String | Role type                                        |
| service_name        | String | Service name                                     |
| pair_name           | String | Pair name.                                       |

| Parameter                  | Type   | Description                                |
|----------------------------|--------|--------------------------------------------|
| relation_pairs             | String | Associated pair                            |
| support_decom              | String | Whether Decom is supported                 |
| support_reinstall          | String | Whether reinstallation is supported        |
| support_collect_stack_info | String | Whether stack information can be collected |

## Example Request

Query all nodes in a specified node group of a cluster.

```
/v2/174ee662a7e24cc99bfc858c4558dbf6/clusters/f0a91b14-8884-4ba1-9e8d-0a21086c0ab4/nodes?node_group=master_node_default_group
```

## Example Response

**Status code: 200**

List node information.

```
{
  "nodes": [ {
    "node_name": "node-master1QxCW",
    "resource_id": "fc1ed6bb-e3d8-4dc8-8162-ef673bff6b7b",
    "node_group_name": "master_node_default_group",
    "node_type": "Master",
    "billing_type": "on-quantity",
    "deployment_type": "SERVER",
    "server_info": {
      "server_id": "3687f50a-cd4e-4c67-8858-5d6555c8834c",
      "server_name": "f28fb043-ecbc-401e-936e-fb321ca8d40a_node-master1QxCW",
      "server_type": "ECS",
      "data_volumes": [ {
        "type": "SATA",
        "size": 200,
        "count": 1
      } ],
      "root_volume": {
        "type": "SATA",
        "size": 100,
        "count": 1
      },
      "cpu_type": "X86",
      "internal_ip": "192.168.10.142"
    },
    "tags": [ {
      "key": "1",
      "value": "2"
    } ],
    "node_status": "started"
  } ],
  "node_total": 1
}
```



## Status Codes

| Status Codes | Description            |
|--------------|------------------------|
| 200          | List of queried nodes. |

## Error Codes

For details, see [Error Codes](#).

# 6.2 Job Management APIs

## 6.2.1 Adding and Executing a Job

### Function

This API is used to add and submit a job in an MRS cluster.

#### NOTE

- If you want to use the OBS encryption function, follow instructions in [Using OBS to Encrypt Data for Running Jobs](#) to configure related information and call an API to run a job.
- On the **Dashboard** tab page of the cluster details page, click **Click to synchronize** on the right side of **IAM User Sync** to synchronize IAM users. Then submit a job through this API.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/job-executions

**Table 6-62** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                              |
|------------|-----------|--------|----------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

| Parameter  | Mandatory | Type   | Description                                                                                           |
|------------|-----------|--------|-------------------------------------------------------------------------------------------------------|
| cluster_id | Yes       | String | Cluster ID. For details on how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> . |

## Request Parameters

**Table 6-63** Request parameters

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_type  | Yes       | String | <p>Type of a job.</p> <ul style="list-style-type: none"> <li>MapReduce</li> <li>SparkSubmit</li> <li>SparkPython: Example request of a SparkPython job (Jobs of this type will be converted to SparkSubmit jobs for submission. The job type is displayed as SparkSubmit on the MRS console. Select SparkSubmit when you call an API to query the job list.)</li> <li>HiveScript</li> <li>HiveSql</li> <li>DistCp, importing and exporting data</li> <li>SparkScript</li> <li>SparkSql</li> <li>Flink</li> </ul> <p><b>NOTE</b><br/>Spark, Hive, and Flink jobs can be added to only clusters that include Spark, Hive, and Flink components.</p> |
| job_name  | Yes       | String | <p>Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> <p><b>NOTE</b><br/>Identical job names are allowed but not recommended.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Parameter  | Mandatory | Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------|-----------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| arguments  | No        | Array of strings   | <p>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.</p> <p>The value can contain a maximum of 150,000 characters. Special characters (; &amp;&gt;'&lt;\$!\) are not allowed. This parameter can be left blank.</p> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• For MRS 1.9.2 or later, a file path on OBS can start with <b>obs://</b>. To use this format to submit HiveScript or HiveSQL jobs, choose <b>Components &gt; Hive &gt; Service Configuration</b> on the cluster details page. Set <b>Type</b> to <b>All</b>, and search for <b>core.site.customized.configs</b>. Add the endpoint configuration item <b>fs.obs.endpoint</b> of OBS and enter the endpoint corresponding to OBS in <b>Value</b>. For details, see <a href="#">Endpoints</a>.</li> <li>• For MRS 3.x or later, a file path on OBS can start with <b>obs://</b>. To use this format to submit HiveScript or HiveSQL jobs, choose <b>Components &gt; Hive &gt; Service Configuration</b> on the cluster details page. Set <b>Basic</b> to <b>All</b>, and search for <b>core.site.customized.configs</b>. Add the endpoint configuration item <b>fs.obs.endpoint</b> of OBS and enter the endpoint corresponding to OBS in <b>Value</b>. For details, see <a href="#">Endpoints</a>.</li> </ul> |
| properties | No        | Map<String,String> | <p>Program system parameter.</p> <p>The parameter contains a maximum of 2,048 characters, excluding special characters such as &gt;&lt; '&amp;! , and can be left blank.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Response Parameters

**Status code: 200**

**Table 6-64** Response body parameter

| Parameter         | Type                   | Description                                                                                  |
|-------------------|------------------------|----------------------------------------------------------------------------------------------|
| job_submit_result | JobSubmitResult object | The job execution result. For details about the parameters, see <a href="#">Table 6-65</a> . |

**Table 6-65** JobSubitResult parameters

| Parameter | Type   | Description                                                                                                                                                       |
|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_id    | String | Job ID                                                                                                                                                            |
| state     | String | Job submission status. <ul style="list-style-type: none"> <li><b>COMPLETE:</b> The job is submitted.</li> <li><b>FAILED:</b> Failed to submit the job.</li> </ul> |

**Status code: 400**

**Table 6-66** Response body parameters

| Parameter  | Type   | Description        |
|------------|--------|--------------------|
| error_code | String | The error code.    |
| error_msg  | String | The error message. |

## Example Request

You must have prepared the OBS paths, sample files, endpoints, and AKs/SKs when submitting a request.

- Create a MapReduce job.

POST `https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions`

```
{
  "job_name": "MapReduceTest",
  "job_type": "MapReduce",
  "arguments": [
    "obs://obs-test/program/hadoop-mapreduce-examples-x.x.x.jar",
    "wordcount",
    "obs://obs-test/input/",
    "obs://obs-test/job/mapreduce/output"
  ],
  "properties": {
    "fs.obs.endpoint": "obs endpoint",
    "fs.obs.access.key": "xxx",
    "fs.obs.secret.key": "yyy"
  }
}
```

- ```
}  
}
```
- **Create a SparkSubmit job.**  
POST [https://{endpoint}/v2/{project\\_id}/clusters/{cluster\\_id}/job-executions](https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions)  

```
{  
  "job_name":"SparkSubmitTest",  
  "job_type":"SparkSubmit",  
  "arguments":[  
    "--master",  
    "yarn",  
    "--deploy-mode",  
    "cluster",  
    "--py-files",  
    "obs://obs-test/a.py",  
    "--conf",  
    "spark.yarn.appMasterEnv.PYTHONPATH=/tmp:$PYTHONPATH",  
    "--conf",  
    "spark.yarn.appMasterEnv.aaa=aaaa",  
    "--conf",  
    "spark.executorEnv.aaa=executaaaa",  
    "--properties-file",  
    "obs://obs-test/test-spark.conf",  
    "obs://obs-test/pi.py",  
    "100000"  
  ],  
  "properties":{  
    "fs.obs.access.key":"xxx",  
    "fs.obs.secret.key":"yyy"  
  }  
}
```
  - **Create a HiveScript job.**  
POST [https://{endpoint}/v2/{project\\_id}/clusters/{cluster\\_id}/job-executions](https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions)  

```
{  
  "job_name":"HiveScriptTest",  
  "job_type":"HiveScript",  
  "arguments":[  
    "obs://obs-test/sql/test_script.sql"  
  ],  
  "properties":{  
    "fs.obs.endpoint":"obs endpoint",  
    "fs.obs.access.key":"xxx",  
    "fs.obs.secret.key":"yyy"  
  }  
}
```
  - **Create a HiveSql job.**  
POST [https://{endpoint}/v2/{project\\_id}/clusters/{cluster\\_id}/job-executions](https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions)  

```
{  
  "job_name" : "HiveSqlTest",  
  "job_type" : "HiveSql",  
  "arguments" : [ "DROP TABLE IF EXISTS src_wordcount;\ncreate external table src_wordcount(line  
string) row format delimited fields terminated by "\\n\\n" stored as textfile location \"obs://donotdel-  
gxc/input/\";\ninsert into src_wordcount values(\"v1\")" ],  
  "properties" : {  
    "fs.obs.endpoint" : "obs endpoint",  
    "fs.obs.access.key" : "xxx",  
    "fs.obs.secret.key" : "yyy"  
  }  
}
```
  - **Create a DistCp job.**  
POST [https://{endpoint}/v2/{project\\_id}/clusters/{cluster\\_id}/job-executions](https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions)  

```
{  
  "job_name":"DistCpTest",  
  "job_type":"DistCp",
```

```
"arguments":[
  "obs://obs-test/DistcpJob/",
  "/user/test/sparksql/"
],
"properties":{
  "fs.obs.endpoint":"obs endpoint",
  "fs.obs.access.key":"xxx",
  "fs.obs.secret.key":"yyy"
}
}
```

- **Create a SparkScript job.**

POST [https://{endpoint}/v2/{project\\_id}/clusters/{cluster\\_id}/job-executions](https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions)

```
{
  "job_name":"SparkScriptTest",
  "job_type":"SparkScript",
  "arguments":[
    "op-key1",
    "op-value1",
    "op-key2",
    "op-value2",
    "obs://obs-test/sql/test_script.sql"
  ],
  "properties":{
    "fs.obs.access.key":"xxx",
    "fs.obs.secret.key":"yyy"
  }
}
```

- **Create a SparkSql job.**

POST [https://{endpoint}/v2/{project\\_id}/clusters/{cluster\\_id}/job-executions](https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions)

```
{
  "job_name":"SparkSqlTest",
  "job_type":"SparkSql",
  "arguments":[
    "op-key1",
    "op-value1",
    "op-key2",
    "op-value2",
    "create table student_info3 (id string,name string,gender string,age int,addr string);"
  ],
  "properties":{
    "fs.obs.access.key":"xxx",
    "fs.obs.secret.key":"yyy"
  }
}
```

- **Create a Flink job.**

POST [https://{endpoint}/v2/{project\\_id}/clusters/{cluster\\_id}/job-executions](https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions)

```
{
  "job_name":"FlinkTest",
  "job_type":"Flink",
  "arguments":[
    "run",
    "-d",
    "-ynm",
    "testExcutorejobhdfsbatch",
    "-m",
    "yarn-cluster",
    "hdfs://test/examples/batch/WordCount.jar"
  ],
  "properties":{
    "fs.obs.endpoint":"obs endpoint",
    "fs.obs.access.key":"xxx",
    "fs.obs.secret.key":"yyy"
  }
}
```

- Create a SparkPython job (Jobs of this type will be converted to SparkSubmit jobs for submission. The job type is displayed as SparkSubmit on the MRS console. Select SparkSubmit when you call an API to query the job list.)  
POST `https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions`

```
{
  "job_name": "SparkPythonTest",
  "job_type": "SparkPython",
  "arguments": [ "--master", "yarn", "--deploy-mode", "cluster", "--py-files", "obs://obs-test/a.py", "--conf", "spark.yarn.appMasterEnv.PYTHONPATH=/tmp:$PYTHONPATH", "--conf", "spark.yarn.appMasterEnv.aaa=aaaa", "--conf", "spark.executorEnv.aaa=executoraaa", "--properties-file", "obs://obs-test/test-spark.conf", "obs://obs-test/pi.py", "100000" ],
  "properties": {
    "fs.obs.access.key": "xxx",
    "fs.obs.secret.key": "yyy"
  }
}
```

## Example Response

### Status code: 200

- Example of a successful response

```
{
  "job_submit_result": {
    "job_id": "44b37a20-ffe8-42b1-b42b-78a5978d7e40",
    "state": "COMPLETE"
  }
}
```

### Status code: 400

- Example of a failed response

```
{
  "error_msg": "Hive jobs cannot be submitted.",
  "error_code": "0168"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.2.2 Querying Information About a Job

### Function

This API is used to query information about a specified job in an MRS cluster.

### Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- Format  
GET /v2/{project\_id}/clusters/{cluster\_id}/job-executions/{job\_execution\_id}
- Parameter description

**Table 6-67** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details on how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .
job_execution_id	Yes	String	The job ID. For details about how to obtain the job ID, see <a href="#">Obtaining a Job ID</a> .

## Request Parameters

None

## Response Parameters

**Status code: 202**

**Table 6-68** Response parameter

Parameter	Type	Description
job_detail	JobQueryBean object	Job details. For details about the parameter, see <a href="#">Table 6-69</a> .

**Table 6-69** Job parameters

Parameter	Type	Description
job_id	String	The job ID.



Parameter	Type	Description
user	String	The name of the user who submits a job.
job_name	String	The job name.
job_result	String	The final result of a job. Possible values: <ul style="list-style-type: none"> <li>● <b>FAILED</b>: indicates that the job fails to be executed.</li> <li>● <b>KILLED</b>: indicates that the job is manually terminated during execution.</li> <li>● <b>UNDEFINED</b>: indicates that the job is being executed.</li> <li>● <b>SUCCEDED</b>: indicates that the job has been successfully executed.</li> </ul>
job_state	String	The job execution status. Possible values: <ul style="list-style-type: none"> <li>● <b>FAILED</b>: indicates that the job failed.</li> <li>● <b>KILLED</b>: indicates that the job is terminated.</li> <li>● <b>New</b>: indicates that the job is created.</li> <li>● <b>NEW_SAVING</b>: indicates that the job has been created and is being saved.</li> <li>● <b>SUBMITTED</b>: indicates that the job is submitted.</li> <li>● <b>ACCEPTED</b>: indicates that the job is accepted.</li> <li>● <b>RUNNING</b>: indicates that the job is running.</li> <li>● <b>FINISHED</b>: indicates that the job is completed.</li> </ul>
job_progress	Float	The job execution progress.

Parameter	Type	Description
job_type	String	The job type. Possible values: <ul style="list-style-type: none"> <li>• <b>MapReduce</b></li> <li>• <b>SparkSubmit</b>: Select this value when you call an API to query SparkPython jobs.</li> <li>• <b>HiveScript</b></li> <li>• <b>HiveSql</b></li> <li>• <b>DistCp</b>: imports and exports data.</li> <li>• <b>SparkScript</b></li> <li>• <b>SparkSql</b></li> <li>• <b>Flink</b></li> </ul>
started_time	Long	The time when a job starts to execute, in milliseconds.
submitted_time	Long	The time when a job is submitted, in milliseconds.
finished_time	Long	The time when a job execution is completed, in milliseconds.
elapsed_time	Long	The job execution duration, in milliseconds.
arguments	String	Run parameters.
properties	String	The configuration parameter, which is used to configure <b>-d</b> parameters. The value can contain a maximum of 2,048 characters. Special characters (; &>'<\$!\ \) are not allowed. This parameter can be left blank.
launcher_id	String	The launcher job ID.
app_id	String	The actual job ID.
tracking_url	String	The URL for accessing logs. Currently, only SparkSubmit jobs support this parameter. This parameter accesses the Yarn Web UI via the EIP bound to the cluster. If the EIP is unbound from the cluster on the VPC console, the MRS service data is not updated in a timely manner and the access fails. In this case, you can bind the EIP to the cluster again to rectify the fault.
queue	String	The type of the resource queue used by a job.

**Status code: 500**

**Table 6-70** Response body parameters

Parameter	Type	Description
error_code	String	The error code.
error_msg	String	The error message.

## Example Request

Query information about a job.

```
GET https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}
```

## Example Response

**Status code: 202**

Job information successful queried.

```
{
  "job_detail" : {
    "job_id" : "431b135e-c090-489f-b1db-0abe3822b855",
    "user" : "xxxx",
    "job_name" : "pyspark1",
    "job_result" : "SUCCEEDED",
    "job_state" : "FINISHED",
    "job_progress" : 100,
    "job_type" : "SparkSubmit",
    "started_time" : 1564626578817,
    "submitted_time" : 1564626561541,
    "finished_time" : 1564626664930,
    "elapsed_time" : 86113,
    "queue" : "default",
    "arguments" : "[--class, org.apache.spark.examples.SparkPi, --driver-memory, 512MB, --num-executors, 1,
--executor-cores, 1, --master, yarn-cluster, s3a://obs-test/jobs/spark/spark-examples_2.11-2.1.0.jar, 10000]",
    "launcher_id" : "application_1564622673393_0006",
    "app_id" : "application_1564622673393_0007",
    "properties" : "{}"
  }
}
```

**Status code: 400**

Failed to query information about a job.

```
{
  "error_msg" : "Failed to query the job."
  "error_code" : "0162"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.2.3 Querying a List of Jobs

### Function

This API is used to query the job list in a specified MRS cluster.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
GET /v2/{project\_id}/clusters/{cluster\_id}/job-executions
- Parameter description

**Table 6-71** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .

**Table 6-72** Query parameters

Parameter	Mandatory	Type	Description
job_name	No	String	The job name. The value can contain 1 to 128 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed.
job_id	No	String	The job ID. The value can contain 1 to 64 characters. Only letters, numbers, and hyphens (-) are allowed.

Parameter	Mandatory	Type	Description
user	No	String	The username. The value can contain 1 to 32 characters. Only letters, numbers, hyphens (-), underscores (_), and periods (.) are allowed, but the value cannot start with a number.
job_type	No	String	The job type. Possible values: <ul style="list-style-type: none"> <li>• <b>MapReduce</b></li> <li>• <b>SparkSubmit</b></li> <li>• <b>SparkSubmit</b>: Select this value when you call an API to query SparkPython jobs.</li> <li>• <b>HiveScript</b></li> <li>• <b>HiveSql</b></li> <li>• <b>DistCp</b>: imports and exports data.</li> <li>• <b>SparkScript</b></li> <li>• <b>SparkSql</b></li> <li>• <b>Flink</b></li> </ul>
job_state	No	String	The job execution status. Possible values: <ul style="list-style-type: none"> <li>• <b>FAILED</b>: indicates that the job fails to be executed.</li> <li>• <b>KILLED</b>: indicates that the job is terminated.</li> <li>• <b>New</b>: indicates that the job is created.</li> <li>• <b>NEW_SAVING</b>: indicates that the job has been created and is being saved.</li> <li>• <b>SUBMITTED</b>: indicates that the job is submitted.</li> <li>• <b>ACCEPTED</b>: indicates that the job is accepted.</li> <li>• <b>RUNNING</b>: indicates that the job is running.</li> <li>• <b>FINISHED</b>: indicates that the job is completed.</li> </ul>

Parameter	Mandatory	Type	Description
job_result	No	String	The job execution result. Possible values: <ul style="list-style-type: none"> <li>• <b>FAILED</b>: indicates that the job fails to be executed.</li> <li>• <b>KILLED</b>: indicates that the job is manually terminated during execution.</li> <li>• <b>UNDEFINED</b>: indicates that the job is being executed.</li> <li>• <b>SUCCEEDED</b>: indicates that the job has been successfully executed.</li> </ul>
queue	No	String	The type of the resource queue used by a job. The value can contain 1 to 64 characters. Only letters, numbers, and hyphens (-) are allowed.
limit	No	String	The number of records displayed on each page in the returned result. The default value is <b>10</b> .
offset	No	String	The default offset from which the job list starts to be queried is <b>1</b> .
sort_by	No	String	The sorting mode of returned results. The default value is <b>desc</b> . Possible values: <ul style="list-style-type: none"> <li>• <b>asc</b>: indicates that the returned results are ranked in ascending order.</li> <li>• <b>desc</b>: indicates that the returned results are ranked in descending order.</li> </ul>
submitted_time_begin	No	Long	The UTC timestamp after which a job is submitted, in milliseconds, for example, <b>1562032041362</b> .
submitted_time_end	No	Long	The UTC timestamp before which a job is submitted, in milliseconds, for example, <b>1562032041362</b> .

## Request Parameters

None

## Response Parameters

Status code: 202

**Table 6-73** Response body parameters

Parameter	Type	Description
total_record	Integer	The total number of records.
job_list	Array of JobQueryBean objects	The job list. For details about the parameters, see <a href="#">Table 6-74</a> .

**Table 6-74** JobQueryBean

Parameter	Type	Description
job_id	String	The job ID.
user	String	The name of the user who submits a job.
job_name	String	The job name.
job_result	String	The final result of a job. Possible values: <ul style="list-style-type: none"> <li>● <b>FAILED</b>: indicates that the job fails to be executed.</li> <li>● <b>KILLED</b>: indicates that the job is manually terminated during execution.</li> <li>● <b>UNDEFINED</b>: indicates that the job is being executed.</li> <li>● <b>SUCCEEDED</b>: indicates that the job has been successfully executed.</li> </ul>

Parameter	Type	Description
job_state	String	The job execution status. Possible values: <ul style="list-style-type: none"> <li>● <b>FAILED</b>: indicates that the job fails to be executed.</li> <li>● <b>KILLED</b>: indicates that the job is terminated.</li> <li>● <b>New</b>: indicates that the job is created.</li> <li>● <b>NEW_SAVING</b>: indicates that the job has been created and is being saved.</li> <li>● <b>SUBMITTED</b>: indicates that the job is submitted.</li> <li>● <b>ACCEPTED</b>: indicates that the job is accepted.</li> <li>● <b>RUNNING</b>: indicates that the job is running.</li> <li>● <b>FINISHED</b>: indicates that the job is completed.</li> </ul>
job_progress	Float	The job execution progress.
job_type	String	The job type. Possible values: <ul style="list-style-type: none"> <li>● <b>MapReduce</b></li> <li>● <b>SparkSubmit</b>: Select this value when you call an API to query SparkPython jobs.</li> <li>● <b>HiveScript</b></li> <li>● <b>HiveSql</b></li> <li>● <b>DistCp</b>: imports and exports data.</li> <li>● <b>SparkScript</b></li> <li>● <b>SparkSql</b></li> <li>● <b>Flink</b></li> </ul>
started_time	Long	The time when a job starts to execute, in milliseconds.
submitted_time	Long	The time when a job is submitted, in milliseconds.
finished_time	Long	The time when a job execution is completed, in milliseconds.
elapsed_time	Long	The job execution duration, in milliseconds.
arguments	String	The running parameters.



Parameter	Type	Description
properties	String	The configuration parameter, which is used to configure <b>-d</b> parameters. The value can contain a maximum of 2,048 characters. Special characters (; &>'<\$!\ \) are not allowed. This parameter can be left blank.
launcher_id	String	The actual job ID.
app_id	String	The actual job ID.
tracking_url	String	The URL for accessing logs. Currently, only SparkSubmit jobs support this parameter. This parameter accesses the Yarn Web UI via the EIP bound to the cluster. If the EIP is unbound from the cluster on the VPC console, the MRS service data is not updated in a timely manner and the access fails. In this case, you can bind the EIP to the cluster again to rectify the fault.
queue	String	The type of the resource queue used by a job.

**Status code: 500**

**Table 6-75** Response body parameters

Parameter	Type	Description
error_code	String	The error code.
error_msg	String	The error message.

## Example Response

**Status code: 202**

Querying a list of jobs is successful.

```
{
  "total_record" : 2,
  "job_list" : [ {
    "job_id" : "981374c1-85da-44ee-be32-edfb4fba776c",
    "user" : "xxx",
    "job_name" : "SparkSubmitTset",
    "job_result" : "UNDEFINED",
    "job_state" : "ACCEPTED",
    "job_progress" : 0,
    "job_type" : "SparkSubmit",
    "started_time" : 0,
    "submitted_time" : 1564714763119,
  }
]
```

```
"finished_time" : 0,
"elapsed_time" : 0,
"queue" : "default",
"arguments" : "[--class, --driver-memory, --executor-cores, --master, yarn-cluster, s3a://obs-test/hadoop-
mapreduce-examples-3.1.1.jar, dddd]",
"launcher_id" : "application_1564622673393_0613",
"properties" : { }
}, {
"job_id" : "c54c8aa0-c277-4f83-8acc-521d85cfa32b",
"user" : "xxxx",
"job_name" : "SparkSubmitTset2",
"job_result" : "UNDEFINED",
"job_state" : "ACCEPTED",
"job_progress" : 0,
"job_type" : "SparkSubmit",
"started_time" : 0,
"submitted_time" : 1564714020099,
"finished_time" : 0,
"elapsed_time" : 0,
"queue" : "default",
"arguments" : "[--conf, yujjsjhe, --driver-memory, yueujdjjd, --master, \nyarn-cluster, \ns3a://obs-test/
hadoop-mapreduce-examples-3.1.1.jar]",
"launcher_id" : "application_1564622673393_0611",
"properties" : { }
} ]
}
```

**Status code: 500**

Failed to query a list of jobs.

```
{
  "error_msg": "Failed to query the job list."
  "error_code" : "0166"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.2.4 Terminating a Job

### Function

This API is used to terminate a specified job in an MRS cluster.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- Format  
POST /v2/{project\_id}/clusters/{cluster\_id}/job-executions/{job\_execution\_id}/kill
- Parameter description

**Table 6-76** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .
job_execution_id	Yes	String	The job ID. For details about how to obtain the job ID, see <a href="#">Obtaining a Job ID</a> .

## Request Parameters

None

## Response Parameters

None

## Example Request

Terminate a job.

```
POST https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/kill
```

## Example Response

**Status code: 400**

Failed to terminate a specified job in an MRS cluster.

```
{
  "error_msg": "Failed to terminate the job."
  "error_code": "0175"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.2.5 Obtaining SQL Results

### Function

This API is used to obtain results returned after the SQL statements for querying SparkSQL and SparkScript jobs in an MRS cluster are executed.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
GET /v2/{project\_id}/clusters/{cluster\_id}/job-executions/{job\_execution\_id}/sql-result
- Parameter description

**Table 6-77** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .
job_execution_id	Yes	String	The job ID. For details about how to obtain the job ID, see <a href="#">Obtaining a Job ID</a> .

### Request Parameters

None

## Response Parameters

**Status code: 202**

**Table 6-78** Response parameter

Parameter	Type	Description
sql_results	Object	The query results of a SQL statement.

**Status code: 500**

**Table 6-79** Response body parameters

Parameter	Type	Description
error_code	String	The error code.
error_msg	String	The error message.

## Example Request

Example request for obtaining the SQL result

```
GET https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/sql-result
```

## Example Response

**Status code: 202**

The SQL statement execution result is obtained.

```
{
  "sql_results" : {
    "0" : [ {
      "result" : "succeed"
    } ],
    "1" : [ {
      "database" : "default",
      "isTemporary" : "false",
      "tableName" : "src_wordcount"
    } ],
    "2" : [ {
      "result" : "succeed"
    } ],
    "3" : [ {
      "result" : "succeed"
    } ],
    "4" : [ {
      "name" : "a",
      "id" : 1
    }, {
      "name" : "b",
      "id" : 2
    } ]
  }
}
```

**Status code: 500**

Failed to obtain the SQL statement execution result.

```
{
  "error_msg": "Failed to obtain SQL job results.",
  "error_code": "0172"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.2.6 Deleting Jobs in Batches

### Function

This API is used to delete jobs in batches.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/job-executions/batch-delete

**Table 6-80** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .

## Request Parameters

**Table 6-81** Request parameter

Parameter	Mandatory	Type	Description
job_id_list	No	Array of strings	The job ID list. For details about how to obtain the job ID, see <a href="#">Obtaining a Job ID</a> .

## Response Parameters

None

## Example Request

Delete jobs in batches.

```
POST https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/job-executions/batch-delete
{
  "job_id_list": [ "48c45725-b699-4aa9-9bfd-f7ff87eb6fe8", "af846665-dd32-4349-a8b5-561e109c383c" ]
}
```

## Example Response

None

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.3 Auto Scaling APIs

### 6.3.1 Viewing Auto Scaling Policies

#### Function

This API is used to view all auto scaling policies of a specified cluster.

#### Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

GET /v2/{project\_id}/autoscaling-policy/{cluster\_id}

**Table 6-82** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .

## Request Parameters

None

## Response Parameters

**Status code: 200**

**Table 6-83** Response body parameter

Parameter	Type	Description
[Array element]	Array of AutoScalingPolicyV2 objects	The auto scaling policy list. For details, see <a href="#">Table 6-84</a> .



**Table 6-84** AutoScalingPolicyV2

Parameter	Type	Description
node_group_name	String	The node group name. This parameter is mandatory. If <b>resource_pool_name</b> is <b>default</b> , the auto scaling policies are created by node group. If <b>resource_pool_name</b> is not <b>default</b> , policies of the resource pool corresponding to the node group are created.
resource_pool_name	String	The resource plan name. This parameter is mandatory. If the cluster version does not support auto scaling for a specified resource pool, set this parameter to <b>default</b> . If this parameter is not <b>default</b> , the auto scaling policy of the specified resource pool is deleted.
auto_scaling_policy	AutoScalingPolicyInfo object	The auto scaling policy. For details, see <a href="#">Table 6-85</a> .

**Table 6-85** AutoScalingPolicyInfo

Parameter	Type	Description
auto_scaling_enable	Boolean	Whether to enable the auto scaling policy.
min_capacity	Integer	The minimum number of nodes reserved in the node group. The value ranges from 0 to 500.
max_capacity	Integer	The maximum number of nodes in the node group. The value ranges from 0 to 500.
resources_plans	Array of ResourcesPlan objects	The resource plan list. If this parameter is left blank, the resource plan is disabled. When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>rules</b> must be configured. For details about this parameter, see <a href="#">Table 6-86</a> .

Parameter	Type	Description
rules	Array of Rule objects	List of auto scaling rules. When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>resources_plans</b> must be configured. For details about this parameter, see <a href="#">Table 6-87</a> .
tags	Array of Tag objects	Auto scaling tags. For details, see <a href="#">Table 6-89</a> .

**Table 6-86** ResourcesPlan

Parameter	Type	Description
period_type	String	The cycle type of a resource plan. Currently, only the following cycle type is supported: daily
start_time	String	The start time of a resource plan. The value is in the format of <b>hour:minute</b> , indicating that the time ranges from 00:00 to 23:59.
end_time	String	The end time of a resource plan. The value is in the same format as that of <b>start_time</b> . The interval between <b>end_time</b> and <b>start_time</b> must be greater than or equal to 30 minutes.
min_capacity	Integer	The minimum number of reserved nodes in a node group in a resource plan. The value ranges from 0 to 500.
max_capacity	Integer	The maximum number of reserved nodes in a node group in a resource plan. The value ranges from 0 to 500.
effective_days	Array of strings	Effective date of a resource plan. If this parameter is left blank, it indicates that the resource plan takes effect every day. The options are as follows: <b>MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</b>

**Table 6-87** Rule

Parameter	Type	Description
name	String	The name of an auto scaling rule. The name can contain only 1 to 64 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed. Rule names must be unique in a node group.
description	String	The description about an auto scaling rule. It contains a maximum of 1,024 characters.
adjustment_type	String	Auto scaling rule adjustment type. Possible values: <ul style="list-style-type: none"> <li>• <b>scale_out</b>: cluster scale-out</li> <li>• <b>scale_in</b>: cluster scale-in</li> </ul>
cool_down_minutes	Integer	The cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute. The value ranges from 0 to 10080. One week is equal to 10,080 minutes.
scaling_adjustment	Integer	The number of cluster nodes that can be adjusted at a time. The value ranges from 1 to 100.
trigger	Trigger object	The condition for triggering a rule. For details, see <a href="#">Table 6-88</a> .

**Table 6-88** Trigger

Parameter	Type	Description
metric_name	String	The metric name. This triggering condition makes a judgment according to the value of the metric. A metric name contains a maximum of 64 characters.
metric_value	String	The metric threshold to trigger a rule. The value must be an integer or a number with two decimal places.

Parameter	Type	Description
comparison_operator	String	The metric judgment logic operator. Possible values: <ul style="list-style-type: none"> <li>• <b>LT</b>: less than</li> <li>• <b>GT</b>: greater than</li> <li>• <b>LTOE</b>: less than or equal to</li> <li>• <b>GTOE</b>: greater than or equal to</li> </ul>
evaluation_periods	Integer	The number of consecutive five-minute periods, during which a metric threshold is reached. The value ranges from 1 to 288.

**Table 6-89** Tag

Parameter	Type	Description
key	String	Tag key <ul style="list-style-type: none"> <li>• The tag key can contain a maximum of 36 Unicode characters and cannot be an empty string.</li> <li>• The tag key cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\_  /).</li> <li>• The tag key of a resource must be unique.</li> </ul>
value	String	Tag value <ul style="list-style-type: none"> <li>• The value can contain 0 to 43 unicode characters.</li> <li>• The tag value cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\_  /).</li> </ul>

**Status code: 400**

**Table 6-90** Response body parameters

Parameter	Type	Description
error_code	String	Error codes.
error_msg	String	The error message.

## Example Request

None

## Example Response

**Status code: 200**

Auto scaling policies are displayed.

```
{
  "auto_scaling_policies" : [ {
    "node_group_name" : "task_node_analysis_group",
    "resource_pool_name" : "default",
    "auto_scaling_policy" : {
      "auto_scaling_enable" : true,
      "min_capacity" : 0,
      "max_capacity" : 1,
      "resources_plans" : [ {
        "period_type" : "daily",
        "effective_days" : [ "SUNDAY" ],
        "start_time" : "12:00",
        "end_time" : "13:00",
        "min_capacity" : 2,
        "max_capacity" : 3
      } ],
      "rules" : [ {
        "name" : "default-expand-1",
        "description" : "",
        "adjustment_type" : "scale_out",
        "cool_down_minutes" : 5,
        "scaling_adjustment" : 1,
        "trigger" : {
          "metric_name" : "YARNAppRunning",
          "metric_value" : 100,
          "comparison_operator" : "GTOE",
          "evaluation_periods" : 1
        }
      } ]
    }
  } ]
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.3.2 Updating an Auto Scaling Policy

### Function

This API is used to update an auto scaling policy.

### Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

## URI

PUT /v2/{project\_id}/autoscaling-policy/{cluster\_id}

**Table 6-91** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID.
cluster_id	Yes	String	Cluster ID.

## Request Parameters

**Table 6-92** Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	Node group name. This parameter is mandatory. If <b>resource_pool_name</b> is set to <b>default</b> , the auto scaling policies are created by node group. If <b>resource_pool_name</b> is not set to <b>default</b> , policies of the resource pool corresponding to the node group are created.
resource_pool_name	Yes	String	Pool name. This parameter is mandatory. If the cluster version does not support auto scaling for a specified resource pool, set this parameter to <b>default</b> . If this parameter is not <b>default</b> , the auto scaling policy of the specified resource pool is deleted.
auto_scaling_policy	No	AutoScalingPolicyInfo object	Auto scaling rule. For details about the parameters, see <a href="#">Table 6-93</a> .

**Table 6-93** AutoScalingPolicyInfo

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	Whether to enable the auto scaling policy.
min_capacity	Yes	Integer	Minimum number of nodes allowed in the node group. Value range: [0, 500]
max_capacity	Yes	Integer	The maximum number of nodes in the node group. Value range: [0, 500]
resources_plans	No	Array of ResourcesPlan objects	Resource plans. For details, see <a href="#">Table 6-94</a> . If this parameter is left blank, the resource plan is disabled. When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>rules</b> must be configured.
rules	No	Array of Rule objects	Auto scaling rules. When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>resources_plans</b> must be configured. For details about this parameter, see <a href="#">Table 6-95</a> .
tags	No	Array of Tag objects	Tags of an auto scaling rule. For details about the parameters, see <a href="#">Table 6-97</a> .

**Table 6-94** ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	The cycle type of a resource plan. Currently, only the following cycle type is supported: daily
start_time	Yes	String	The start time of a resource plan. The value is in the format of <b>hour:minute</b> , indicating that the time ranges from 00:00 to 23:59.

Parameter	Mandatory	Type	Description
end_time	Yes	String	The end time of a resource plan. The value is in the same format as that of <b>start_time</b> . The interval between <b>end_time</b> and <b>start_time</b> must be greater than or equal to 30 minutes.
min_capacity	Yes	Integer	The minimum number of reserved nodes in a node group in a resource plan. Value range: [0, 500]
max_capacity	Yes	Integer	The maximum number of reserved nodes in a node group in a resource plan. Value range: [0, 500]
effective_days	No	Array of strings	The effective date of a resource plan. If this parameter is left blank, it indicates that the resource plan takes effect every day. The options are as follows: <b>MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</b>

**Table 6-95** Rule

Parameter	Mandatory	Type	Description
name	Yes	String	The name of an auto scaling rule. A cluster name can contain only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. Rule names must be unique in a node group.
description	No	String	The description about an auto scaling rule. It contains a maximum of 1,024 characters.



Parameter	Mandatory	Type	Description
adjustment_type	Yes	String	The adjustment type of an auto scaling rule. The options are as follows: <ul style="list-style-type: none"> <li>• <b>scale_out</b>: cluster scale-out</li> <li>• <b>scale_in</b>: cluster scale-in</li> </ul>
cool_down_minutes	Yes	Integer	The cluster cooling time after an auto scaling rule is triggered, in minutes, during which period no auto scaling operation is performed. Value range: [0,10080]. One week is equal to 10,080 minutes.
scaling_adjustment	Yes	Integer	The number of cluster nodes that can be adjusted at a time. Value range: [1, 100]
trigger	Yes	Trigger object	Condition for triggering a rule. For details, see <a href="#">Table 6-96</a> .

**Table 6-96** Trigger

Parameter	Mandatory	Type	Description
metric_name	Yes	String	Metric name. This triggering condition makes a judgment according to the value of the metric. A metric name contains a maximum of 64 characters.
metric_value	Yes	String	The metric threshold to trigger a rule. The value must be an integer or a number with two decimal places.
comparison_operator	No	String	The metric judgment operator. The options are as follows: <ul style="list-style-type: none"> <li>• <b>LT</b>: less than</li> <li>• <b>GT</b>: greater than</li> <li>• <b>LTOE</b>: less than or equal to</li> <li>• <b>GTOE</b>: greater than or equal to</li> </ul>

Parameter	Mandatory	Type	Description
evaluation_periods	Yes	Integer	The number of consecutive five-minute periods, during which a metric threshold is reached. The value ranges from 1 to 288.

**Table 6-97** Tag

Parameter	Mandatory	Type	Description
key	Yes	String	The tag key. <ul style="list-style-type: none"> <li>It contains a maximum of 36 Unicode characters and cannot be an empty string.</li> <li>The tag key cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\, /).</li> <li>The tag key of a resource must be unique.</li> </ul>
value	Yes	String	The tag value. <ul style="list-style-type: none"> <li>The value can contain 0 to 43 unicode characters.</li> <li>The tag value cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\, /).</li> </ul>

## Response Parameters

None

## Example Request

Update an auto scaling policy, set the number of node groups in **node\_group\_1** to 0 to 5, and configure two scaling rules. If the value of **YARNAppRunning** is greater than or equal to 75 for one 5-minute period, add one Task node. If the value of **YARNAppRunning** is less than or equal to 25 for one 5-minute period, stop one Task node. The cooling duration is 20 minutes, and the **aaa=bbb** label is added to the involved nodes.

```
/v2/{project_id}/autoscaling-policy/{cluster_id}
```

```
{
  "node_group_name" : "ttt",
```

```
"auto_scaling_policy" : {
  "auto_scaling_enable" : true,
  "min_capacity" : 0,
  "max_capacity" : 5,
  "rules" : [ {
    "name" : "default-expand-1",
    "adjustment_type" : "scale_out",
    "cool_down_minutes" : 20,
    "scaling_adjustment" : 1,
    "trigger" : {
      "metric_name" : "YARNAppRunning",
      "metric_value" : 75,
      "comparison_operator" : "GT",
      "evaluation_periods" : 1
    }
  }, {
    "name" : "default-shrink-1",
    "adjustment_type" : "scale_in",
    "cool_down_minutes" : 20,
    "scaling_adjustment" : 1,
    "trigger" : {
      "metric_name" : "YARNAppRunning",
      "metric_value" : 25,
      "comparison_operator" : "LT",
      "evaluation_periods" : 1
    }
  } ],
  "resources_plans" : [ {
    "period_type" : "daily",
    "start_time" : "06:00",
    "end_time" : "20:00",
    "min_capacity" : "0",
    "max_capacity" : "2",
    "effective_days" : null
  } ],
  "tags" : [ {
    "key" : "aaa",
    "value" : "bbb"
  } ]
},
"resource_pool_name" : "default"
}
```

## Example Response

None

## Status Codes

For details, see [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.3.3 Deleting an AS policy

### Function

This API is used to delete an auto scaling policy.

## Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

## URI

DELETE /v2/{project\_id}/autoscaling-policy/{cluster\_id}

**Table 6-98** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID.
cluster_id	Yes	String	Cluster ID.

## Request Parameters

**Table 6-99** Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	Node group name. If <b>resource_pool_name</b> is set to <b>default</b> , the auto scaling policies are deleted by node group. If <b>resource_pool_name</b> is not set to <b>default</b> , policies of the resource pool corresponding to the node group are deleted.
resource_pool_name	Yes	String	Pool name. If the cluster version does not support auto scaling for a specified resource pool, set this parameter to <b>default</b> . If this parameter is not <b>default</b> , the auto scaling policy of the specified resource pool is deleted.

## Response Parameters

None

## Example Request

- Delete the auto scaling policies of the **node\_group\_2** node group.  
/v2/174ee662a7e24cc99bfc858c4558dbf6/autoscaling-policy/daf42ff0-05bc-4a1e-afbf-42b3131a1295  

```
{
  "node_group_name": "node_group_2"
}
```
- Delete the resource pool policy **resource\_1** node group policies.  
/v2/174ee662a7e24cc99bfc858c4558dbf6/autoscaling-policy/daf42ff0-05bc-4a1e-afbf-42b3131a1295  

```
{
  "node_group_name": "node_group_2",
  "resource_pool_name": "resource_1"
}
```

## Example Response

None

## Status Codes

For details, see [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.3.4 Creating an AS policy

### Function

This API is used to create an AS policy.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

POST /v2/{project\_id}/autoscaling-policy/{cluster\_id}

**Table 6-100** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID.
cluster_id	Yes	String	Cluster ID.

## Request Parameters

**Table 6-101** Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	Node group name. This parameter is mandatory. If <b>resource_pool_name</b> is set to <b>default</b> , the auto scaling policies are created by node group. If <b>resource_pool_name</b> is not set to <b>default</b> , policies of the resource pool corresponding to the node group are created.
resource_pool_name	Yes	String	Pool name. This parameter is mandatory. If the cluster version does not support auto scaling for a specified resource pool, set this parameter to <b>default</b> . If this parameter is not <b>default</b> , the auto scaling policy of the specified resource pool is deleted.
auto_scaling_policy	No	AutoScalingPolicyInfo object	Auto scaling rule. For details about the parameters, see <a href="#">Table 6-102</a> .

**Table 6-102** AutoScalingPolicyInfo

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	Whether to enable the auto scaling policy.
min_capacity	Yes	Integer	Minimum number of nodes allowed in the node group. Value range: [0, 500]
max_capacity	Yes	Integer	The maximum number of nodes in the node group. Value range: [0, 500]

Parameter	Mandatory	Type	Description
resources_plans	No	Array of ResourcesPlan objects	Resource plans. For details, see <a href="#">Table 6-103</a> . If this parameter is left blank, the resource plan is disabled. When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>rules</b> must be configured.
rules	No	Array of Rule objects	Auto scaling rules. For details, see <a href="#">Table 6-104</a> . When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>resources_plans</b> must be configured.
tags	No	Array of Tag objects	Tags of an auto scaling rule. For details about the parameters, see <a href="#">Table 6-106</a> .

**Table 6-103** ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	The cycle type of a resource plan. Currently, only the following cycle type is supported: daily
start_time	Yes	String	The start time of a resource plan. The value is in the format of <b>hour:minute</b> , indicating that the time ranges from 00:00 to 23:59.
end_time	Yes	String	The end time of a resource plan. The value is in the same format as that of <b>start_time</b> . The interval between <b>end_time</b> and <b>start_time</b> must be greater than or equal to 30 minutes.
min_capacity	Yes	Integer	The minimum number of reserved nodes in a node group in a resource plan. Value range: [0, 500]

Parameter	Mandatory	Type	Description
max_capacity	Yes	Integer	The maximum number of reserved nodes in a node group in a resource plan. Value range: [0, 500]
effective_days	No	Array of strings	The effective date of a resource plan. If this parameter is left blank, it indicates that the resource plan takes effect every day. The options are as follows: <b>MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</b>

**Table 6-104** Rule

Parameter	Mandatory	Type	Description
name	Yes	String	The name of an auto scaling rule. A cluster name can contain only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. Rule names must be unique in a node group.
description	No	String	The description about an auto scaling rule. It contains a maximum of 1,024 characters.
adjustment_type	Yes	String	The adjustment type of an auto scaling rule. The options are as follows: <ul style="list-style-type: none"> <li>● <b>scale_out</b>: cluster scale-out</li> <li>● <b>scale_in</b>: cluster scale-in</li> </ul>
cool_down_minutes	Yes	Integer	The cluster cooling time after an auto scaling rule is triggered, in minutes, during which period no auto scaling operation is performed. The value ranges from 0 to 10080. One week is equal to 10,080 minutes.



Parameter	Mandatory	Type	Description
scaling_adjustment	Yes	Integer	The number of cluster nodes that can be adjusted at a time. Value range: [1, 100]
trigger	Yes	Trigger object	Condition for triggering a rule. For details, see <a href="#">Table 6-105</a> .

**Table 6-105** Trigger

Parameter	Mandatory	Type	Description
metric_name	Yes	String	Metric name. This triggering condition makes a judgment according to the value of the metric. A metric name contains a maximum of 64 characters.
metric_value	Yes	String	The metric threshold to trigger a rule. The value must be an integer or a number with two decimal places.
comparison_operator	No	String	The metric judgment operator. The options are as follows: <ul style="list-style-type: none"> <li>• <b>LT</b>: less than</li> <li>• <b>GT</b>: greater than</li> <li>• <b>LTOE</b>: less than or equal to</li> <li>• <b>GTOE</b>: greater than or equal to</li> </ul>
evaluation_periods	Yes	Integer	The number of consecutive five-minute periods, during which a metric threshold is reached. The value ranges from 1 to 288.

**Table 6-106** Tag

Parameter	Mandatory	Type	Description
key	Yes	String	<p>The tag key.</p> <ul style="list-style-type: none"> <li>It contains a maximum of 36 Unicode characters and cannot be an empty string.</li> <li>The tag key cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\\,/).</li> <li>The tag key of a resource must be unique.</li> </ul>
value	Yes	String	<p>The tag value.</p> <ul style="list-style-type: none"> <li>The value can contain 0 to 43 unicode characters.</li> <li>The tag value cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\\,/).</li> </ul>

## Response Parameters

None

## Example Request

Create an auto scaling policy, set the number of node groups in **node\_group\_1** to 0 to 5, and configure two scaling rules. If the value of **YARNAppRunning** is greater than or equal to 75 for one 5-minute period, add one Task node. If the value of **YARNAppRunning** is less than or equal to 25 for one 5-minute period, stop one Task node. The cooling duration is 20 minutes, and the **aaa=bbb** label is added to the involved nodes.

```
/v2/{project_id}/autoscaling-policy/{cluster_id}
```

```
{
  "node_group_name": "node_group_1",
  "auto_scaling_policy": {
    "auto_scaling_enable": true,
    "min_capacity": 0,
    "max_capacity": 5,
    "rules": [ {
      "name": "default-expand-1",
      "adjustment_type": "scale_out",
      "cool_down_minutes": 20,
      "scaling_adjustment": 1,
      "trigger": {
        "metric_name": "YARNAppRunning",
        "metric_value": "75",
        "comparison_operator": "GT",
```

```
    "evaluation_periods" : 1
  }
}, {
  "name" : "default-shrink-1",
  "adjustment_type" : "scale_in",
  "cool_down_minutes" : 20,
  "scaling_adjustment" : 1,
  "trigger" : {
    "metric_name" : "YARNAppRunning",
    "metric_value" : "25",
    "comparison_operator" : "LT",
    "evaluation_periods" : 1
  }
}],
"resources_plans" : [ {
  "period_type" : "daily",
  "start_time" : "06:00",
  "end_time" : "20:00",
  "min_capacity" : "0",
  "max_capacity" : "2",
  "effective_days" : [ "MONDAY" ]
} ],
"tags" : [ {
  "key" : "aaa",
  "value" : "bbb"
} ]
},
"resource_pool_name" : "default"
}
```

## Example Response

None

## Status Codes

For details, see [Status Codes](#).

## Error Codes

See [Error Codes](#).

# 6.4 Cluster HDFS File API

## 6.4.1 Obtaining the List of Files from a Specified Directory

### Function

This API is used to obtain the list of files from a specified directory in an MRS cluster.

### Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

GET /v2/{project\_id}/clusters/{cluster\_id}/files

**Table 6-107** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .
path	Yes	String	<p>The file directory. For example, to access the <code>/tmp/test</code> directory list, the value must be a directory. The overall URI is as follows:</p> <pre><code>/v2/{project_id}/clusters/{cluster_id}/files?path=%2Ftmp%2Ftest</code></pre> <p>A single-level directory must comply with the following rules:</p> <ul style="list-style-type: none"> <li>• The directory path cannot be left blank.</li> <li>• The value cannot start or end with a period (.).</li> <li>• The value cannot contain the following characters: <code>/*? "&lt;&gt;  \;&amp;,'!{} [] \$%+</code></li> <li>• The value cannot exceed 255 characters.</li> </ul>
offset	No	String	Pagination parameter. The file list is queried from the offset. The default value is 1.

Parameter	Mandatory	Type	Description
limit	No	String	The pagination parameter, indicating the maximum number of records on a page. The default value is <b>100</b> and the maximum value is <b>1000</b> .
sort_key	No	String	The list is sorted by this attribute. The default value is <b>path_suffix</b> . The following attributes are supported: <ul style="list-style-type: none"> <li>• <b>path_suffix</b>: file or directory name</li> <li>• <b>length</b>: file size</li> <li>• <b>modification_time</b>: modification time</li> </ul>
order	No	String	The list sorting mode. Possible values: <ul style="list-style-type: none"> <li>• <b>desc</b>: Files are displayed in descending order.</li> <li>• <b>asc</b>: Files are displayed in ascending order.</li> </ul> The default value is <b>desc</b> .

## Request Parameters

None

## Response Parameters

**Table 6-108** Response parameters

Parameter	Type	Description
total_count	Integer	The total number of files, which is irrelevant to pagination.
files	Array of FileStatusV2 objects	The file list. For details, see <a href="#">Table 6-109</a> .

**Table 6-109** FileStatus parameters

Parameter	Type	Description
path_suffix	String	The file name extension in the current directory. For example, if you obtain the <code>/tmp/test</code> file in the <code>/tmp</code> directory, the value of <code>path_suffix</code> is <code>test</code> .
owner	String	The file owner.
group	String	The file owner group.
permission	String	The permission information.
replication	Integer	The number of replicas.
block_size	Integer	The block size.
length	Integer	The file length.
type	String	The file type. Possible values: <ul style="list-style-type: none"> <li><b>FILE</b>: file</li> <li><b>DIRECTORY</b>: directory</li> </ul>
children_num	Integer	The number of files in the directory.
access_time	Long	The file access time.
modification_time	Long	The file modification time.

## Example Request

```
GET /v2/{project_id}/clusters/{cluster_id}/files?
path={directory}&offset={offset}&limit={limit}&sort_key={sort_key}&order={order}
```

## Example Response

**Status code: 200**

Obtaining the file list of a specified directory is successful.

```
{
  "total_count": 2,
  "files": [
    {
      "access_time": 0,
      "block_size": 0,
      "children_num": 0,
      "group": "hadoop",
      "length": 0,
      "modification_time": 1587179516623,
      "owner": "hdfs",
      "path_suffix": "app-logs",
      "permission": "777",
      "replication": 0,
      "type": "DIRECTORY"
    },
    {
      "access_time": 1587267212761,
```

```
"block_size": 134217728,  
"children_num": 0,  
"group": "hadoop",  
"length": 23666188,  
"modification_time": 1587222156003,  
"owner": "root",  
"path_suffix": "data-m-00000",  
"permission": "644",  
"replication": 3,  
"type": "FILE"  
}  
]  
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

# 6.5 SQL APIs

## 6.5.1 Submitting a SQL Statement

### Function

This API is used to submit and execute a SQL statement in an MRS cluster.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
POST /v2/{project\_id}/clusters/{cluster\_id}/sql-execution
- Parameter description

**Table 6-110** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .

## Request Parameters

**Table 6-111** Request parameters

Parameter	Mandatory	Type	Description
sql_type	Yes	String	The SQL type. Currently, only the SQL of the presto type is supported. <b>NOTE</b> <ul style="list-style-type: none"> <li>The SQL of the presto type can be submitted and executed only in clusters containing the Presto component.</li> <li>This parameter is available only for MRS 2.0.6.1, 2.1.0.7, 3.1.2, or later clusters.</li> </ul>
sql_content	Yes	String	The SQL statement to be executed. <b>NOTE</b> Currently, only a single SQL statement can be executed at a time, and the statement cannot contain a semicolon (;).
database	No	String	The database where the SQL statement to be executed is located. The default value is <b>default</b> .
archive_path	No	String	The directory for storing the dumped SQL execution results. <b>NOTE</b> Only the <b>SELECT</b> statement dumps query results. Currently, the query results can be dumped only to OBS.



## Response Parameters

Status code: 200

Table 6-112 Response parameters

Parameter	Type	Description
id	String	The execution ID of a SQL statement. The ID is generated only when the <b>SELECT</b> , <b>SHOW</b> , or <b>DESC</b> statement is executed. For other operations, the ID is empty.
message	String	The error message.
statement	String	The SQL statement to be executed.
status	String	The execution status of a SQL statement. Possible values: <ul style="list-style-type: none"> <li>• <b>QUEUED</b></li> <li>• <b>WAITING_FOR_RESOURCES</b></li> <li>• <b>PLANNING</b></li> <li>• <b>STARTING</b></li> <li>• <b>RUNNING</b></li> <li>• <b>FINISHING</b></li> <li>• <b>FINISHED</b></li> <li>• <b>FAILED</b></li> </ul>
result_location	String	The path for archiving the final results of the SQL query statement. <b>NOTE</b> Only the <b>SELECT</b> statement dumps the SQL execution results to <b>result_location</b> .
content	Array<Array<String>>	The execution result of the SQL statement. <b>NOTE</b> Only non- <b>SELECT</b> statements return results in <b>content</b> . If the SQL statement does not return results, <b>content</b> is empty.

Status code: 400

Table 6-113 Response body parameters

Parameter	Type	Description
error_code	String	The error code.
error_msg	String	The error message.

## Example Request

Submit a Presto SQL statement.

```
POST https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/sql-execution
{
  "sql_type": "presto",
  "sql_content": "show tables",
  "database": "default",
  "archive_path": "obs://my-bucket/path"
}
```

## Example Response

**Status code: 200**

The SQL statement is submitted successfully.

```
{
  "id": "20190909_011820_00151_xxxxx",
  "statement": "show tables",
  "status": "FINISHED",
  "result_location": "obs://my_bucket/uuid_date/xxxx.csv",
  "content": [ [ "t1", null ], [ null, "t2" ], [ null, "t3" ] ]
}
```

**Status code: 400**

Failed to submit the SQL statement.

```
{
  "error_code": "MRS.0011",
  "message": "Failed to submit SQL to the executor. The cluster ID is xxxx"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.5.2 Querying SQL Results

### Function

This API is used to query the execution result of a SQL statement in an MRS cluster.

### Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- Format  
GET /v2/{project\_id}/clusters/{cluster\_id}/sql-execution/{sql\_id}
- Parameter description

**Table 6-114** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	Cluster ID. For details on how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .
sql_id	Yes	String	The execution ID of a SQL statement, that is, <b>sql_id</b> in the return result in <a href="#">Submitting a SQL Statement</a> .

## Request Parameters

None

## Response Parameters

**Status code: 200**

**Table 6-115** Response parameters

Parameter	Type	Description
id	String	The execution ID of a SQL statement. The ID is generated only when the <b>SELECT</b> , <b>SHOW</b> , or <b>DESC</b> statement is executed. For other operations, the ID is empty.
message	String	The error message.
statement	String	The SQL statement to be executed.

Parameter	Type	Description
status	String	The execution status of a SQL statement. Possible values: <ul style="list-style-type: none"> <li>• <b>QUEUED</b></li> <li>• <b>WAITING_FOR_RESOURCES</b></li> <li>• <b>PLANNING</b></li> <li>• <b>STARTING</b></li> <li>• <b>RUNNING</b></li> <li>• <b>FINISHING</b></li> <li>• <b>FINISHED</b></li> <li>• <b>FAILED</b></li> </ul>
result_location	String	The path for archiving the final results of the SQL query statement. <b>NOTE</b> Only the <b>SELECT</b> statement dumps the SQL execution results to <b>result_location</b> .
content	Array	The execution result of the SQL statement. <b>NOTE</b> Only non-SELECT statements return results in <b>content</b> . If the SQL statement does not return results, <b>content</b> is empty.

**Status code: 400**

**Table 6-116** Response body parameters

Parameter	Type	Description
error_code	String	The error code.
error_msg	String	The error message.

## Example Request

Example request for querying the SQL execution result

```
GET https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}
```

## Example Response

**Status code: 200**

Querying the SQL execution result is successful.

```
{
  "id": "20190909_011820_00151_xxxxx",
```

```
"statement" : "show tables",
"status" : "FINISHED",
"result_location" : "obs://my_bucket/uuid_date/xxxx.csv",
"content" : [ [ "t1", null ], [ null, "t2" ], [ null, "t3" ] ]
}
```

**Status code: 400**

Failed to query the SQL execution result.

```
{
"error_code" : "MRS.0011",
"message": "Failed to submit SQL to the executor. The cluster ID is xxxx"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.5.3 Canceling a SQL Execution Task

### Function

This API is used to cancel the execution task of a SQL statement in an MRS cluster.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/sql-execution/{sql\_id}/cancel

**Table 6-117** URI parameters

Parameter	Mandatory	Description
project_id	Yes	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .

Parameter	Mandatory	Description
sql_id	Yes	The execution ID of a SQL statement, that is, <b>sql_id</b> in the return result in <a href="#">Submitting a SQL Statement</a> .

## Request Parameters

None

## Response Parameters

Status code: 200

**Table 6-118** Response parameters

Parameter	Type	Description
message	String	The error message.
status	String	The result of canceling the execution of the SQL statement. Possible values: By default, <b>SUCCEED</b> is returned. <b>SUCCEED</b> is also returned when the task is completed. Only when the running SQL statement fails to be canceled, <b>FAILED</b> is returned. Enumerated values: <ul style="list-style-type: none"> <li>• <b>SUCCEED</b></li> <li>• <b>FAILED</b></li> </ul>

Status code: 400

**Table 6-119** Response body parameter

Parameter	Type	Description
message	String	The error message.
status	String	The result of canceling the execution of the SQL statement. Possible values: By default, <b>SUCCEED</b> is returned. <b>SUCCEED</b> is also returned when the task is completed. Only when the running SQL statement fails to be canceled, <b>FAILED</b> is returned. Enumerated values: <ul style="list-style-type: none"> <li>• <b>SUCCEED</b></li> <li>• <b>FAILED</b></li> </ul>

## Example Request

Cancel a SQL execution task.

```
POST https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}/cancel
```

## Example Response

**Status code: 200**

Canceling a SQL execution task is successful.

```
{  
  "status": "SUCCEED"  
}
```

**Status code: 400**

Failed to cancel a SQL execution task.

```
{  
  "status": "FAILED",  
  "message": "Cancel sql error"  
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

# 6.6 Agency Management

## 6.6.1 Querying the Mapping Between a User (Group) and an IAM Agency

### Function

This API is used to obtain details about the mapping between a user or user group and an IAM agency.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- URI format  
GET /v2/{project\_id}/clusters/{cluster\_id}/agency-mapping
- Parameter description

**Table 6-120** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .

## Request Parameters

None

## Response Parameters

Status code: 200

**Table 6-121** Response parameter

Parameter	Type	Description
agency_mappings	Array of AgencyMapping objects	The mapping between users or user groups and agencies. For details, see <a href="#">Table 6-122</a> .

**Table 6-122** agency\_mappings parameters

Parameter	Type	Description
agency	String	The name of the IAM agency bound to the mapping.
identifier_type	String	The agency type. Possible values: <ul style="list-style-type: none"> <li>• <b>User</b>: indicates that the mapping is for users. The user name list is displayed in <b>identifiers</b>.</li> <li>• <b>Group</b>: indicates that the mapping is for user groups. The user group name list is displayed in <b>identifiers</b>.</li> </ul>



Parameter	Type	Description
identifiers	Array of String	The list of users or user groups mapped to the IAM agency. Log in to the IAM management console and choose <b>Users</b> or <b>User Groups</b> in the left navigation pane to obtain the user or user group name list.
agency_id	String	The unique ID of the agency bound to the mapping. Log in to the IAM management console and choose <b>Agencies</b> in the left navigation pane. On the <b>Agencies</b> page that is displayed, move your cursor over the agency name to obtain the agency ID.

**Status code: 400**

**Table 6-123** Response parameters

Parameter	Type	Description
error_code	String	The error code.
error_msg	String	The error message.

## Example Request

None

## Example Response

**Status code: 200**

Querying the mapping between a user or user group and an IAM agency is successful.

```
{
  "agency_mappings": [ {
    "agency": "agency01",
    "identifier_type": "User",
    "identifiers": [ "user01" ],
    "agency_id": "092adc623c00d2ea4fdac01d4b637f0b"
  }, {
    "agency": "agency02",
    "identifier_type": "User",
    "identifiers": [ "user02" ],
    "agency_id": "065239307e00d3ae4f80c01d4bdafdfd"
  }, {
    "agency": "groupAgency",
    "identifier_type": "Group",
    "identifiers": [ "group01", "group02", "group03" ],
    "agency_id": "08467a446200d5ac4ff9c01d56670c3b"
  }
]
```

```
}]  
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

# 6.6.2 Updating the Mapping Between a User (Group) and an IAM Agency

## Function

This API is used to update the mapping between a user or user group and an IAM agency.

## Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- URI format  
PUT /v2/{project\_id}/clusters/{cluster\_id}/agency-mapping
- Parameter description

**Table 6-124** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
cluster_id	Yes	String	The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .

## Request Parameters

**Table 6-125** Request parameter

Parameter	Mandatory	Type	Description
agency_mappings	Yes	Array of AgencyMapping objects	The mapping between users or user groups and agencies. For details, see <a href="#">Table 6-126</a> .

**Table 6-126** agency\_mappings parameters

Parameter	Mandatory	Type	Description
agency	Yes	String	The name of the IAM agency bound to the mapping.
identifier_type	Yes	String	The agency type. Possible values: <ul style="list-style-type: none"> <li><b>User</b>: indicates that the mapping is for users. The user name list is displayed in <b>identifiers</b>.</li> <li><b>Group</b>: indicates that the mapping is for user groups. The user group name list is displayed in <b>identifiers</b>.</li> </ul>
identifiers	Yes	Array of String	The list of users or user groups mapped to the IAM agency. Log in to the IAM management console and choose <b>Users</b> or <b>User Groups</b> in the left navigation pane to obtain the user or user group name list.
agency_id	Yes	String	The ID of the agency bound to the mapping. Log in to the IAM management console and choose <b>Agencies</b> in the left navigation pane. On the <b>Agencies</b> page that is displayed, move your cursor over the agency name to obtain the agency ID.

## Response Parameters

Status code: 200

**Table 6-127** Response body parameter

Parameter	Type	Description
result	String	The operation result of the request for updating a mapping. Possible values: <ul style="list-style-type: none"> <li>• <b>succeeded</b>: The operation is successful.</li> <li>• <b>failed</b>: The operation failed.</li> </ul>

**Status code: 400**

**Table 6-128** Response body parameters

Parameter	Type	Description
error_code	String	The error code.
error_msg	String	The error description.

## Example Request

Example request for updating the mapping between a user or user group and an IAM agency

```
PUT https://{endpoint}/v2/{project_id}/clusters/{cluster_id}/agency-mapping
{
  "agency_mappings": [ {
    "agency": "agency01",
    "identifier_type": "User",
    "identifiers": [ "test" ],
    "agency_id": "xxxx"
  } ]
}
```

## Example Response

**Status code: 200**

Updating the mapping between a user or user group and an IAM agency is successful.

```
{
  "result": "succeeded"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.7 Data Connection Management

### 6.7.1 Creating a Data Connection

#### Function

This API is used to create a data connection.

#### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

#### URI

POST /v2/{project\_id}/data-connectors

**Table 6-129** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID

#### Request Parameters

**Table 6-130** Request body parameters

Parameter	Mandatory	Type	Description
data_connect or	Yes	object	Data connection. For details, see <a href="#">Table 6-131</a> .

**Table 6-131** DataConnector

Parameter	Mandatory	Type	Description
connector_name	Yes	String	Data connection name

Parameter	Mandatory	Type	Description
source_type	Yes	String	The data connection type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>RDS_POSTGRES</b>: RDS PostgreSQL database</li> <li>• <b>RDS_MYSQL</b>: RDS MySQL database</li> <li>• <b>gaussdb-mysql</b>: GaussDB(for MySQL)</li> </ul>
source_info	Yes	String	Data source information, in JSON format. The content varies from one data connection to another. For details about the <b>source_info</b> content for each data source, see the example request.

## Response Parameters

Status code: 201

**Table 6-132** Response body parameter

Parameter	Type	Description
connector_id	String	Data connection ID returned after a data connection is created.

## Example Request

Create a data connection.

```
/v2/ac66f1a5-e8f4-4399-8ec6-2c8cb1aefda7/data-connector
{
  "data_connector" : {
    "connector_name" : "mrs-mysql-connector",
    "source_type" : "RDS_MYSQL",
    "source_info" : "{\"db_name\":\"default\",\"user_name\":\"mrs\",\"password\":\"*****\",\"rds_instance_id\":\"3aa55f1cb4c3491686936130f21e9f16in01\"}"
  }
}
```

## Example Response

Status code: 201

Data connection created successfully.

```
{
  "connector_id" : "0822239c1e80d4502f82c008937da9b5"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.7.2 Querying the Data Connection List

### Function

This API is used to query the data connection list.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

GET /v2/{project\_id}/data-connectors

**Table 6-133** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID

**Table 6-134** Query parameters

Parameter	Mandatory	Type	Description
connector_id	No	String	Connection ID
source_type	No	String	Data source type
connector_name	No	String	Data connection name
limit	No	Integer	Number of results returned on each page
offset	No	Integer	Start offset of the pagination query
available	No	Boolean	Whether the data connection is available

## Request Parameters

None

## Response Parameters

Status code: 200

**Table 6-135** Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of data connections
data_connectors	Array of DataConnectorDetail objects	Data connection details list. For details, see <a href="#">Table 6-136</a> .

**Table 6-136** DataConnectorDetail

Parameter	Type	Description
connector_name	String	Data connection name
source_type	String	Data connection type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>RDS_POSTGRES</b>: RDS PostgreSQL database</li> <li>• <b>RDS_MYSQL</b>: RDS MySQL database</li> <li>• <b>gaussdb-mysql</b>: GaussDB(for MySQL)</li> </ul>
source_info	String	Data source information, in JSON format. The content varies from one data connection to another. For details about the <b>source_info</b> content for each data source, see the example request.
connector_id	String	Data connection ID
create_time	Long	Creation time
last_update_time	Long	Last update time
create_by	String	Creator ID
create_user	String	Creator
tenant_id	String	Tenant ID



Parameter	Type	Description
last_update_by	String	ID of the user who last updated the information
status	Integer	Data connection status. The options are as follows: <ul style="list-style-type: none"> <li>• -1: deleted</li> <li>• 0: normal</li> <li>• -2: abnormal</li> <li>• 1: in use</li> </ul>
used_clusters	String	Clusters using the connection
encrypt_type	Integer	Encryption type

## Example Request

None

## Example Response

**Status code: 200**

Successful query of the data connection list

```
{
  "total_count": 1,
  "data_connectors": [ {
    "connector_id": "7d169c6c-ab50-4a56-a5d2-240ca20aabda",
    "connector_name": "test",
    "create_time": 1681270961,
    "source_type": "RDS_MYSQL",
    "source_info": "{\"rds_instance_id\":\"3ce1ae8af8cd43b2bcbd18b7541bb8b0in01\",\"jdbc_url\":\"jdbc:mysql://192.XXX.XXX.XXX:3306/test?socketTimeout=60000\",\"db_name\":\"test\",\"user_name\":\"root\",\"driver_path\":\"mrs-public/drivers/mysql-connector-java-5.1.47.jar\"}",
    "last_update_time": 1681270961,
    "create_by": "148c89d603e048b291f5940935b38f46",
    "create_user": "xxx",
    "tenant_id": "xxxxe662a7e24cc99bfc858c4558dbf6",
    "last_update_by": "148c89d603e048b291f5940935b38f46",
    "status": 0,
    "encrypt_type": 2
  } ]
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.7.3 Updating a Data Connection

### Function

This API is used to update a data connection.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

PUT /v2/{project\_id}/data-connectors/{connector\_id}

**Table 6-137** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID
connector_id	Yes	String	Data connection ID

### Request Parameters

**Table 6-138** Request body parameters

Parameter	Mandatory	Type	Description
data_connector	Yes	DataConnector object	Data connection. For details, see <a href="#">Table 6-139</a> .

**Table 6-139** DataConnector

Parameter	Mandatory	Type	Description
connector_name	Yes	String	Data connection name
source_type	Yes	String	The data connection type. The options are as follows: <ul style="list-style-type: none"><li>• <b>RDS_POSTGRES</b>: RDS PostgreSQL database</li><li>• <b>RDS_MYSQL</b>: RDS MySQL database</li><li>• <b>gaussdb-mysql</b>: GaussDB(for MySQL)</li></ul>

Parameter	Mandatory	Type	Description
source_info	Yes	String	Data source information, in JSON format. The content varies from one data connection to another. For details about the <b>source_info</b> content for each data source, see the example request.

## Response Parameters

Status code: 200

**Table 6-140** Response body parameter

Parameter	Type	Description
result	String	The operation result of the mapping update request. Value <b>succeeded</b> indicates that the operation is successful, and value <b>failed</b> indicates that the operation fails.

## Example Request

```
/v2/274ee662a7e24cc99bfc858c4558dbf6/data-connector/7d169c6c-ab50-4a56-a5d2-240ca20aabda
{
  "data_connector" : {
    "connector_name" : "mrs-mysql-connector",
    "source_type" : "RDS_MYSQL",
    "source_info" : "{\"db_name\":\"default\",\"user_name\":\"mrs\",\"password\":\"*****\",\"rds_instance_id\":\"3aa55f1cb4c3491686936130f21e9f16in01\"}"
  }
}
```

## Example Response

Status code: 200

Data connection updated.

```
{
  "result" : "succeed"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 6.7.4 Deleting a Data Connection

### Function

This API is used to delete a data connection.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

DELETE /v2/{project\_id}/data-connectors/{connector\_id}

**Table 6-141** URI parameters

Parameter	Mandatory	Type	Description
connector_id	Yes	String	Data connection ID
project_id	Yes	String	Project ID

### Request Parameters

None

### Response Parameters

None

### Example Request

None

### Example Response

None

### Status Codes

See [Status Codes](#).

### Error Codes

See [Error Codes](#).

## 6.8 Querying Version Metadata

## 6.8.1 Obtaining MRS Version List

### Function

This API is used to obtain the MRS version list.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

GET /v2/{project\_id}/metadata/versions

**Table 6-142** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID

### Request Parameter

None

### Response Parameter

**Status code: 200**

**Table 6-143** Response body parameter

Parameter	Type	Description
cluster_versions	Array of strings	Cluster version list

### Example Request

None

### Example Response

**Status code: 200**

Successful request

```
{
  "cluster_versions" : [ "MRS xxx", "MRS aaa", "MRS bbb" ]
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.8.2 Querying Available Specifications of an MRS Cluster Version

### Function

This API is used to query available specifications of an MRS cluster version.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

GET /v2/{project\_id}/metadata/version/{version\_name}/available-flavor

**Table 6-144** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .
version_name	Yes	String	MRS cluster version. You can query only one version at a time, for example, MRS 3.1.0. If the request client does not support automatic escape, escape the space to %20, for example, MRS %203.1.0.

### Request Parameters

None

### Response Parameters

**Status code: 200**

**Table 6-145** Response body parameter

Parameter	Type	Description
version_name	String	Version Name
available_flavors	Array of AzFlavors objects	Flavors supported by different AZs. For details, see <a href="#">Table 6-146</a> .

**Table 6-146** AzFlavors

Parameter	Type	Description
az_code	String	AZ code
az_name	String	AZ name
master	Array of Flavor objects	Specifications supported by the master node. For details, see <a href="#">Table 6-147</a> .
core	Array of Flavor objects	Specifications supported by the core node. For details, see <a href="#">Table 6-147</a> .
task	Array of Flavor objects	Specifications supported by the task node. For details, see <a href="#">Table 6-147</a> .

**Table 6-147** Flavor

Parameter	Type	Description
flavor_name	String	Specification name

## Example Request

Obtain details about MRS 3.1.0 metadata.

```
/v2/f77c10d14a544393a24e5f0bf53202b6/metadata/version/MRS%203.1.0/available-flavor
```

## Example Response

**Status code: 200**

Available specifications of the current version

```
{
  "version_name" : "MRS 3.1.0",
  "available_flavors" : [ {
    "az_code" : "cn-north-7a",
    "az_name" : "AZ 1",
    "master" : [ {
      "flavor_name" : "Si3.4xlarge.4"
    }, {
      "flavor_name" : "m3.8xlarge.8"
    }
  ]
}
```

```

    }, {
      "flavor_name": "c6.xlarge.4"
    }, {
      "flavor_name": "m6.xlarge.8"
    }
  ],
  "core": [ {
    "flavor_name": "Si3.4xlarge.4"
  }, {
    "flavor_name": "m3.8xlarge.8"
  }, {
    "flavor_name": "c6.xlarge.4"
  }, {
    "flavor_name": "m6.xlarge.8"
  }
  ],
  "task": [ {
    "flavor_name": "Si3.4xlarge.4"
  }, {
    "flavor_name": "m3.8xlarge.8"
  }, {
    "flavor_name": "c6.xlarge.4"
  }, {
    "flavor_name": "m6.xlarge.8"
  }
  ]
}, {
  "az_code": "cn-north-7b",
  "az_name": "AZ 2",
  "master": [ {
    "flavor_name": "Si3.4xlarge.4"
  }, {
    "flavor_name": "m3.8xlarge.8"
  }, {
    "flavor_name": "c6.xlarge.4"
  }, {
    "flavor_name": "m6.xlarge.8"
  }
  ],
  "core": [ {
    "flavor_name": "Si3.4xlarge.4"
  }, {
    "flavor_name": "m3.8xlarge.8"
  }, {
    "flavor_name": "c6.xlarge.4"
  }, {
    "flavor_name": "m6.xlarge.8"
  }
  ],
  "task": [ {
    "flavor_name": "Si3.4xlarge.4"
  }, {
    "flavor_name": "m3.8xlarge.8"
  }, {
    "flavor_name": "c6.xlarge.4"
  }, {
    "flavor_name": "m6.xlarge.8"
  }
  ]
}, {
  "az_code": "cn-north-7c",
  "az_name": "AZ 3",
  "master": [ {
    "flavor_name": "Si3.4xlarge.4"
  }, {
    "flavor_name": "m3.8xlarge.8"
  }, {
    "flavor_name": "c6.xlarge.4"
  }, {
    "flavor_name": "m6.xlarge.8"
  }
  ],
  "core": [ {
    "flavor_name": "Si3.4xlarge.4"
  }, {
    "flavor_name": "m3.8xlarge.8"

```



```

    }, {
      "flavor_name" : "c6.xlarge.4"
    }, {
      "flavor_name" : "m6.xlarge.8"
    }
  ],
  "task" : [ {
    "flavor_name" : "Si3.4xlarge.4"
  }, {
    "flavor_name" : "m3.8xlarge.8"
  }, {
    "flavor_name" : "c6.xlarge.4"
  }, {
    "flavor_name" : "m6.xlarge.8"
  }
  ]
}

```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

# 6.9 IAM Synchronization

## 6.9.1 Obtaining Synchronized IAM Users and User Groups

### Function

This API is used to obtain synchronized IAM users and user groups.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

GET /v2/{project\_id}/clusters/{cluster\_id}/iam-sync-user

**Table 6-148** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID
cluster_id	Yes	String	Cluster ID

### Request Parameters

None

## Response Parameters

**Status code: 200**

**Table 6-149** Response body parameter

Parameter	Type	Description
user_names	Array of strings	Synchronized users
group_names	Array of strings	Synchronized user groups

## Example Request

None

## Example Response

**Status code: 200**

Synchronized users and user groups

```
{
  "user_names": [ "user1", "user2" ],
  "group_names": [ "group1" ]
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.9.2 Synchronizing an IAM User and User Group

### Function

This API is used to synchronize IAM users and user groups to Manager. If a user is specified, the IAM user group the user belongs to is also synchronized.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/iam-sync-user

**Table 6-150** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID
cluster_id	Yes	String	Cluster ID

## Request Parameters

**Table 6-151** Request body parameters

Parameter	Mandatory	Type	Description
is_all_sync	No	Boolean	Full synchronization. Value <b>true</b> indicates full synchronization, and <b>false</b> indicates synchronization of specified users and user groups. The default value is <b>false</b> .
group_names	No	Array of strings	IAM user groups to be synchronized
user_names	No	Array of strings	IAM users to be synchronized

## Response Parameters

Status code: 202

**Table 6-152** Response body parameter

Parameter	Type	Description
state	String	Request result

## Example Request

Synchronize group1 and user1 to the Manager. The IAM user group associated with user1 is automatically synchronized to Manager.

```
v2/ff8080828997cb24018a1b2db3440b80/clusters/f7f45c04-4303-411c-9b71-d2cb730dd162/iam-sync-user
{
  "is_all_sync" : false,
  "group_names" : [ "groups1" ],
  "user_names" : [ "user1", "user2" ]
}
```

## Example Response

**Status code: 202**

Result the synchronization request.

```
{  
  "state": "synchronizing"  
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.9.3 Cancelling Synchronization of Specified Users and User Groups

### Function

This API is used to cancel synchronization of specified users and user groups.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

DELETE /v2/{project\_id}/clusters/{cluster\_id}/iam-sync-user

**Table 6-153** URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID
cluster_id	Yes	String	Cluster ID

## Request Parameters

**Table 6-154** Request body parameters

Parameter	Mandatory	Type	Description
group_names	No	Array of strings	IAM user group whose synchronization is to be canceled
user_names	No	Array of strings	IAM users whose synchronization is to be canceled

## Response Parameters

**Status code: 202**

**Table 6-155** Response body parameter

Parameter	Type	Description
state	String	Request result

## Example Request

```
/v2/ff8080828997cb24018a1b2db3440b80/clusters/f7f45c04-4303-411c-9b71-d2cb730dd162/iam-sync-user
{
  "group_names": [ "groups1" ],
  "user_names": [ "user1", "user2" ]
}
```

## Example Response

**Status code: 202**

Request result

```
{
  "state": "synchronizing"
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.10 Tag Management APIs

### 6.10.1 Enabling or Disabling the Default Tag of a Cluster

#### Function

This API is used to enable or disable default cluster tags. If this function is enabled, nodes in a cluster are added with the default cluster tag.

#### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

#### URI

POST /v2/{project\_id}/clusters/{cluster\_id}/tags/switch

Table 6-156 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID
cluster_id	Yes	String	Cluster ID

#### Request Parameters

Table 6-157 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Operation, which can be create or delete.

#### Response Parameters

None

#### Example Request

- Disable default tags for a cluster.  

```
/v2/174ee662a7e24cc99bfc858c455xxxxx/clusters/848cd341-cbc8-4ac3-82f2-35cd58c077d9/tags/switch  
{
```

- ```
"action" : "delete"
}
```
- Enable default tags for a cluster.  
`/v2/174ee662a7e24cc99bfc858c455xxxxx/clusters/848cd341-cbc8-4ac3-82f2-35cd58c077d9/tags/switch`

```
{
  "action" : "create"
}
```

## Example Response

**Status code: 400**

Response for a failed request

```
{
  "error_code" : "MRS.00000000",
  "error_msg" : "Insufficient tag quotas."
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.10.2 Querying the Status of Default Cluster Tags

### Function

This API is used to query the status of default cluster tags.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

GET `/v2/{project_id}/clusters/{cluster_id}/tags/status`

**Table 6-158** URI parameters

| Parameter  | Mandatory | Type   | Description |
|------------|-----------|--------|-------------|
| project_id | Yes       | String | Project ID  |
| cluster_id | Yes       | String | Cluster ID  |

## Request Parameters

None

## Response Parameters

**Status code: 200**

**Table 6-159** Response body parameter

| Parameter           | Type    | Description                              |
|---------------------|---------|------------------------------------------|
| status              | String  | Tag status                               |
| default_tags_enable | Boolean | Whether default cluster tags are enabled |

## Example Request

None

## Example Response

**Status code: 200**

Response for a successful request

```
{  
  "status" : "succeed",  
  "default_tags_enable" : "true"  
}
```

**Status code: 400**

Response for a failed request

```
{  
  "error_code" : "MRS.00000000",  
  "error_msg" : "An error occurred when requesting the tag."  
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

## 6.10.3 Querying Tag Quotas

### Function

This API is used to query tag quotas.



## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

## URI

GET /v2/{project\_id}/clusters/{cluster\_id}/tags/quota

**Table 6-160** URI parameters

| Parameter  | Mandatory | Type   | Description |
|------------|-----------|--------|-------------|
| project_id | Yes       | String | Project ID  |
| cluster_id | Yes       | String | Cluster ID  |

**Table 6-161** Query parameters

| Parameter                | Mandatory | Type    | Description                                |
|--------------------------|-----------|---------|--------------------------------------------|
| auto_scaling_policy_tags | No        | Boolean | Whether to query auto scaling policy tags. |

## Request Parameters

None

## Response Parameters

**Status code: 200**

**Table 6-162** Response body parameter

| Parameter       | Type    | Description      |
|-----------------|---------|------------------|
| total_quota     | Integer | Total quotas     |
| available_quota | Integer | Available quotas |

## Example Request

None

## Example Response

**Status code: 200**

Response for a successful request

```
{  
  "total_quota" : 10,  
  "available_quota" : 2  
}
```

**Status code: 400**

Response for a failed request

```
{  
  "error" : {  
    "error_code" : "MRS.00000000",  
    "error_msg" : "The cluster does not exist."  
  }  
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

# 7 API V1.1

## 7.1 Cluster Management APIs

### 7.1.1 Creating a Cluster and Executing a Job

#### Function

This API is used to create an MRS cluster and submit a job in the cluster. This API is incompatible with Sahara.

You are advised to use the V2 APIs for [creating a cluster](#) and [creating a cluster and submitting a job](#).

A maximum of 10 clusters can be concurrently created. You can set the `enterprise_project_id` parameter to perform fine-grained authorization for resources.

Before using the API, you need to obtain the resources listed in [Table 7-1](#).

**Table 7-1** Obtaining resources

| Resource | How to Obtain                                                                                                                                                     |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VPC      | See operation instructions in <a href="#">Querying VPCs</a> and <a href="#">Creating a VPC</a> in the <i>VPC API Reference</i> .                                  |
| Subnet   | See operation instructions in <a href="#">Querying Subnets</a> and <a href="#">Creating a Subnet</a> in the <i>VPC API Reference</i> .                            |
| Key Pair | See operation instructions in <a href="#">Querying SSH Key Pairs</a> and <a href="#">Creating and Importing an SSH Key Pair</a> in the <i>ECS API Reference</i> . |
| Zone     | See <a href="#">Endpoints</a> for details about regions and AZs.                                                                                                  |
| Version  | Currently, MRS 1.9.2, 3.1.0, 3.1.5, 3.1.2-LTS.3, and 3.2.0-LTS.1 are supported.                                                                                   |

| Resource  | How to Obtain                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Component | <ul style="list-style-type: none"> <li>● MRS 3.2.0-LTS.1 supports the following components:               <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and Guardian.</li> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, Flume, and Guardian.</li> <li>- A custom cluster contains the following components: CDL, Hadoop, Spark2x, HBase, Hive, Hue, IoTDB, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse, and Guardian.</li> </ul> </li> <li>● MRS 3.1.5 supports the following components:               <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Sqoop, and Guardian</li> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Sqoop, Guardian, Kafka, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Kafka, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, Kudu, Sqoop, and Guardian.</li> </ul> </li> <li>● MRS 3.1.2-LTS.3 supports the following components:               <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, and Tez.</li> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse.</li> </ul> </li> <li>● MRS 3.1.0 supports the following components:               <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, and Kudu.</li> </ul> </li> </ul> |

| Resource | How to Obtain                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|          | <ul style="list-style-type: none"> <li>- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Kafka, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Kafka, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, and Kudu.</li> <li>• MRS 3.0.5 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, and Alluxio.</li> <li>- A streaming cluster contains the following components: Kafka, Storm, Flume, ZooKeeper, and Ranger.</li> <li>- A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Alluxio, Kafka, Storm, and Flume.</li> <li>- A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Kafka, Storm, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, Kudu, and Alluxio.</li> </ul> </li> <li>• MRS 2.1.0 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Presto, Hadoop, Spark, HBase, Hive, Hue, Loader, Tez, Impala, Kudu, and Flink.</li> <li>- A streaming cluster contains the following components: Kafka, Storm, and Flume.</li> </ul> </li> <li>• MRS 1.9.2 supports the following components: <ul style="list-style-type: none"> <li>- An analysis cluster contains the following components: Presto, Hadoop, Spark, HBase, OpenTSDB, Hive, Hue, Loader, Tez, Flink, Alluxio, and Ranger.</li> <li>- A streaming cluster contains the following components: Kafka, KafkaManager, Storm, and Flume.</li> </ul> </li> </ul> |

## Constraints

- You can log in to a cluster using either a password or a key pair.
- To use the password mode, you need to configure the password of user **root** for accessing the cluster node, that is, **cluster\_master\_secret**.
- To use the key pair mode, you need to configure the key pair name, that is, **node\_public\_cert\_name**.

- Disk parameters can be represented either by **volume\_type** and **volume\_size**, or by multi-disk parameters (**master\_data\_volume\_type**, **master\_data\_volume\_size**, **master\_data\_volume\_count**, **core\_data\_volume\_type**, **core\_data\_volume\_size**, and **core\_data\_volume\_count**).

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

POST /v1.1/{project\_id}/run-job-flow

Table 7-2 URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                  |
|------------|-----------|--------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

## Request Parameters

Table 7-3 Request parameters

| Parameter       | Mandatory | Type   | Description                                                                                                                                                                                      |
|-----------------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cluster_version | Yes       | String | The cluster version.<br>Possible values: <ul style="list-style-type: none"><li>• MRS 1.9.2</li><li>• MRS 3.1.0</li><li>• MRS 3.1.2-LTS.3</li><li>• MRS 3.1.5</li><li>• MRS 3.2.0-LTS.1</li></ul> |
| cluster_name    | Yes       | String | The cluster name, which must be unique. A cluster name can contain only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.                                  |

| Parameter       | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------|-----------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| master_node_num | No        | Integer | The number of master nodes. If cluster HA is enabled, set this parameter to <b>2</b> . If cluster HA is disabled, set this parameter to <b>1</b> . This parameter cannot be set to <b>1</b> in MRS 3.x.                                                                                                                                                                      |
| core_node_num   | No        | Integer | The number of core nodes.<br>Value range: [1, 500]<br>The default maximum number of core nodes is <b>500</b> . If more than 500 core nodes are required, apply for a higher quota.                                                                                                                                                                                           |
| billing_type    | Yes       | Integer | The cluster billing mode. <ul style="list-style-type: none"> <li>● <b>11</b>: The cluster is billed on a yearly/monthly basis.</li> <li>● <b>12</b>: The cluster is billed on a pay-per-use basis.</li> </ul>                                                                                                                                                                |
| data_center     | Yes       | String  | The information about the region where the cluster is located. For details, see <a href="#">Endpoints</a> .                                                                                                                                                                                                                                                                  |
| vpc             | Yes       | String  | The name of the VPC where the subnet is located. Obtain the VPC name by performing the following operations on the VPC management console: <ol style="list-style-type: none"> <li>1. Log in to the management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> </ol> On the <b>Virtual Private Cloud</b> page, obtain the VPC name from the list. |

| Parameter        | Mandatory | Type                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|-----------|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| master_node_size | No        | String                          | The instance specifications of master nodes, for example, <b>c3.4xlarge.2.linux.bigdata</b> . The host specifications supported by MRS are determined by CPU, memory, and disk space. For details about instance specifications, see <a href="#">ECS Specifications Used by MRS</a> and <a href="#">BMS Specifications Used by MRS</a> . You are advised to obtain the specifications supported by the corresponding version in the corresponding region from the cluster creation page on the MRS console. |
| core_node_size   | No        | String                          | The instance specifications of core nodes, for example, <b>c3.4xlarge.2.linux.bigdata</b> . For details about instance specifications, see <a href="#">ECS Specifications Used by MRS</a> and <a href="#">BMS Specifications Used by MRS</a> . You are advised to obtain the specifications supported by the corresponding version in the corresponding region from the cluster creation page on the MRS console.                                                                                           |
| component_list   | Yes       | Array of component_list objects | The list of service components to be installed. For details about the parameters, see <a href="#">Table 7-4</a> .                                                                                                                                                                                                                                                                                                                                                                                           |



| Parameter         | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| available_zone_id | Yes       | String | <p>The AZ ID.</p> <ul style="list-style-type: none"> <li>• CN-Hong Kong AZ1 (ap-southeast-1a):<br/>8902e05a7ee04542a6a73246fdc46b0</li> <li>• CN-Hong Kong AZ2 (ap-southeast-1b):<br/>e9554f5c6fb84eeeb29ab766436b6454</li> <li>• AP-Bangkok AZ1 (ap-southeast-2a):<br/>11d18bfa9d57488b8f96680013667546</li> <li>• AP-Bangkok AZ2 (ap-southeast-2b):<br/>09d6a3cddb643a5aa8837600c9af32c</li> <li>• AP-Singapore AZ1 (ap-southeast-3a):<br/>82cc0d8877374316b669613539efd0d9</li> <li>• AP-Singapore AZ2 (ap-southeast-3b):<br/>77394a8450e147779666771f796e9f03</li> <li>• AP-Singapore AZ3 (ap-southeast-3c):<br/>dba8d1bd3d9146659e2b5a38c09b19a4</li> </ul> |
| vpc_id            | Yes       | String | <p>The ID of the VPC where the subnet is located.</p> <p>Obtain the VPC ID by performing the following operations on the VPC management console:</p> <ol style="list-style-type: none"> <li>1. Log in to the management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> </ol> <p>On the <b>Virtual Private Cloud</b> page, obtain the VPC ID from the list.</p>                                                                                                                                                                                                                                                                      |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| subnet_id | Yes       | String | <p>The subnet ID.</p> <p>Obtain the subnet ID by performing the following operations on the VPC management console:</p> <ol style="list-style-type: none"> <li>1. Log in to the management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> <li>3. Locate the row that contains the target VPC and click the number in the <b>Subnets</b> column to view the subnet information.</li> <li>4. Click the subnet name to obtain the network ID.</li> </ol> <p>At least one of <b>subnet_id</b> and <b>subnet_name</b> must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. <b>subnet_id</b> is recommended.</p> |

| Parameter          | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| subnet_name        | Yes       | String | <p>The subnet name.</p> <p>Obtain the subnet name by performing the following operations on the VPC management console:</p> <ol style="list-style-type: none"> <li>1. Log in to the management console.</li> <li>2. Choose <b>Virtual Private Cloud &gt; My VPCs</b>.</li> <li>3. Locate the row that contains the target VPC and click the number in the <b>Subnets</b> column to obtain the subnet name.</li> </ol> <p>At least one of <b>subnet_id</b> and <b>subnet_name</b> must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. If only <b>subnet_name</b> is configured and subnets with the same name exist in the VPC, the first subnet name in the VPC is used when a cluster is created. <b>subnet_id</b> is recommended.</p> |
| security_groups_id | No        | String | <p>The ID of the security group configured for the cluster.</p> <ul style="list-style-type: none"> <li>• If this parameter is left blank, MRS automatically creates a security group, whose name starts with <b>mrs_{cluster_name}</b>.</li> <li>• If this parameter is not left blank, a fixed security group is used to create a cluster. The transferred ID must be the security group ID owned by the current tenant. The security group must include an inbound rule in which all protocols and all ports are allowed and the source is the IP address of the specified node on the management plane.</li> </ul>                                                                                                                                                                                       |

| Parameter   | Mandatory | Type                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------|-----------|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| add_jobs    | No        | Array of add_jobs objects | Jobs can be submitted when a cluster is created. Currently, only one job can be created. For details about the parameters, see <a href="#">Table 7-5</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| volume_size | No        | Integer                   | <p>The data disk storage space of Master and Core nodes. To increase the data storage capacity, you can add disks when creating a cluster. Select a proper disk storage space based on the following application scenarios:</p> <ul style="list-style-type: none"> <li>• Storage-compute decoupling: Data is stored in the OBS system. Costs of clusters are relatively low but computing performance is poor. The clusters can be deleted at any time. It is recommended when data computing is infrequently performed.</li> <li>• Storage-compute integration: Data is stored in the HDFS system. Costs of clusters are relatively high but computing performance is good. The clusters cannot be deleted in a short term. It is recommended when data computing is frequently performed.</li> </ul> <p>Value range: 100 GB to 32,000 GB. When assigning a value to this parameter, you only need to pass a number, without the unit. This parameter is not recommended. For details, see the description of the <b>volume_type</b> parameter.</p> |

| Parameter                | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| volume_type              | No        | String  | The data disk storage type of master and core nodes. Currently, SATA, SAS, SSD, and GPSSD are supported. Disk parameters can be represented by <b>volume_type</b> and <b>volume_size</b> , or multi-disk parameters. If the <b>volume_type</b> and <b>volume_size</b> parameters coexist with the multi-disk parameters, the system reads the <b>volume_type</b> and <b>volume_size</b> parameters first. You are advised to use the multi-disk parameters. Possible values: <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O</li> <li>• <b>SAS</b>: high I/O</li> <li>• <b>SSD</b>: ultra-high I/O</li> <li>• <b>GPSSD</b>: general-purpose SSD</li> </ul> |
| master_data_volume_type  | No        | String  | This parameter is a multi-disk parameter, indicating the data disk storage type of the master node. Currently, SATA, SAS, SSD, and GPSSD are supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| master_data_volume_size  | No        | Integer | This parameter is a multi-disk parameter, indicating the data disk storage space of master nodes. To increase the data storage capacity, you can add disks when creating a cluster.<br>Value range: 100 GB to 32,000 GB. When assigning a value to this parameter, you only need to pass a number, without the unit.                                                                                                                                                                                                                                                                                                                                                        |
| master_data_volume_count | No        | Integer | This parameter is a multi-disk parameter, indicating the number of data disks of the master nodes. The value can only be <b>1</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| core_data_volume_type    | No        | String  | This parameter is a multi-disk parameter, indicating the data disk storage type of core nodes. Currently, SATA, SAS, SSD, and GPSSD are supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

| Parameter              | Mandatory | Type                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------|-----------|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| core_data_volume_size  | No        | Integer                           | This parameter is a multi-disk parameter, indicating the data disk storage space of core nodes. To increase the data storage capacity, you can add disks when creating a cluster.<br><br>Value range: 100 GB to 32,000 GB. When assigning a value to this parameter, you only need to pass a number, without the unit.                                                                                                                  |
| core_data_volume_count | No        | Integer                           | This parameter is a multi-disk parameter, indicating the number of data disks of the core nodes.<br>Value range: 1 to 20                                                                                                                                                                                                                                                                                                                |
| task_node_groups       | No        | Array of task_node_groups objects | The list of task nodes. For details about the parameters, see <a href="#">Table 7-6</a> .                                                                                                                                                                                                                                                                                                                                               |
| bootstrap_scripts      | No        | Array of Bootstrap Script objects | The Bootstrap action script information. For details about the parameters, see <a href="#">Table 7-8</a> .                                                                                                                                                                                                                                                                                                                              |
| node_public_cert_name  | No        | String                            | The name of a key pair. You can use a key pair to log in to a cluster node. If <b>login_mode</b> is set to <b>1</b> , the request body contains the <b>node_public_cert_name</b> field.                                                                                                                                                                                                                                                 |
| cluster_admin_secret   | No        | String                            | Password of the MRS Manager administrator. <ul style="list-style-type: none"> <li>• Must contain 8 to 26 characters.</li> <li>• Cannot be the username or the username spelled backwards.</li> <li>• Must contain every type of the following: <ul style="list-style-type: none"> <li>- Lowercase letters</li> <li>- Uppercase letters</li> <li>- Numbers</li> <li>- Special characters (!@\$%^_+=+[{ } ; , / ?)</li> </ul> </li> </ul> |

| Parameter             | Mandatory | Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------|-----------|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cluster_master_secret | No        | String               | <p>The password of user <b>root</b> for logging in to a cluster node. If <b>login_mode</b> is set to <b>0</b>, the request body contains the <b>cluster_master_secret</b> field.</p> <p>A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> <li>• Must be 8 to 26 characters long.</li> <li>• Must contain every type of the following: uppercase letters, lowercase letters, numbers, and special characters (!@\$%^&amp;*_+[]:;./?), but must not contain spaces.</li> <li>• Cannot be the username or the username spelled backwards.</li> </ul>                                              |
| safe_mode             | Yes       | Integer              | <p>The running mode of an MRS cluster. Possible values:</p> <ul style="list-style-type: none"> <li>• <b>0</b>: normal cluster. In a normal cluster, Kerberos authentication is disabled, and users can use all functions provided by the cluster.</li> <li>• <b>1</b>: security cluster. In a security cluster, Kerberos authentication is enabled, and common users cannot use the file management and job management functions of an MRS cluster or view cluster resource usage and the job records of Hadoop and Spark. To use more functions, the users must obtain the relevant permissions from the MRS Manager administrator.</li> </ul> |
| tags                  | No        | Array of tag objects | <p>The cluster tags. For details about the parameters, see <a href="#">Table 7-9</a>.</p> <ul style="list-style-type: none"> <li>• A cluster allows a maximum of 10 tags. A tag name (key) must be unique in a cluster.</li> <li>• A tag key or value cannot contain the following special characters: =*&lt;&gt;\, /</li> </ul>                                                                                                                                                                                                                                                                                                                |

| Parameter             | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cluster_type          | No        | Integer | <p>The cluster type. Possible values:</p> <ul style="list-style-type: none"> <li>• <b>0</b>: analysis cluster</li> <li>• <b>1</b>: streaming cluster</li> </ul> <p>The default value is <b>0</b>.</p> <p>Note: Currently, hybrid clusters cannot be created using APIs.</p>                                                                                                                                                                                                   |
| log_collection        | No        | Integer | <p>Whether to collect logs when cluster creation fails.</p> <p>The default value is <b>1</b>, indicating that OBS buckets are created only for collecting logs when an MRS cluster fails to create.</p> <p>Enumerated values:</p> <ul style="list-style-type: none"> <li>• <b>0</b>: Do not collect logs.</li> <li>• <b>1</b>: Collect logs.</li> </ul>                                                                                                                       |
| enterprise_project_id | No        | String  | <p>The enterprise project ID.</p> <p>When you create a cluster, associate the enterprise project ID with the cluster.</p> <p>The default value is <b>0</b>, indicating the <b>default</b> enterprise project.</p> <p>To obtain the enterprise project ID, see the <b>id</b> value in the <b>enterprise_project field data structure</b> table in "Querying the Enterprise Project List" in <i>Enterprise Management API Reference</i>.</p>                                    |
| login_mode            | No        | Integer | <p>Cluster login mode. The default value is <b>1</b>.</p> <ul style="list-style-type: none"> <li>• If <b>login_mode</b> is set to <b>0</b>, the request body contains the <b>cluster_master_secret</b> field.</li> <li>• If <b>login_mode</b> is set to <b>1</b>, the request body contains the <b>node_public_cert_name</b> field.</li> </ul> <p>Enumerated values:</p> <ul style="list-style-type: none"> <li>• <b>0</b>: password</li> <li>• <b>1</b>: key pair</li> </ul> |



| Parameter   | Mandatory | Type                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------|-----------|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_groups | No        | Array of NodeGroupV11 objects | The list of nodes. For details about the parameters, see <a href="#">Table 7-10</a> .<br>Note: Configure either this parameter or the following parameters:<br><b>master_node_num,</b><br><b>master_node_size,</b><br><b>core_node_num, core_node_size,</b><br><b>master_data_volume_type,</b><br><b>master_data_volume_size,</b><br><b>master_data_volume_count,</b><br><b>core_data_volume_type,</b><br><b>core_data_volume_size,</b><br><b>core_data_volume_count,</b><br><b>volume_type, volume_size,</b><br><b>task_node_groups</b> |

**Table 7-4** ComponentAmbV11

| Parameter      | Mandatory | Type   | Description                                                                                   |
|----------------|-----------|--------|-----------------------------------------------------------------------------------------------|
| component_name | Yes       | String | The component name. For details, see the component information in <a href="#">Table 7-1</a> . |

**Table 7-5** AddJobsReqV11

| Parameter | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_type  | Yes       | Integer | <p>Job type code.</p> <ul style="list-style-type: none"> <li>● 1: MapReduce</li> <li>● 2: Spark</li> <li>● 3: Hive Script</li> <li>● 4: HiveQL (not supported currently)</li> <li>● 5: DistCp for importing and exporting data (not supported currently)</li> <li>● 6: Spark Script</li> <li>● 7: Spark SQL for submitting Spark SQL statements (not supported currently)</li> </ul> <p><b>NOTE</b><br/>Spark and Hive jobs can be created only in clusters that where Spark and Hive are installed.</p> |
| job_name  | Yes       | String  | <p>The job name. The value can contain 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> <p><b>NOTE</b><br/>Identical job names are allowed but not recommended.</p>                                                                                                                                                                                                                                                                                           |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| jar_path  | No        | String | <p>The path of the JAR or SQL file for program execution. The value must meet the following requirements:</p> <ul style="list-style-type: none"> <li>The value contains a maximum of 1,023 characters. It cannot contain special characters (; &amp;&gt;,&lt;'\$) and cannot be left blank or all spaces.</li> <li>Files can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> <li>OBS: The path must start with <b>s3a://</b>. Files or programs encrypted by KMS are not supported.</li> <li>HDFS: The path starts with a slash (/).</li> </ul> </li> <li>Spark Script must end with <b>.sql</b> while MapReduce and Spark Jar must end with <b>.jar.sql</b> and <b>jar</b> are case-insensitive.</li> </ul> |
| arguments | No        | String | <p>The 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.</p> <p>The parameter can contain 0 to 150,000 characters, but special characters (; &amp;&gt;'&lt;\$) are not allowed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| input     | No        | String | <p>The data input path.</p> <p>Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> <li>OBS: The path must start with <b>s3a://</b>. Files or programs encrypted by KMS are not supported.</li> <li>HDFS: The path starts with a slash (/).</li> </ul> <p>The value can contain 0 to 1,023 characters, but special characters (; &amp;&gt;'&lt;\$) are not allowed.</p>                                                                                                                                                                                                                                                                                                                             |

| Parameter                   | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| output                      | No        | String  | <p>The data output path.</p> <p>Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> <li>• OBS: The path must start with <b>s3a://</b>.</li> <li>• HDFS: The path starts with a slash (/).</li> </ul> <p>If the specified path does not exist, the system will automatically create it.</p> <p>The value can contain 0 to 1,023 characters, but special characters (; &amp;&gt;'&lt;\$) are not allowed.</p> |
| job_log                     | No        | String  | <p>The path for storing job logs that record job running status.</p> <p>Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> <li>• OBS: The path must start with <b>s3a://</b>.</li> <li>• HDFS: The path starts with a slash (/).</li> </ul> <p>The value can contain 0 to 1,023 characters, but special characters (; &amp;&gt;'&lt;\$) are not allowed.</p>                                               |
| shutdown_cluster            | No        | Boolean | <p>Whether to terminate the cluster after the job execution is complete.</p> <ul style="list-style-type: none"> <li>• <b>true</b>: Terminate the cluster.</li> <li>• <b>false</b>: Do not terminate the cluster.</li> </ul>                                                                                                                                                                                                                                                      |
| file_action                 | No        | String  | <p>The action to be performed on a file. Possible values:</p> <ul style="list-style-type: none"> <li>• <b>import</b></li> <li>• <b>export</b></li> </ul>                                                                                                                                                                                                                                                                                                                         |
| submit_job_once_cluster_run | Yes       | Boolean | <ul style="list-style-type: none"> <li>• <b>true</b>: Submit a job during cluster creation.</li> <li>• <b>false</b>: Submit a job after the cluster is created.</li> </ul> <p>Set this parameter to <b>true</b> in this example.</p>                                                                                                                                                                                                                                             |

| Parameter        | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hql              | No        | String | The HQL script statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| hive_script_path | No        | String | <p>The SQL program path. This parameter is required by Spark Script and Hive Script jobs only, and the value must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• The value contains a maximum of 1,023 characters. It cannot contain special characters (; &amp;&gt;&lt;'\$) and cannot be left blank or all spaces.</li> <li>• Files can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> <li>– OBS: The path must start with <b>s3a://</b>. Files or programs encrypted by KMS are not supported.</li> <li>– HDFS: The path starts with a slash (/).</li> </ul> </li> <li>• Ends with <b>.sql</b>. <b>sql</b> is case-insensitive.</li> </ul> |

**Table 7-6** TaskNodeGroup

| Parameter | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_num  | Yes       | Integer | The number of task nodes. The value ranges from 0 to 500. The maximum number of core and task nodes is 500.                                                                                                                                                                                                                                                                                                                                       |
| node_size | Yes       | String  | <p>The instance specifications of task nodes, for example, <b>c3.4xlarge.2.linux.bigdata</b><b>c6.4xlarge.4.linux.mrs</b>. For details about instance specifications, see <a href="#">ECS Specifications Used by MRS</a> and <a href="#">BMS Specifications Used by MRS</a>.</p> <p>Obtain the instance specifications of the corresponding version in the corresponding region from the cluster creation page of the MRS management console.</p> |

| Parameter           | Mandatory | Type                       | Description                                                                                                                                                                                                                                               |
|---------------------|-----------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| data_volume_type    | Yes       | String                     | The data disk storage type of task nodes. Possible values: <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O</li> <li>• <b>SAS</b>: high I/O</li> <li>• <b>SSD</b>: ultra-high I/O</li> <li>• <b>GPSSD</b>: general-purpose SSD</li> </ul> |
| data_volume_count   | Yes       | Integer                    | The number of data disks of a task node.<br>Value range: 0 to 20                                                                                                                                                                                          |
| data_volume_size    | Yes       | Integer                    | The data disk storage space of a task node.<br>Value range: 100 GB to 32,000 GB<br>You only need to enter a number and do not need to enter the unit GB.                                                                                                  |
| auto_scaling_policy | No        | auto_scaling_policy object | The auto scaling policy. For details, see <a href="#">Table 7-7</a> .                                                                                                                                                                                     |

**Table 7-7** AutoScalingPolicy

| Parameter           | Mandatory | Type                            | Description                                                                                                                                                                                                                                                    |
|---------------------|-----------|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| auto_scaling_enable | Yes       | Boolean                         | Whether to enable the auto scaling policy.                                                                                                                                                                                                                     |
| min_capacity        | Yes       | Integer                         | The minimum number of nodes reserved in the node group.<br>Value range: [0, 500]                                                                                                                                                                               |
| max_capacity        | Yes       | Integer                         | The maximum number of nodes in the node group.<br>Value range: [0, 500]                                                                                                                                                                                        |
| resources_plans     | No        | Array of resources_plan objects | The resource plan list. For details, see <a href="#">Table 7-11</a> . If this parameter is left blank, the resource plan is disabled.<br><br>When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>rules</b> must be configured. |

| Parameter    | Mandatory | Type                          | Description                                                                                                                                                                                                  |
|--------------|-----------|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| exec_scripts | No        | Array of scale_script objects | The list of custom scaling automation scripts. For details, see <a href="#">Table 7-14</a> . If this parameter is left blank, a hook script is disabled.                                                     |
| rules        | No        | Array of rules objects        | The list of auto scaling rules. For details, see <a href="#">Table 7-12</a> .<br>When <b>auto_scaling_enable</b> is set to <b>true</b> , either this parameter or <b>resources_plans</b> must be configured. |

**Table 7-8** BootstrapScript

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                    |
|-----------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | Yes       | String | The name of a bootstrap action script, which must be unique in a cluster.<br>The value can contain only numbers, letters, spaces, hyphens (-), and underscores (_) and cannot start with a space.<br>The value can contain 1 to 64 characters. |

| Parameter     | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|-----------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri           | Yes       | String           | <p>The path of a Bootstrap action script. Set this parameter to an OBS bucket path or a local VM path.</p> <ul style="list-style-type: none"> <li>OBS bucket path: Enter a script path. For example, enter the path of the public sample script provided by MRS. Example: <b>s3a://bootstrap/presto/presto-install.sh</b>. If <b>dualroles</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>dualroles</b>. If <b>worker</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>worker</b>. Based on the Presto usage habit, you are advised to install <b>dualroles</b> on the active master nodes and <b>worker</b> on the Core nodes.</li> <li>Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul> |
| parameters    | No        | String           | The bootstrap action script parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| nodes         | Yes       | Array of strings | <p>The type of a node where the bootstrap action script is executed. The value can be <b>Master</b>, <b>Core</b>, or <b>Task</b>.</p> <p>The node type must be represented in lowercase letters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| active_master | No        | Boolean          | <p>Whether the bootstrap action script runs only on active master nodes.</p> <p>The default value is <b>false</b>, indicating that the bootstrap action script can run on all master nodes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



| Parameter              | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| before_component_start | No        | Boolean | Time when the bootstrap action script is executed. Currently, the following two options are available: <b>Before component start</b> and <b>After component start</b> . The default value is <b>false</b> , indicating that the bootstrap action script is executed after the component is started.                                                                                                                                                                                                                                                                                                   |
| fail_action            | Yes       | String  | Whether to continue executing subsequent scripts and creating a cluster after the Bootstrap action script fails to be executed. <ul style="list-style-type: none"> <li>● <b>continue</b>: Continue to execute subsequent scripts.</li> <li>● <b>errorout</b>: Stop the action.</li> </ul> The default value is <b>errorout</b> , indicating that the action is stopped. <p><b>NOTE</b><br/>You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p> |
| start_time             | No        | Long    | The execution time of one bootstrap action script.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| state                  | No        | String  | The running status of one bootstrap action script. <ul style="list-style-type: none"> <li>● PENDING</li> <li>● IN_PROGRESS</li> <li>● SUCCESS</li> <li>● FAILURE</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Parameter     | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|-----------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| action_stages | No        | Array of strings | <p>Select the time when the bootstrap action script is executed.</p> <ul style="list-style-type: none"> <li>• <b>BEFORE_COMPONENT_FIRST_START</b>: before initial component starts</li> <li>• <b>AFTER_COMPONENT_FIRST_START</b>: after initial component starts</li> <li>• <b>BEFORE_SCALE_IN</b>: before scale-in</li> <li>• <b>AFTER_SCALE_IN</b>: after scale-in</li> <li>• <b>BEFORE_SCALE_OUT</b>: before scale-out</li> <li>• <b>AFTER_SCALE_OUT</b>: after scale-out</li> </ul> |

**Table 7-9** Tag

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| key       | Yes       | String | <p>The tag key.</p> <ul style="list-style-type: none"> <li>• It contains a maximum of 36 Unicode characters and cannot be an empty string.</li> <li>• The tag key cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt; \./).</li> <li>• The tag key of a resource must be unique.</li> </ul> |
| value     | Yes       | String | <p>The value.</p> <ul style="list-style-type: none"> <li>• The value contains a maximum of 43 Unicode characters and can be an empty string.</li> <li>• The tag value cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt; \./).</li> </ul>                                                  |

**Table 7-10** NodeGroupV11

| Parameter        | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| group_name       | Yes       | String  | The node group name. Possible values: <ul style="list-style-type: none"> <li>• <b>master_node_default_group</b></li> <li>• <b>core_node_analysis_group</b></li> <li>• <b>core_node_streaming_group</b></li> <li>• <b>task_node_analysis_group</b></li> <li>• <b>task_node_streaming_group</b></li> </ul>                                                                                                    |
| node_num         | Yes       | Integer | The number of nodes. The value ranges from 0 to 500, and the default value is 0. The total number of core and task nodes cannot exceed 500.                                                                                                                                                                                                                                                                 |
| node_size        | Yes       | String  | The instance specifications of cluster nodes, for example, <b>c3.4xlarge.2.linux.bigdata</b> . For details about instance specifications, see <a href="#">ECS Specifications Used by MRS</a> and <a href="#">BMS Specifications Used by MRS</a> . Obtain the instance specifications of the corresponding version in the corresponding region from the cluster creation page of the MRS management console. |
| root_volume_size | No        | String  | The system disk storage space of a node.                                                                                                                                                                                                                                                                                                                                                                    |
| root_volume_type | No        | String  | The system disk storage type of a node. Possible values: <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O</li> <li>• <b>SAS</b>: high I/O</li> <li>• <b>SSD</b>: ultra-high I/O</li> <li>• <b>GPSSD</b>: general-purpose SSD</li> </ul>                                                                                                                                                     |
| data_volume_type | No        | String  | The data disk storage type of a node. Possible values: <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O</li> <li>• <b>SAS</b>: high I/O</li> <li>• <b>SSD</b>: ultra-high I/O</li> <li>• <b>GPSSD</b>: general-purpose SSD</li> </ul>                                                                                                                                                       |

| Parameter           | Mandatory | Type                       | Description                                                                                                                                                                                                                                     |
|---------------------|-----------|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| data_volume_count   | No        | Integer                    | The number of data disks of a node.<br>Value range: 0 to 20                                                                                                                                                                                     |
| data_volume_size    | No        | Integer                    | The data disk storage space of a node.<br>Value range: 100 GB to 32,000 GB                                                                                                                                                                      |
| auto_scaling_policy | No        | auto_scaling_policy object | The auto scaling rule information. This parameter is available only when <b>group_name</b> is set to <b>task_node_analysis_group</b> or <b>task_node_streaming_group</b> .<br>For details about the parameters, see <a href="#">Table 7-7</a> . |

**Table 7-11** ResourcesPlan

| Parameter    | Mandatory | Type    | Description                                                                                                                                                                                         |
|--------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| period_type  | Yes       | String  | Cycle type of a resource plan. Currently, only the following cycle type is supported:<br><b>daily</b>                                                                                               |
| start_time   | Yes       | String  | The start time of a resource plan. The value is in the format of <b>hour:minute</b> , indicating that the time ranges from 00:00 to 23:59.                                                          |
| end_time     | Yes       | String  | End time of a resource plan. The value is in the same format as that of <b>start_time</b> . The interval between <b>end_time</b> and <b>start_time</b> must be greater than or equal to 30 minutes. |
| min_capacity | Yes       | Integer | Minimum number of the preserved nodes in a node group in a resource plan.<br>Value range: [0, 500]                                                                                                  |
| max_capacity | Yes       | Integer | Maximum number of the preserved nodes in a node group in a resource plan.<br>Value range: [0, 500]                                                                                                  |

**Table 7-12 Rule**

| Parameter          | Mandatory | Type           | Description                                                                                                                                                                                      |
|--------------------|-----------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name               | Yes       | String         | Name of an auto scaling rule.<br>It contains only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.<br>Rule names must be unique in a node group.          |
| description        | No        | String         | Description about an auto scaling rule.<br>It contains a maximum of 1024 characters.                                                                                                             |
| adjustment_type    | Yes       | String         | Auto scaling rule adjustment type.<br>Possible values: <ul style="list-style-type: none"> <li>● <b>scale_out</b>: cluster scale-out</li> <li>● <b>scale_in</b>: cluster scale-in</li> </ul>      |
| cool_down_minutes  | Yes       | Integer        | Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.<br>Value range: 0 to 10,080. One week is equal to 10,080 minutes. |
| scaling_adjustment | Yes       | Integer        | Number of nodes that can be adjusted once.<br>Value range: [1, 100]                                                                                                                              |
| trigger            | Yes       | trigger object | Condition for triggering a rule. For details, see <a href="#">Table 7-13</a> .                                                                                                                   |

**Table 7-13 Trigger**

| Parameter   | Mandatory | Type   | Description                                                                                                                                            |
|-------------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| metric_name | Yes       | String | Metric name.<br>This triggering condition makes a judgment according to the value of the metric.<br>A metric name contains a maximum of 64 characters. |

| Parameter           | Mandatory | Type    | Description                                                                                                                                                                                                                                                    |
|---------------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| metric_value        | Yes       | String  | Metric threshold to trigger a rule<br>The value must be an integer or a number with two decimal places.                                                                                                                                                        |
| comparison_operator | No        | String  | Metric judgment logic operator.<br>Possible values: <ul style="list-style-type: none"> <li>• <b>LT</b>: less than</li> <li>• <b>GT</b>: greater than</li> <li>• <b>LTOE</b>: less than or equal to</li> <li>• <b>GTOE</b>: greater than or equal to</li> </ul> |
| evaluation_periods  | Yes       | Integer | Number of consecutive five-minute periods, during which a metric threshold is reached<br>Value range: 1 to 288                                                                                                                                                 |

**Table 7-14** ScaleScript

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | Yes       | String | Name of a custom automation script. It must be unique in a same cluster.<br><br>The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and cannot start with a space.<br><br>The value can contain 1 to 64 characters.                                                                                                            |
| uri       | Yes       | String | Path of a custom automation script. Set this parameter to an OBS bucket path or a local VM path. <ul style="list-style-type: none"> <li>• OBS bucket path: Enter a script path manually, for example, <b>s3a://XXX/scale.sh</b>.</li> <li>• Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul> |

| Parameter     | Mandatory | Type            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------|-----------|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| parameters    | No        | String          | <p>Parameters of a custom automation script.</p> <ul style="list-style-type: none"> <li>Multiple parameters are separated by space.</li> <li>The following predefined system parameters can be transferred: <ul style="list-style-type: none"> <li><i>\${mrs_scale_node_num}</i>: Number of the nodes to be added or removed</li> <li><i>\${mrs_scale_type}</i>: Scaling type. The value can be <b>scale_out</b> or <b>scale_in</b>.</li> <li><i>\${mrs_scale_node_hostnames}</i>: Host names of the nodes to be added or removed</li> <li><i>\${mrs_scale_node_ips}</i>: IP addresses of the nodes to be added or removed</li> <li><i>\${mrs_scale_rule_name}</i>: Name of the rule that triggers auto scaling</li> </ul> </li> <li>Other user-defined parameters are used in the same way as those of common shell scripts. Parameters are separated by space.</li> </ul> |
| nodes         | Yes       | Array of string | Type of a node where the custom automation script is executed. The node type can be Master, Core, or Task.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| active_master | No        | Boolean         | <p>Whether the custom automation script runs only on the active master node.</p> <p>The default value is <b>false</b>, indicating that the custom automation script can run on all Master nodes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

| Parameter    | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| action_stage | Yes       | String | Time when a script is executed.<br>The following four options are supported: <ul style="list-style-type: none"> <li>• <b>before_scale_out</b>: before scale-out</li> <li>• <b>before_scale_in</b>: before scale-in</li> <li>• <b>after_scale_out</b>: after scale-out</li> <li>• <b>after_scale_in</b>: after scale-in</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| fail_action  | Yes       | String | Whether to continue to execute subsequent scripts and create a cluster after the custom automation script fails to be executed. <ul style="list-style-type: none"> <li>• <b>continue</b>: Continue to execute subsequent scripts.</li> <li>• <b>errorout</b>: Stop the action.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the custom automation script is executed successfully.</li> <li>• The scale-in operation cannot be undone. Therefore, <b>fail_action</b> must be set to <b>continue</b> for the scripts that are executed after scale-in.</li> </ul> |

## Response Parameters

Status code: 200

**Table 7-15** Response body parameters

| Parameter  | Type    | Description                                                                                                                                                    |
|------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cluster_id | String  | Cluster ID, which is returned by the system after the cluster is created.                                                                                      |
| result     | Boolean | Operation result. <ul style="list-style-type: none"> <li>• <b>true</b>: The operation is successful.</li> <li>• <b>false</b>: The operation failed.</li> </ul> |



| Parameter | Type   | Description                         |
|-----------|--------|-------------------------------------|
| msg       | String | System message, which can be empty. |

## Example Request

- Use the **node\_groups** parameter group to create a cluster with the HA function enabled. The cluster version is MRS 3.2.0-LTS.1.

POST https://{endpoint}/v1.1/{project\_id}/run-job-flow

```
{
  "billing_type" : 12,
  "data_center" : "",
  "available_zone_id" : "0e7a368b6c54493e94ad32666b47e23e",
  "cluster_name" : "mrs_HEbK",
  "cluster_version" : "MRS 3.2.0-LTS.1",
  "safe_mode" : 0,
  "cluster_type" : 0,
  "component_list" : [ {
    "component_name" : "Hadoop"
  }, {
    "component_name" : "Spark2x"
  }, {
    "component_name" : "HBase"
  }, {
    "component_name" : "Hive"
  }, {
    "component_name" : "Zookeeper"
  }, {
    "component_name" : "Tez"
  }, {
    "component_name" : "Hue"
  }, {
    "component_name" : "Loader"
  }, {
    "component_name" : "Flink"
  } ],
  "vpc" : "vpc-4b1c",
  "vpc_id" : "4a365717-67be-4f33-80c5-98e98a813af8",
  "subnet_id" : "67984709-e15e-4e86-9886-d76712d4e00a",
  "subnet_name" : "subnet-4b44",
  "security_groups_id" : "4820eace-66ad-4f2c-8d46-cf340e3029dd",
  "enterprise_project_id" : "0",
  "tags" : [ {
    "key" : "key1",
    "value" : "value1"
  }, {
    "key" : "key2",
    "value" : "value2"
  } ],
  "node_groups" : [ {
    "group_name" : "master_node_default_group",
    "node_num" : 2,
    "node_size" : "s3.xlarge.2.linux.bigdata",
    "root_volume_size" : 480,
    "root_volume_type" : "SATA",
    "data_volume_type" : "SATA",
    "data_volume_count" : 1,
    "data_volume_size" : 600
  }, {
    "group_name" : "core_node_analysis_group",
    "node_num" : 3,
    "node_size" : "s3.xlarge.2.linux.bigdata",
    "root_volume_size" : 480,
    "root_volume_type" : "SATA",
```

```
"data_volume_type" : "SATA",
"data_volume_count" : 1,
"data_volume_size" : 600
}, {
"group_name" : "task_node_analysis_group",
"node_num" : 2,
"node_size" : "s3.xlarge.2.linux.bigdata",
"root_volume_size" : 480,
"root_volume_type" : "SATA",
"data_volume_type" : "SATA",
"data_volume_count" : 0,
"data_volume_size" : 600,
"auto_scaling_policy" : {
"auto_scaling_enable" : true,
"min_capacity" : 1,
"max_capacity" : "3",
"resources_plans" : [ {
"period_type" : "daily",
"start_time" : "9:50",
"end_time" : "10:20",
"min_capacity" : 2,
"max_capacity" : 3
}, {
"period_type" : "daily",
"start_time" : "10:20",
"end_time" : "12:30",
"min_capacity" : 0,
"max_capacity" : 2
}
]
},
"exec_scripts" : [ {
"name" : "before_scale_out",
"uri" : "s3a://XXX/zeppelin_install.sh",
"parameters" : "${mrs_scale_node_num} ${mrs_scale_type} xxx",
"nodes" : [ "master", "core", "task" ],
"active_master" : "true",
"action_stage" : "before_scale_out",
"fail_action" : "continue"
}, {
"name" : "after_scale_out",
"uri" : "s3a://XXX/storm_rebalance.sh",
"parameters" : "${mrs_scale_node_hostnames} ${mrs_scale_node_ips}",
"nodes" : [ "master", "core", "task" ],
"active_master" : "true",
"action_stage" : "after_scale_out",
"fail_action" : "continue"
}
]
},
"rules" : [ {
"name" : "default-expand-1",
"adjustment_type" : "scale_out",
"cool_down_minutes" : 5,
"scaling_adjustment" : 1,
"trigger" : {
"metric_name" : "YARNMemoryAvailablePercentage",
"metric_value" : "25",
"comparison_operator" : "LT",
"evaluation_periods" : 10
}
}
], {
"name" : "default-shrink-1",
"adjustment_type" : "scale_in",
"cool_down_minutes" : 5,
"scaling_adjustment" : 1,
"trigger" : {
"metric_name" : "YARNMemoryAvailablePercentage",
"metric_value" : "70",
"comparison_operator" : "GT",
"evaluation_periods" : 10
}
}
}
]
```

```

    }
  ],
  "login_mode" : 1,
  "cluster_master_secret" : "",
  "cluster_admin_secret" : "",
  "log_collection" : 1,
  "add_jobs" : [ {
    "job_type" : 1,
    "job_name" : "tenji111",
    "jar_path" : "s3a://bigdata/program/hadoop-mapreduce-examples-2.7.2.jar",
    "arguments" : "wordcount",
    "input" : "s3a://bigdata/input/wd_1k/",
    "output" : "s3a://bigdata/output/",
    "job_log" : "s3a://bigdata/log/",
    "shutdown_cluster" : true,
    "file_action" : "",
    "submit_job_once_cluster_run" : true,
    "hql" : "",
    "hive_script_path" : ""
  } ],
  "bootstrap_scripts" : [ {
    "name" : "Modify os config",
    "uri" : "s3a://XXX/modify_os_config.sh",
    "parameters" : "param1 param2",
    "nodes" : [ "master", "core", "task" ],
    "active_master" : "false",
    "before_component_start" : "true",
    "start_time" : "1667892101",
    "state" : "IN_PROGRESS",
    "fail_action" : "continue",
    "action_stages" : [ "BEFORE_COMPONENT_FIRST_START", "BEFORE_SCALE_IN" ]
  }, {
    "name" : "Install zeppelin",
    "uri" : "s3a://XXX/zeppelin_install.sh",
    "parameters" : "",
    "nodes" : [ "master" ],
    "active_master" : "true",
    "before_component_start" : "false",
    "start_time" : "1667892101",
    "state" : "IN_PROGRESS",
    "fail_action" : "continue",
    "action_stages" : [ "AFTER_SCALE_IN", "AFTER_SCALE_OUT" ]
  } ]
}

```

- Create a cluster with the HA function enabled without using the **node\_groups** parameter group. The cluster version is MRS 3.2.0-LTS.1.

```

POST https://{endpoint}/v1.1/{project_id}/run-job-flow

{
  "billing_type" : 12,
  "data_center" : "",
  "master_node_num" : 2,
  "master_node_size" : "s3.2xlarge.2.linux.bigdata",
  "core_node_num" : 3,
  "core_node_size" : "s3.2xlarge.2.linux.bigdata",
  "available_zone_id" : "0e7a368b6c54493e94ad32666b47e23e",
  "cluster_name" : "newcluster",
  "vpc" : "vpc1",
  "vpc_id" : "5b7db34d-3534-4a6e-ac94-023cd36aaf74",
  "subnet_id" : "815bece0-fd22-4b65-8a6e-15788c99ee43",
  "subnet_name" : "subnet",
  "security_groups_id" : "845bece1-fd22-4b45-7a6e-14338c99ee43",
  "tags" : [ {
    "key" : "key1",
    "value" : "value1"
  }, {
    "key" : "key2",
    "value" : "value2"
  } ],
}

```

```
"cluster_version" : "MRS 3.2.0-LTS.1",
"cluster_type" : 0,
"master_data_volume_type" : "SATA",
"master_data_volume_size" : 600,
"master_data_volume_count" : 1,
"core_data_volume_type" : "SATA",
"core_data_volume_size" : 600,
"core_data_volume_count" : 2,
"node_public_cert_name" : "SSHkey-bba1",
"safe_mode" : 0,
"log_collection" : 1,
"task_node_groups" : [ {
  "node_num" : 2,
  "node_size" : "s3.xlarge.2.linux.bigdata",
  "data_volume_type" : "SATA",
  "data_volume_count" : 1,
  "data_volume_size" : 600,
  "auto_scaling_policy" : {
    "auto_scaling_enable" : true,
    "min_capacity" : 1,
    "max_capacity" : "3",
    "resources_plans" : [ {
      "period_type" : "daily",
      "start_time" : "9: 50",
      "end_time" : "10: 20",
      "min_capacity" : 2,
      "max_capacity" : 3
    }, {
      "period_type" : "daily",
      "start_time" : "10: 20",
      "end_time" : "12: 30",
      "min_capacity" : 0,
      "max_capacity" : 2
    }
  ]
}, {
  "exec_scripts" : [ {
    "name" : "before_scale_out",
    "uri" : "s3a://XXX/zeppelin_install.sh",
    "parameters" : "${mrs_scale_node_num}${mrs_scale_type}xxx",
    "nodes" : [ "master", "core", "task" ],
    "active_master" : "true",
    "action_stage" : "before_scale_out",
    "fail_action" : "continue"
  }, {
    "name" : "after_scale_out",
    "uri" : "s3a://XXX/storm_rebalance.sh",
    "parameters" : "${mrs_scale_node_hostnames}${mrs_scale_node_ips}",
    "nodes" : [ "master", "core", "task" ],
    "active_master" : "true",
    "action_stage" : "after_scale_out",
    "fail_action" : "continue"
  }
],
  "rules" : [ {
    "name" : "default-expand-1",
    "adjustment_type" : "scale_out",
    "cool_down_minutes" : 5,
    "scaling_adjustment" : 1,
    "trigger" : {
      "metric_name" : "YARNMemoryAvailablePercentage",
      "metric_value" : "25",
      "comparison_operator" : "LT",
      "evaluation_periods" : 10
    }
  }, {
    "name" : "default-shrink-1",
    "adjustment_type" : "scale_in",
    "cool_down_minutes" : 5,
    "scaling_adjustment" : 1,
    "trigger" : {
      "metric_name" : "YARNMemoryAvailablePercentage",
```

```

        "metric_value" : "70",
        "comparison_operator" : "GT",
        "evaluation_periods" : 10
    }
}
}],
"component_list" : [ {
    "component_name" : "Hadoop"
}, {
    "component_name" : "Spark"
}, {
    "component_name" : "HBase"
}, {
    "component_name" : "Hive"
} ],
"add_jobs" : [ {
    "job_type" : 1,
    "job_name" : "tenji111",
    "jar_path" : "s3a://bigdata/program/hadoop-mapreduce-examples-2.7.2.jar",
    "arguments" : "wordcount",
    "input" : "s3a://bigdata/input/wd_1k/",
    "output" : "s3a://bigdata/output/",
    "job_log" : "s3a://bigdata/log/",
    "shutdown_cluster" : true,
    "file_action" : "",
    "submit_job_once_cluster_run" : true,
    "hql" : "",
    "hive_script_path" : ""
} ],
"bootstrap_scripts" : [ {
    "name" : "Modifyosconfig",
    "uri" : "s3a://XXX/modify_os_config.sh",
    "parameters" : "param1param2",
    "nodes" : [ "master", "core", "task" ],
    "active_master" : "false",
    "before_component_start" : "true",
    "start_time" : "1667892101",
    "state" : "IN_PROGRESS",
    "fail_action" : "continue",
    "action_stages" : [ "BEFORE_COMPONENT_FIRST_START", "BEFORE_SCALE_IN" ]
}, {
    "name" : "Installzeppelin",
    "uri" : "s3a://XXX/zeppelin_install.sh",
    "parameters" : "",
    "nodes" : [ "master" ],
    "active_master" : "true",
    "before_component_start" : "false",
    "start_time" : "1667892101",
    "state" : "IN_PROGRESS",
    "fail_action" : "continue",
    "action_stages" : [ "AFTER_SCALE_IN", "AFTER_SCALE_OUT" ]
} ]
}

```

- Use the **node\_groups** parameter group to create a cluster with the HA function disabled. The cluster version is MRS 3.2.0-LTS.1.

```

POST https://{endpoint}/v1.1/{project_id}/run-job-flow

{
  "billing_type" : 12,
  "data_center" : "",
  "available_zone_id" : "0e7a368b6c54493e94ad32666b47e23e",
  "cluster_name" : "mrs_HEbK",
  "cluster_version" : "MRS 3.2.0-LTS.1",
  "safe_mode" : 0,
  "cluster_type" : 0,
  "component_list" : [ {
    "component_name" : "Hadoop"
  }, {

```

```
"component_name" : "Spark2x"
}, {
  "component_name" : "HBase"
}, {
  "component_name" : "Hive"
}, {
  "component_name" : "Zookeeper"
}, {
  "component_name" : "Tez"
}, {
  "component_name" : "Hue"
}, {
  "component_name" : "Loader"
}, {
  "component_name" : "Flink"
} ],
"vpc" : "vpc-4b1c",
"vpc_id" : "4a365717-67be-4f33-80c5-98e98a813af8",
"subnet_id" : "67984709-e15e-4e86-9886-d76712d4e00a",
"subnet_name" : "subnet-4b44",
"security_groups_id" : "4820eace-66ad-4f2c-8d46-cf340e3029dd",
"enterprise_project_id" : "0",
"tags" : [ {
  "key" : "key1",
  "value" : "value1"
}, {
  "key" : "key2",
  "value" : "value2"
} ],
"node_groups" : [ {
  "group_name" : "master_node_default_group",
  "node_num" : 1,
  "node_size" : "s3.xlarge.2.linux.bigdata",
  "root_volume_size" : 480,
  "root_volume_type" : "SATA",
  "data_volume_type" : "SATA",
  "data_volume_count" : 1,
  "data_volume_size" : 600
}, {
  "group_name" : "core_node_analysis_group",
  "node_num" : 1,
  "node_size" : "s3.xlarge.2.linux.bigdata",
  "root_volume_size" : 480,
  "root_volume_type" : "SATA",
  "data_volume_type" : "SATA",
  "data_volume_count" : 1,
  "data_volume_size" : 600
} ],
"login_mode" : 1,
"cluster_master_secret" : "",
"cluster_admin_secret" : "",
"log_collection" : 1,
"add_jobs" : [ {
  "job_type" : 1,
  "job_name" : "tenji111",
  "jar_path" : "s3a://bigdata/program/hadoop-mapreduce-examples-2.7.2.jar",
  "arguments" : "wordcount",
  "input" : "s3a://bigdata/input/wd_1k/",
  "output" : "s3a://bigdata/output/",
  "job_log" : "s3a://bigdata/log/",
  "shutdown_cluster" : true,
  "file_action" : "",
  "submit_job_once_cluster_run" : true,
  "hql" : "",
  "hive_script_path" : ""
} ],
"bootstrap_scripts" : [ {
  "name" : "Modify os config",
  "uri" : "s3a://XXX/modify_os_config.sh",
```

```

"parameters" : "param1 param2",
"nodes" : [ "master", "core", "task" ],
"active_master" : "false",
"before_component_start" : "true",
"start_time" : "1667892101",
"state" : "IN_PROGRESS",
"fail_action" : "continue",
"action_stages" : [ "BEFORE_COMPONENT_FIRST_START", "BEFORE_SCALE_IN" ]
}, {
"name" : "Install zeppelin",
"uri" : "s3a://XXX/zeppelin_install.sh",
"parameters" : "",
"nodes" : [ "master" ],
"active_master" : "true",
"before_component_start" : "false",
"start_time" : "1667892101",
"state" : "IN_PROGRESS",
"fail_action" : "continue",
"action_stages" : [ "AFTER_SCALE_IN", "AFTER_SCALE_OUT" ]
}]
}

```

- Create a cluster with the HA function disabled without using the **node\_groups** parameter group. The cluster version is MRS 3.2.0-LTS.1.  
POST [https://endpoint/v1.1/{project\\_id}/run-job-flow](https://endpoint/v1.1/{project_id}/run-job-flow)

```

{
  "billing_type" : 12,
  "data_center" : "",
  "master_node_num" : 1,
  "master_node_size" : "s3.2xlarge.2.linux.bigdata",
  "core_node_num" : 1,
  "core_node_size" : "s3.2xlarge.2.linux.bigdata",
  "available_zone_id" : "0e7a368b6c54493e94ad32666b47e23e",
  "cluster_name" : "newcluster",
  "vpc" : "vpc1",
  "vpc_id" : "5b7db34d-3534-4a6e-ac94-023cd36aaf74",
  "subnet_id" : "815bece0-fd22-4b65-8a6e-15788c99ee43",
  "subnet_name" : "subnet",
  "security_groups_id" : "",
  "enterprise_project_id" : "0",
  "tags" : [ {
    "key" : "key1",
    "value" : "value1"
  }, {
    "key" : "key2",
    "value" : "value2"
  } ],
  "cluster_version" : "MRS 3.2.0-LTS.1",
  "cluster_type" : 0,
  "master_data_volume_type" : "SATA",
  "master_data_volume_size" : 600,
  "master_data_volume_count" : 1,
  "core_data_volume_type" : "SATA",
  "core_data_volume_size" : 600,
  "core_data_volume_count" : 1,
  "login_mode" : 1,
  "node_public_cert_name" : "SSHkey-bba1",
  "safe_mode" : 0,
  "cluster_admin_secret" : "*****",
  "log_collection" : 1,
  "component_list" : [ {
    "component_name" : "Hadoop"
  }, {
    "component_name" : "Spark2x"
  }, {
    "component_name" : "HBase"
  }, {
    "component_name" : "Hive"
  }, {

```

```

"component_name" : "Zookeeper"
}, {
"component_name" : "Tez"
}, {
"component_name" : "Hue"
}, {
"component_name" : "Loader"
}, {
"component_name" : "Flink"
}],
"add_jobs" : [ {
"job_type" : 1,
"job_name" : "tenji111",
"jar_path" : "s3a://bigdata/program/hadoop-mapreduce-examples-XXX.jar",
"arguments" : "wordcount",
"input" : "s3a://bigdata/input/wd_1k/",
"output" : "s3a://bigdata/output/",
"job_log" : "s3a://bigdata/log/",
"shutdown_cluster" : false,
"file_action" : "",
"submit_job_once_cluster_run" : true,
"hql" : "",
"hive_script_path" : ""
}],
"bootstrap_scripts" : [ {
"name" : "Install zeppelin",
"uri" : "s3a://XXX/zeppelin_install.sh",
"parameters" : "",
"nodes" : [ "master" ],
"active_master" : "false",
"before_component_start" : "false",
"start_time" : "1667892101",
"state" : "IN_PROGRESS",
"fail_action" : "continue",
"action_stages" : [ "AFTER_SCALE_IN", "AFTER_SCALE_OUT" ]
} ]
}

```

## Example Response

**Status code: 200**

The cluster is created.

```

{
"cluster_id" : "da1592c2-bb7e-468d-9ac9-83246e95447a",
"result" : true,
"msg" : ""
}

```

## Status Codes

[Table 7-16](#) describes the status code.

**Table 7-16** Status code

| Status Code | Description                   |
|-------------|-------------------------------|
| 200         | The cluster has been created. |

See [Status Codes](#).



## Error Codes

See [Error Codes](#).

## 7.1.2 Resizing a Cluster

### Function

This API is used to scale out or scale in Core or Task nodes in a cluster that has been created. After an MRS cluster is created, the number of Master nodes cannot be adjusted. That is, Master nodes cannot be scaled in or out. This API is incompatible with Sahara.

Only clusters in the **Running** state can be scaled out or in.

The APIs described in this section support only streaming, analysis, and hybrid clusters.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
PUT /v1.1/{project\_id}/cluster\_infos/{cluster\_id}
- Parameter description

**Table 7-17** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                             |
|------------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .                |
| cluster_id | Yes       | String | Cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining the MRS Cluster Information</a> . |

## Request Parameters

**Table 7-18** Request parameters

| Parameter       | Mandatory | Type               | Description                                                                                        |
|-----------------|-----------|--------------------|----------------------------------------------------------------------------------------------------|
| service_id      | No        | String             | Service ID. This parameter is reserved for extension. You do not need to set this parameter.       |
| plan_id         | No        | String             | Plan ID. This parameter is reserved for extension. You do not need to set this parameter.          |
| parameters      | Yes       | Object             | Core parameters. For details, see <a href="#">Table 7-19</a> .                                     |
| previous_values | No        | Map<String,String> | This parameter is an extended API and needs to be reserved. You do not need to set this parameter. |

**Table 7-19** parameters description

| Parameter  | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| order_id   | No        | String | Order ID obtained by the system during scale-out or scale-in. You do not need to set the parameter.                                                                                                                                                                                                                                                                                                           |
| scale_type | Yes       | String | <ul style="list-style-type: none"> <li><b>scale_in</b>: cluster scale-in</li> <li><b>scale_out</b>: cluster scale-out</li> </ul>                                                                                                                                                                                                                                                                              |
| node_id    | Yes       | String | ID of the newly added or removed node. The parameter value is fixed to <b>node_orderadd</b> .                                                                                                                                                                                                                                                                                                                 |
| node_group | No        | String | <p>Node group to be scaled out or in</p> <ul style="list-style-type: none"> <li>If the value of <b>node_group</b> is <b>core_node_default_group</b>, the node group is a Core node group.</li> <li>If the value of <b>node_group</b> is <b>task_node_default_group</b>, the node group is a Task node group.</li> </ul> <p>If it is left blank, the default value <b>core_node_default_group</b> is used.</p> |

| Parameter      | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|-----------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| task_node_info | No        | Object  | <p>Task node specifications. For more parameter description, see <a href="#">Table 7-20</a>.</p> <ul style="list-style-type: none"> <li>When the number of Task nodes is 0, this parameter is used to specify Task node specifications.</li> <li>When the number of Task nodes is greater than 0, this parameter is unavailable.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| instances      | Yes       | Integer | <p>Number of nodes to be added or removed</p> <ul style="list-style-type: none"> <li>The maximum number of nodes to be added is 500 minus the number of Core and Task nodes. For example, the current number of Core nodes is 3, the number of nodes to be added must be less than or equal to 497.<br/>A maximum of 500 Core and Task nodes are supported by default. If more than 500 Core and Task nodes are required, contact technical support engineers or call a background API to modify the database.</li> <li>Nodes can be deleted for cluster scale-out when the number of Core nodes is greater than 3 or the number of Task nodes is greater than 0. For example, if there are 5 Core nodes and 5 Task nodes in a cluster, only 2 (5 minus 3) Core nodes are available for deletion and 5 or fewer than 5 Task nodes can be deleted.</li> </ul> |

| Parameter              | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------|-----------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| skip_bootstrap_scripts | No        | String           | This parameter is valid only when a bootstrap action is configured during cluster creation and takes effect during scale-out. It indicates whether the bootstrap action specified during cluster creation is performed on nodes added during scale-out. The default value is <b>false</b> , indicating that the bootstrap action is performed.                                                                                                                                                                         |
| scale_without_start    | No        | Boolean          | Whether to start components on the added nodes after cluster scale-out <ul style="list-style-type: none"> <li>• <b>true</b>: Do not start components after scale-out.</li> <li>• <b>false</b>: Start components after scale-out.</li> </ul>                                                                                                                                                                                                                                                                            |
| server_ids             | No        | Array of strings | ID list of Task nodes to be deleted during task node scale-in. <ul style="list-style-type: none"> <li>• This parameter does not take effect when <b>scale_type</b> is set to <b>scale-out</b>.</li> <li>• If <b>scale_type</b> is set to <b>scale-in</b> and cannot be left blank, the system deletes the specified Task nodes.</li> <li>• When <b>scale_type</b> is set to <b>scale-in</b> and <b>server_ids</b> is left blank, the system automatically deletes the Task nodes based on the system rules.</li> </ul> |

**Table 7-20** task\_node\_info parameters

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                          |
|-----------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_size | Yes       | String | Instance specifications of a Task node, Example:<br>c3.4xlarge.2.linux.bigdata<br>For details about instance specifications, see <a href="#">ECS Specifications Used by MRS</a> and <a href="#">BMS Specifications Used by MRS</a> . |

| Parameter         | Mandatory | Type    | Description                                                                                                                                                                                                                                                                      |
|-------------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| data_volume_type  | Yes       | String  | Data disk storage type of the Task node, supporting SATA, SAS, and SSD currently. <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O</li> <li>• <b>SAS</b>: high I/O</li> <li>• <b>SSD</b>: ultra-high I/O</li> <li>• <b>GPSSD</b>: general-purpose SSD</li> </ul> |
| data_volume_count | Yes       | Integer | Number of data disks of a Task node<br>Value range: 0 to 10                                                                                                                                                                                                                      |
| data_volume_size  | Yes       | Integer | Data disk storage space of a Task node<br>Value range: 100 GB to 32,000 GB.<br>When assigning a value to this parameter, you only need to pass a number, without the unit.                                                                                                       |

## Response Parameters

[Table 7-21](#) describes the response parameter.

**Table 7-21** Response parameter

| Parameter | Type   | Description                                                                                                                                                                                                   |
|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| result    | String | Operation result <ul style="list-style-type: none"> <li>• <b>succeeded</b>: The operation is successful.</li> <li>• <b>Error Codes</b> describes the error codes returned upon operation failures.</li> </ul> |

## Example Request

- Add one core node to the MRS cluster.

```
PUT /v1.1/{project_id}/cluster_infos/{cluster_id}
```

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_out",
    "node_id": "node_orderadd",
    "node_group": "core_node_default_group",
    "instances": "1",
```

```
    "skip_bootstrap_scripts":false,
    "scale_without_start":false
  },
  "previous_values": { }
}
```

- If the number of Task nodes is greater than 0, add one Task node to the MRS cluster.

PUT /v1.1/{project\_id}/cluster\_infos/{cluster\_id}

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_out",
    "node_id": "node_orderadd",
    "node_group": "task_node_default_group",
    "instances": "1",
    "skip_bootstrap_scripts":false,
    "scale_without_start":false
  },
  "previous_values": { }
}
```

- If the number of Task nodes is greater than 0, add a Task node of the s3.xlarge.2.linux.bigdata specification to the MRS cluster.

PUT /v1.1/{project\_id}/cluster\_infos/{cluster\_id}

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_out",
    "node_id": "node_orderadd",
    "node_group": "task_node_default_group",
    "task_node_info": {
      "node_size": "s3.xlarge.2.linux.bigdata",
      "data_volume_type":"SATA",
      "data_volume_count":2,
      "data_volume_size":600
    },
    "instances": "1",
    "scale_without_start":false
  },
  "previous_values": { }
}
```

- Remove a core node from the MRS cluster.

PUT /v1.1/{project\_id}/cluster\_infos/{cluster\_id}

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_in",
    "node_id": "node_orderadd",
    "node_group": "core_node_default_group",
    "instances": "1"
  },
  "previous_values": { }
}
```

- Remove a Task node from the MRS cluster.

```
PUT /v1.1/{project_id}/cluster_infos/{cluster_id}
```

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_in",
    "node_id": "node_orderadd",
    "node_group": "task_node_default_group",
    "instances": "1"
  },
  "previous_values": { }
}
```

- Remove a specified Task node from the MRS cluster.

```
PUT /v1.1/{project_id}/cluster_infos/{cluster_id}
```

```
{
  "service_id": "",
  "plan_id": "",
  "parameters": {
    "order_id": "",
    "scale_type": "scale_in",
    "node_id": "node_orderadd",
    "node_group": "task_node_default_group",
    "instances": "2",
    "server_ids": ["c9573435-7814-4b2c-9131-ad78b814414c",
"a4951009-6a0f-4e7b-9c81-9d4bd1f8c537"]
  },
  "previous_values": { }
}
```

## Example Response

**Status code: 200**

The core or task node has been added or removed.

```
{
  "result": "succeeded"
}
```

## Status Codes

- [Table 7-22](#) describes the status code.

**Table 7-22** Status code

| Status Code | Description                                        |
|-------------|----------------------------------------------------|
| 200         | The core or task nodes have been added or removed. |

- See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 7.1.3 Querying a Cluster List

### Function

This API is used to query a list of clusters created by a user. This API is incompatible with Sahara.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
GET /v1.1/{project\_id}/cluster\_infos
- Parameter description

**Table 7-23** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                  |
|------------|-----------|--------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

**Table 7-24** Query parameters

| Parameter   | Mandatory | Type   | Description                                                                                                            |
|-------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------|
| pageSize    | No        | String | Maximum number of clusters displayed on a page<br>The default value is <b>10</b> . The value range is [1, 2147483646]. |
| currentPage | No        | String | Current page number. The default value is <b>1</b> .                                                                   |
| clusterName | No        | String | The cluster name.                                                                                                      |



| Parameter    | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| clusterState | No        | String | <p>Query clusters by status.</p> <ul style="list-style-type: none"> <li>● <b>existing</b>: Query existing clusters, including all clusters except those in the <b>Deleted</b> state and the yearly/monthly clusters in the <b>Order processing</b> or <b>Preparing</b> state.</li> <li>● <b>history</b>: Quer historical clusters, including all the deleted clusters, clusters that fail to delete, clusters whose VMs fail to delete, and clusters whose database updates fail to delete.</li> <li>● <b>starting</b>: Query a list of clusters that are being started.</li> <li>● <b>running</b>: Query a list of running clusters.</li> <li>● <b>terminated</b>: Query a list of terminated clusters.</li> <li>● <b>failed</b>: Query a list of failed clusters.</li> <li>● <b>abnormal</b>: Query a list of abnormal clusters.</li> <li>● <b>terminating</b>: Query a list of clusters that are being terminated.</li> <li>● <b>frozen</b>: Query a list of frozen clusters.</li> <li>● <b>scaling-out</b>: Query a list of clusters that are being scaled out.</li> <li>● <b>scaling-in</b>: Query a list of clusters that are being scaled in.</li> </ul> |
| tags         | No        | String | <p>You can search for a cluster by its tag. If you specify multiple tags, the relationship between them is AND.</p> <ul style="list-style-type: none"> <li>● The format of the <b>tags</b> parameter is <b>tags=k1*v1,k2*v2,k3*v3</b>.</li> <li>● When the values of some tags are null, the format is <b>tags=k1,k2,k3*v3</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Parameter           | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| enterpriseProjectId | No        | String | The enterprise project ID used to query clusters in a specified enterprise project.<br>The default value is <b>0</b> , indicating the <b>default</b> enterprise project.<br>To obtain the enterprise project ID, see the <b>id</b> value in the <b>enterprise_project field data structure</b> table in "Querying the Enterprise Project List" in <i>Enterprise Management API Reference</i> . |

## Request Parameters

None

## Response Parameters

**Table 7-25** Response parameters

| Parameter    | Type                     | Description                                                       |
|--------------|--------------------------|-------------------------------------------------------------------|
| clusterTotal | Integer                  | Total number of clusters in a list                                |
| clusters     | Array of Cluster objects | Cluster parameters. For details, see <a href="#">Table 7-26</a> . |

**Table 7-26** cluster parameters

| Parameter     | Type   | Description                                   |
|---------------|--------|-----------------------------------------------|
| clusterId     | String | Cluster ID.                                   |
| clusterName   | String | Cluster name.                                 |
| masterNodeNum | String | Number of Master nodes deployed in a cluster. |
| coreNodeNum   | String | Number of Core nodes deployed in a cluster.   |
| totalNodeNum  | String | Total number of nodes deployed in a cluster.  |

| Parameter      | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| clusterState   | String | Cluster status. Valid values include: <ul style="list-style-type: none"> <li>• <b>starting</b>: The cluster is being started.</li> <li>• <b>running</b>: The cluster is running.</li> <li>• <b>terminated</b>: The cluster has been terminated.</li> <li>• <b>failed</b>: The cluster fails.</li> <li>• <b>abnormal</b>: The cluster is abnormal.</li> <li>• <b>terminating</b>: The cluster is being terminated.</li> <li>• <b>frozen</b>: The cluster has been frozen.</li> <li>• <b>scaling-out</b>: The cluster is being scaled out.</li> <li>• <b>scaling-in</b>: The cluster is being scaled in.</li> </ul> |
| createAt       | String | Cluster creation time, which is a 10-bit timestamp.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| updateAt       | String | Cluster update time, which is a 10-bit timestamp.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| billingType    | String | Cluster billing mode. The options are as follows: <ul style="list-style-type: none"> <li>• Yearly/monthly</li> <li>• Pay-per-use</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| dataCenter     | String | Cluster work region.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| vpc            | String | VPC name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| vpcId          | String | VPC ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| duration       | String | Cluster subscription duration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| fee            | String | Cluster creation fee, which is automatically calculated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| hadoopVersion  | String | Hadoop version.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| masterNodeSize | String | Instance specifications of a Master node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| coreNodeSize   | String | Instance specifications of a Core node.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| componentList  | Array  | Component list. For details, see <a href="#">Table 7-27</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| externalIp     | String | External IP address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

| Parameter           | Type    | Description                                                                                                                                                                                                                               |
|---------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| externalAlternatelp | String  | Backup external IP address.                                                                                                                                                                                                               |
| internallp          | String  | Internal IP address.                                                                                                                                                                                                                      |
| deploymentId        | String  | Cluster deployment ID.                                                                                                                                                                                                                    |
| remark              | String  | Cluster remarks.                                                                                                                                                                                                                          |
| orderId             | String  | Cluster creation order ID.                                                                                                                                                                                                                |
| azId                | String  | AZ ID.                                                                                                                                                                                                                                    |
| masterNodeProductId | String  | Product ID of a Master node.                                                                                                                                                                                                              |
| masterNodeSpecId    | String  | Specification ID of a Master node.                                                                                                                                                                                                        |
| coreNodeProductId   | String  | Product ID of a Core node.                                                                                                                                                                                                                |
| coreNodeSpecId      | String  | Specification ID of a Core node.                                                                                                                                                                                                          |
| azName              | String  | AZ name.                                                                                                                                                                                                                                  |
| azCode              | String  | AZ name (en).                                                                                                                                                                                                                             |
| availabilityZoneId  | String  | The AZ.                                                                                                                                                                                                                                   |
| instanceId          | String  | Instance ID.                                                                                                                                                                                                                              |
| vnc                 | String  | URI for remotely logging in to an ECS.                                                                                                                                                                                                    |
| tenantId            | String  | Project ID.                                                                                                                                                                                                                               |
| volumeSize          | Integer | Disk storage space.                                                                                                                                                                                                                       |
| volumeType          | String  | Disk type.                                                                                                                                                                                                                                |
| subnetId            | String  | Subnet ID.                                                                                                                                                                                                                                |
| enterpriseProjectId | String  | Enterprise project ID.                                                                                                                                                                                                                    |
| clusterType         | Integer | Cluster type. The options are as follows: <ul style="list-style-type: none"> <li>● 0: analysis cluster</li> <li>● 1: streaming cluster</li> <li>● 2: hybrid cluster</li> <li>● 3: custom cluster</li> <li>● 4: Offline cluster</li> </ul> |
| subnetName          | String  | Subnet name.                                                                                                                                                                                                                              |
| securityGroupsId    | String  | Security group ID.                                                                                                                                                                                                                        |

| Parameter              | Type                             | Description                                                                                                                                |
|------------------------|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| slaveSecurityGroupslid | String                           | Security group ID of a non-Master node. Currently, one MRS cluster uses only one security group. Therefore, this field has been discarded. |
| bootstrapScripts       | Array of BootstrapScript objects | The bootstrap action script information.<br>For details about this parameter, see <a href="#">Table 7-28</a> .                             |

| Parameter | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| stageDesc | String | <p>Cluster progress description. (The description is in English by default. When the X-Language request header is set to <b>zh-cn</b>, this field returns Chinese.) The cluster installation progress includes:</p> <ul style="list-style-type: none"> <li>• Verifying cluster parameters: Cluster parameters are being verified.</li> <li>• Applying for cluster resources: Cluster resources are being applied for.</li> <li>• Creating VMs: The VMs are being created.</li> <li>• Initializing VMs: The VMs are being initialized.</li> <li>• Installing MRS Manager: MRS Manager is being installed.</li> <li>• Deploying the cluster: The cluster is being deployed.</li> <li>• Cluster installation failed: Failed to install the cluster.</li> </ul> <p>The cluster scale-out progress includes:</p> <ul style="list-style-type: none"> <li>• Preparing for scale-out: Cluster scale-out is being prepared.</li> <li>• Creating VMs: The VMs are being created.</li> <li>• Initializing VMs: The VMs are being initialized.</li> <li>• Adding nodes to the cluster: The nodes are being added to the cluster.</li> <li>• Scale-out failed: Failed to scale out the cluster.</li> </ul> <p>The cluster scale-in progress includes:</p> <ul style="list-style-type: none"> <li>• Preparing for scale-in: Cluster scale-in is being prepared.</li> <li>• Decommissioning instance: The instance is being decommissioned.</li> <li>• Deleting VMs: The VMs are being deleted.</li> <li>• Deleting nodes from the cluster: The nodes are being deleted from the cluster.</li> <li>• Scale-in failed: Failed to scale in the cluster.</li> </ul> |

| Parameter                 | Type            | Description                                                                                                                                                                                                                                     |
|---------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                           |                 | If the cluster installation, scale-out, or scale-in fails, <b>stageDesc</b> will display the failure cause.                                                                                                                                     |
| isMrsManagerFinish        | Boolean         | Whether MRS Manager installation is finished during cluster creation. <ul style="list-style-type: none"> <li>• <b>true</b>: MRS Manager installation is finished.</li> <li>• <b>false</b>: MRS Manager installation is not finished.</li> </ul> |
| safeMode                  | Integer         | Running mode of an MRS cluster. <ul style="list-style-type: none"> <li>• <b>0</b>: Normal cluster</li> <li>• <b>1</b>: Security cluster</li> </ul>                                                                                              |
| clusterVersion            | String          | Cluster version.                                                                                                                                                                                                                                |
| nodePublicCertificateName | String          | Name of the key file.                                                                                                                                                                                                                           |
| masterNodeIp              | String          | IP address of a Master node.                                                                                                                                                                                                                    |
| privateIpFirst            | String          | Preferred private IP address.                                                                                                                                                                                                                   |
| errorInfo                 | String          | Error message.                                                                                                                                                                                                                                  |
| tags                      | String          | The tag information.                                                                                                                                                                                                                            |
| chargingStartTime         | String          | Start time of billing.                                                                                                                                                                                                                          |
| logCollection             | Integer         | Whether to collect logs when cluster installation fails. <ul style="list-style-type: none"> <li>• <b>0</b>: Do not collect logs.</li> <li>• <b>1</b>: Collect logs.</li> </ul>                                                                  |
| taskNodeGroups            | List<NodeGroup> | List of Task nodes. For more parameter description, see <a href="#">Table 7-29</a> .                                                                                                                                                            |
| nodeGroups                | List<NodeGroup> | List of Master, Core and Task nodes. For more parameter description, see <a href="#">Table 7-29</a> .                                                                                                                                           |
| masterDataVolumeType      | String          | Data disk storage type of the Master node. Currently, SATA, SAS, and SSD are supported.                                                                                                                                                         |
| masterDataVolumeSize      | Integer         | Data disk storage space of the Master node To increase data storage capacity, you can add disks at the same time when creating a cluster.<br>Value range: 100 GB to 32,000 GB                                                                   |

| Parameter             | Type    | Description                                                                                                                                                                     |
|-----------------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| masterDataVolumeCount | Integer | Number of data disks of the Master node<br>The value can be set to <b>1</b> only.                                                                                               |
| coreDataVolumeType    | String  | Data disk storage type of the Core node. Currently, SATA, SAS, and SSD are supported.                                                                                           |
| coreDataVolumeSize    | Integer | Data disk storage space of the Core node. To increase data storage capacity, you can add disks at the same time when creating a cluster.<br>Value range: 100 GB to 32,000 GB    |
| coreDataVolumeCount   | Integer | Number of data disks of the Core node.<br>Value range: 1 to 10                                                                                                                  |
| periodType            | Integer | Whether the subscription type is yearly or monthly. <ul style="list-style-type: none"> <li>• <b>0</b>: monthly subscription</li> <li>• <b>1</b>: yearly subscription</li> </ul> |



| Parameter    | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scale        | String | <p>Status of node changes. If this parameter is left blank, no change operation is performed on a cluster node. (The status is in English by default. When the X-Language request header is set to <b>zh-cn</b>, this field returns Chinese text.) The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>Scaling-out</b>: The cluster is being scaled out.</li> <li>• <b>Scaling-in</b>: The cluster is being scaled in.</li> <li>• <b>scaling-error</b>: The cluster is in the running state and fails to be scaled in or out or the specifications fail to be scaled up for the last time.</li> <li>• <b>scaling-up</b>: The master node specifications are being scaled up.</li> <li>• <b>scaling_up_first</b>: The standby master node specifications are being scaled up.</li> <li>• <b>scaled_up_first</b>: The standby master node specifications have been scaled up.</li> <li>• <b>scaled-up-success</b>: The master node specifications have been scaled up.</li> </ul> |
| eipId        | String | Unique ID of the cluster EIP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| eipAddress   | String | IPv4 address of the cluster EIP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| eipv6Address | String | IPv6 address of the cluster EIP. This parameter is not returned when an IPv4 address is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 7-27** componentAmb

| Parameter     | Type   | Description                                                                                                                               |
|---------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------|
| componentId   | String | <p>Component ID</p> <p>For example, the <b>component_id</b> of Hadoop is MRS 3.0.2_001, MRS 2.1.0_001, MRS 1.9.2_001, MRS 1.8.10_001.</p> |
| componentName | String | Component name                                                                                                                            |

| Parameter        | Type   | Description           |
|------------------|--------|-----------------------|
| componentVersion | String | Component version     |
| componentDesc    | String | Component description |

**Table 7-28** bootstrapScripts

| Parameter  | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name       | String           | <p>The name of a bootstrap action script, which must be unique in a cluster.</p> <p>The value can contain only numbers, letters, spaces, hyphens (-), and underscores (_) and cannot start with a space.</p> <p>The value can contain 1 to 64 characters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| uri        | String           | <p>The path of a bootstrap action script. Set this parameter to an OBS bucket path or a local VM path.</p> <ul style="list-style-type: none"> <li>OBS bucket path: Enter a script path. For example, enter the path of the public sample script provided by MRS. Example: <b>s3a://bootstrap/presto/presto-install.sh</b>. If <b>dualroles</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>dualroles</b>. If <b>worker</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>worker</b>. Based on the Presto usage habit, you are advised to install <b>dualroles</b> on the active master nodes and <b>worker</b> on the core nodes.</li> <li>Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul> |
| parameters | String           | The bootstrap action script parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| nodes      | Array of strings | Type of a node where the Bootstrap action script is executed. The value can be <b>master</b> , <b>core</b> , or <b>task</b> . The node type must be represented in lowercase letters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Parameter              | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active_master          | Boolean | Whether the bootstrap action script runs only on active master nodes. The default value is <b>false</b> , indicating that the bootstrap action script can run on all master nodes.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| fail_action            | String  | Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to execute.<br>The default value is <b>errorout</b> , indicating that the action is stopped.<br>Note: You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.<br>Enumerated values: <ul style="list-style-type: none"> <li>• <b>continue</b>: Continue to execute subsequent scripts.</li> <li>• <b>errorout</b>: Stop the action.</li> </ul> |
| before_component_start | Boolean | Time when the bootstrap action script is executed. Currently, the following two options are available: <b>Before component start</b> and <b>After component start</b> . The default value is <b>false</b> , indicating that the bootstrap action script is executed after the component is started.                                                                                                                                                                                                                                                                                                          |
| start_time             | Long    | The execution time of one bootstrap action script.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| state                  | String  | The running status of one bootstrap action script. <ul style="list-style-type: none"> <li>• PENDING</li> <li>• IN_PROGRESS</li> <li>• SUCCESS</li> <li>• FAILURE</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Parameter     | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|---------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| action_stages | Array of strings | <p>Select the time when the bootstrap action script is executed.</p> <ul style="list-style-type: none"> <li>• <b>BEFORE_COMPONENT_FIRST_START</b>: before initial component starts</li> <li>• <b>AFTER_COMPONENT_FIRST_START</b>: after initial component starts</li> <li>• <b>BEFORE_SCALE_IN</b>: before scale-in</li> <li>• <b>AFTER_SCALE_IN</b>: after scale-in</li> <li>• <b>BEFORE_SCALE_OUT</b>: before scale-out</li> <li>• <b>AFTER_SCALE_OUT</b>: after scale-out</li> </ul> |

**Table 7-29** NodeGroupV10

| Parameter                  | Type    | Description                                                                                                                                                      |
|----------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GroupName                  | String  | Node group name.                                                                                                                                                 |
| NodeNum                    | Integer | Number of nodes. The value ranges from 0 to 500. The minimum number of Master and Core nodes is 1 and the total number of Core and Task nodes cannot exceed 500. |
| NodeSize                   | String  | Instance specifications of a node.                                                                                                                               |
| NodeSpecId                 | String  | Instance specification ID of a node                                                                                                                              |
| NodeProductId              | String  | Instance product ID of a node                                                                                                                                    |
| VmProductId                | String  | VM product ID of a node                                                                                                                                          |
| VmSpecCode                 | String  | VM specifications of a node                                                                                                                                      |
| RootVolumeSize             | Integer | System disk size of a node. This parameter is not configurable and its default value is <b>40 GB</b> .                                                           |
| RootVolumeProductId        | String  | System disk product ID of a node                                                                                                                                 |
| RootVolumeType             | String  | System disk type of a node                                                                                                                                       |
| RootVolumeResourceSpecCode | String  | System disk product specifications of a node                                                                                                                     |
| RootVolumeResourceType     | String  | System disk product type of a node                                                                                                                               |

| Parameter                  | Type    | Description                                                                                                                                                                                                                      |
|----------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DataVolumeType             | String  | Data disk storage type of a node. Currently, SATA, SAS, and SSD are supported. <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O</li> <li>• <b>SAS</b>: high I/O</li> <li>• <b>SSD</b>: ultra-high I/O</li> </ul> |
| DataVolumeCount            | Integer | Number of data disks of a node.                                                                                                                                                                                                  |
| DataVolumeSize             | Integer | Data disk storage space of a node.                                                                                                                                                                                               |
| DataVolumeProductId        | String  | Data disk product ID of a node                                                                                                                                                                                                   |
| DataVolumeResourceSpecCode | String  | Data disk product specifications of a node                                                                                                                                                                                       |
| DataVolumeResourceType     | String  | Data disk product type of a node                                                                                                                                                                                                 |

## Example Response

**Status code: 200**

Querying the cluster list is successful.

```
{
  "clusterTotal": "1",
  "clusters": [ {
    "clusterId": "bc134369-294c-42b7-a707-b2036ba38524",
    "clusterName": "mrs_D0zW",
    "masterNodeNum": "2",
    "coreNodeNum": "3",
    "clusterState": "terminated",
    "createAt": "1498272043",
    "updateAt": "1498636753",
    "chargingStartTime": "1498273733",
    "logCollection": "1",
    "billingType": "Metered",
    "dataCenter": "cn-hk1",
    "vpc": null,
    "duration": "0",
    "fee": null,
    "hadoopVersion": null,
    "masterNodeSize": null,
    "coreNodeSize": null,
    "componentList": [ {
      "componentId": "MRS 3.1.0_001",
      "componentName": "Hadoop",
      "componentVersion": "3.1.1",
      "componentDesc": "A framework that allows for the distributed processing of large data sets across clusters."
    }, {
      "componentId": "MRS 3.1.0_003",
      "componentName": "HBase",
      "componentVersion": "2.1.1",
      "componentDesc": "A scalable, distributed database that supports structured data storage for large tables."
    }
  ]
}
}
```

```
    }, {
      "componentId" : "MRS 3.1.0_002",
      "componentName" : "Spark",
      "componentVersion" : "2.3.2",
      "componentDesc" : "A fast and general engine for large-scale data processing."
    }, {
      "componentId" : "MRS 3.1.0_004",
      "componentName" : "Hive",
      "componentVersion" : "3.1.0",
      "componentDesc" : "A data warehouse infrastructure that provides data summarization and ad hoc
querying."
    } ],
    "externalIp" : null,
    "externalAlternateIp" : null,
    "internalIp" : null,
    "deploymentId" : null,
    "remark" : "",
    "orderId" : null,
    "azId" : null,
    "azCode" : null,
    "masterNodeProductId" : null,
    "masterNodeSpecId" : null,
    "coreNodeProductId" : null,
    "coreNodeSpecId" : null,
    "azName" : "AZ1",
    "instanceId" : null,
    "vnc" : "v2/5a3314075bfa49b9ae360f4ecd333695/servers/e2cda891-232e-4703-995e-3b1406add01d/
action",
    "tenantId" : null,
    "volumeSize" : "0",
    "volumeType" : null,
    "subnetId" : null,
    "subnetName" : null,
    "securityGroupsId" : null,
    "slaveSecurityGroupsId" : null,
    "bootstrapScripts" : [ {
      "name" : "test1-success",
      "uri" : "s3a://bootscrip/script/simple/basic_success.sh",
      "parameters" : "",
      "nodes" : [ "master", "core" ],
      "active_master" : true,
      "fail_action" : "errorout",
      "before_component_start" : true,
      "state" : "SUCCESS",
      "start_time" : 1527681083,
      "action_stages" : [ "AFTER_SCALE_IN", "AFTER_SCALE_OUT" ]
    } ],
    "isMrsManagerFinish" : false,
    "stageDesc" : "Installing MRS Manager",
    "safeMode" : "0",
    "clusterVersion" : null,
    "nodePublicCertName" : null,
    "masterNodeIp" : "unknown",
    "privateIpFirst" : null,
    "errorInfo" : "",
    "clusterType" : "0",
    "nodeGroups" : [ {
      "GroupName" : "master_node_default_group",
      "NodeNum" : "1",
      "NodeSize" : "s3.xlarge.2.linux.bigdata",
      "NodeSpecId" : "cdc6035a249a40249312f5ef72a23cd7",
      "VmProductId" : "",
      "VmSpecCode" : null,
      "NodeProductId" : "dc970349d128460e960a0c2b826c427c",
      "RootVolumeSize" : "40",
      "RootVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
      "RootVolumeType" : "SATA",
      "RootVolumeResourceSpecCode" : "",
      "RootVolumeResourceType" : ""
    } ],
```

```
"DataVolumeType": "SATA",
"DataVolumeCount": "1",
"DataVolumeSize": "100",
"DataVolumeProductId": "16c1dcf0897249758b1ec276d06e0572",
"DataVolumeResourceSpecCode": "",
"DataVolumeResourceType": ""
}, {
  "GroupName": "core_node_analysis_group",
  "NodeNum": "1",
  "NodeSize": "s3.xlarge.2.linux.bigdata",
  "NodeSpecId": "cdc6035a249a40249312f5ef72a23cd7",
  "VmProductId": "",
  "VmSpecCode": null,
  "NodeProductId": "dc970349d128460e960a0c2b826c427c",
  "RootVolumeSize": "40",
  "RootVolumeProductId": "16c1dcf0897249758b1ec276d06e0572",
  "RootVolumeType": "SATA",
  "RootVolumeResourceSpecCode": "",
  "RootVolumeResourceType": "",
  "DataVolumeType": "SATA",
  "DataVolumeCount": "1",
  "DataVolumeSize": "100",
  "DataVolumeProductId": "16c1dcf0897249758b1ec276d06e0572",
  "DataVolumeResourceSpecCode": "",
  "DataVolumeResourceType": ""
}, {
  "GroupName": "task_node_analysis_group",
  "NodeNum": "1",
  "NodeSize": "s3.xlarge.2.linux.bigdata",
  "NodeSpecId": "cdc6035a249a40249312f5ef72a23cd7",
  "VmProductId": "",
  "VmSpecCode": null,
  "NodeProductId": "dc970349d128460e960a0c2b826c427c",
  "RootVolumeSize": "40",
  "RootVolumeProductId": "16c1dcf0897249758b1ec276d06e0572",
  "RootVolumeType": "SATA",
  "RootVolumeResourceSpecCode": "",
  "RootVolumeResourceType": "",
  "DataVolumeType": "SATA",
  "DataVolumeCount": "1",
  "DataVolumeSize": "100",
  "DataVolumeProductId": "16c1dcf0897249758b1ec276d06e0572",
  "DataVolumeResourceSpecCode": "",
  "DataVolumeResourceType": ""
}],
"taskNodeGroups": [ {
  "GroupName": "task_node_default_group",
  "NodeNum": "1",
  "NodeSize": "s3.xlarge.2.linux.bigdata",
  "NodeSpecId": "cdc6035a249a40249312f5ef72a23cd7",
  "VmProductId": "",
  "VmSpecCode": null,
  "NodeProductId": "dc970349d128460e960a0c2b826c427c",
  "RootVolumeSize": "40",
  "RootVolumeProductId": "16c1dcf0897249758b1ec276d06e0572",
  "RootVolumeType": "SATA",
  "RootVolumeResourceSpecCode": "",
  "RootVolumeResourceType": "",
  "DataVolumeType": "SATA",
  "DataVolumeCount": "1",
  "DataVolumeSize": "100",
  "DataVolumeProductId": "16c1dcf0897249758b1ec276d06e0572",
  "DataVolumeResourceSpecCode": "",
  "DataVolumeResourceType": ""
} ],
"masterDataVolumeType": "SATA",
"masterDataVolumeSize": "200",
"masterDataVolumeCount": "1",
"coreDataVolumeType": "SATA",
```

```
"coreDataVolumeSize" : "100",  
"coreDataVolumeCount" : "1",  
"periodType" : "0"  
}]  
}
```

## Status Codes

[Table 7-30](#) describes the status code.

**Table 7-30** Status code

| Status Code | Description                                    |
|-------------|------------------------------------------------|
| 200         | The cluster list information has been queried. |

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 7.1.4 Querying Cluster Details

### Function

This API is used to query details about a specified cluster. This API is incompatible with Sahara.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
GET /v1.1/{project\_id}/cluster\_infos/{cluster\_id}
- Parameter description



**Table 7-31** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                       |
|------------|-----------|--------|-------------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .          |
| cluster_id | Yes       | String | Cluster ID For details about how to obtain the value, see <a href="#">Obtaining the MRS Cluster Information</a> . |

## Request Parameters

None

## Response Parameters

**Table 7-32** Response body parameter

| Parameter | Type           | Description                                                       |
|-----------|----------------|-------------------------------------------------------------------|
| cluster   | Cluster object | Cluster parameters. For details, see <a href="#">Table 7-33</a> . |

**Table 7-33** Response parameters

| Parameter    | Type   | Description                                 |
|--------------|--------|---------------------------------------------|
| clusterId    | String | Cluster ID                                  |
| clusterName  | String | Cluster name                                |
| totalNodeNum | String | Total number of nodes deployed in a cluster |

| Parameter    | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| clusterState | String | <p>Cluster status. (The status is in English by default. When the X-Language request header is set to <b>zh-cn</b>, this field returns Chinese text.) The options are as follows:</p> <ul style="list-style-type: none"><li>• <b>starting</b>: The cluster is being started.</li><li>• <b>running</b>: The cluster is running.</li><li>• <b>terminated</b>: The cluster has been deleted.</li><li>• <b>failed</b>: The cluster is failed.</li><li>• <b>abnormal</b>: The cluster is abnormal.</li><li>• <b>terminating</b>: The cluster is being deleted.</li><li>• <b>frozen</b>: The cluster is frozen.</li><li>• <b>scaling-out</b>: The cluster is being scaled out.</li><li>• <b>scaling-in</b>: The cluster is being scaled in.</li></ul> |

| Parameter | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| stageDesc | String | <p>Cluster progress description. (The description is in English by default. When the X-Language request header is set to <b>zh-cn</b>, this field returns Chinese text.) The cluster installation progress includes:</p> <ul style="list-style-type: none"> <li>• Verifying cluster parameters: Cluster parameters are being verified.</li> <li>• Applying for cluster resources: Cluster resources are being applied for.</li> <li>• Creating VMs: The VMs are being created.</li> <li>• Initializing VMs: The VMs are being initialized.</li> <li>• Installing MRS Manager: MRS Manager is being installed.</li> <li>• Deploying the cluster: The cluster is being deployed.</li> <li>• Cluster installation failed: Failed to install the cluster.</li> </ul> <p>The cluster scale-out progress includes:</p> <ul style="list-style-type: none"> <li>• Preparing for scale-out: Cluster scale-out is being prepared.</li> <li>• Creating VMs: The VMs are being created.</li> <li>• Initializing VMs: The VMs are being initialized.</li> <li>• Adding nodes to the cluster: The nodes are being added to the cluster.</li> <li>• Scale-out failed: Failed to scale out the cluster.</li> </ul> <p>The cluster scale-in progress includes:</p> <ul style="list-style-type: none"> <li>• Preparing for scale-in: Cluster scale-in is being prepared.</li> <li>• Decommissioning instance: The instance is being decommissioned.</li> <li>• Deleting VMs: The VMs are being deleted.</li> <li>• Deleting nodes from the cluster: The nodes are being deleted from the cluster.</li> </ul> |

| Parameter           | Type   | Description                                                                                                                                                                                                           |
|---------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     |        | <ul style="list-style-type: none"> <li>Scale-in failed: Failed to scale in the cluster.</li> </ul> <p>If the cluster installation, scale-out, or scale-in fails, <b>stageDesc</b> will display the failure cause.</p> |
| createAt            | String | Cluster creation time, which is a 10-bit timestamp                                                                                                                                                                    |
| updateAt            | String | Cluster update time, which is a 10-bit timestamp                                                                                                                                                                      |
| chargingStartTime   | String | Start time of billing                                                                                                                                                                                                 |
| billingType         | String | Cluster billing mode. The options are as follows: <ul style="list-style-type: none"> <li><b>11</b>: Yearly/Monthly</li> <li><b>12</b>: Pay-per-use</li> </ul>                                                         |
| dataCenter          | String | Cluster work region                                                                                                                                                                                                   |
| vpc                 | String | VPC name                                                                                                                                                                                                              |
| vpcId               | String | VPC ID                                                                                                                                                                                                                |
| duration            | String | Cluster subscription duration                                                                                                                                                                                         |
| fee                 | String | Cluster creation fee, which is automatically calculated                                                                                                                                                               |
| hadoopVersion       | String | Hadoop version                                                                                                                                                                                                        |
| componentList       | Array  | Component list. For details, see <a href="#">Table 7-34</a> .                                                                                                                                                         |
| externalIp          | String | External IP address                                                                                                                                                                                                   |
| externalAlternateIp | String | Backup external IP address                                                                                                                                                                                            |
| internalIp          | String | Internal IP address                                                                                                                                                                                                   |
| deploymentId        | String | Cluster deployment ID                                                                                                                                                                                                 |
| remark              | String | Cluster remarks                                                                                                                                                                                                       |
| orderId             | String | Cluster creation order ID                                                                                                                                                                                             |
| azId                | String | AZ ID                                                                                                                                                                                                                 |
| azName              | String | AZ name                                                                                                                                                                                                               |
| azCode              | String | AZ name (en).                                                                                                                                                                                                         |
| availabilityZoneId  | String | AZ.                                                                                                                                                                                                                   |

| Parameter             | Type                             | Description                                                                                                                                                                                                                                   |
|-----------------------|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| instanceId            | String                           | Instance ID                                                                                                                                                                                                                                   |
| vnc                   | String                           | URI for remotely logging in to an ECS                                                                                                                                                                                                         |
| tenantId              | String                           | Project ID                                                                                                                                                                                                                                    |
| volumeSize            | Integer                          | Disk storage space                                                                                                                                                                                                                            |
| volumeType            | String                           | Disk type.                                                                                                                                                                                                                                    |
| subnetId              | String                           | Subnet ID                                                                                                                                                                                                                                     |
| subnetName            | String                           | Subnet name                                                                                                                                                                                                                                   |
| securityGroupsId      | String                           | Security group ID                                                                                                                                                                                                                             |
| slaveSecurityGroupsId | String                           | Security group ID of a non-master node. Currently, one MRS cluster uses only one security group. Therefore, this field has been discarded. This field returns the same value as <b>securityGroupsId</b> does for compatibility consideration. |
| bootstrapScripts      | Array of BootstrapScript objects | Bootstrap action script information. For more parameter description, see <a href="#">Table 7-36</a> .                                                                                                                                         |
| safeMode              | Integer                          | Running mode of an MRS cluster <ul style="list-style-type: none"> <li>• 0: Normal cluster</li> <li>• 1: Security cluster</li> </ul>                                                                                                           |
| clusterVersion        | String                           | Cluster version                                                                                                                                                                                                                               |
| nodePublicKeyName     | String                           | Name of the public key file                                                                                                                                                                                                                   |
| masterNodeIp          | String                           | IP address of a Master node                                                                                                                                                                                                                   |
| privateIpFirst        | String                           | Preferred private IP address                                                                                                                                                                                                                  |
| errorMsg              | String                           | Error message                                                                                                                                                                                                                                 |
| tags                  | String                           | Tag information                                                                                                                                                                                                                               |
| masterNodeNum         | String                           | Number of Master nodes deployed in a cluster                                                                                                                                                                                                  |
| coreNodeNum           | String                           | Number of Core nodes deployed in a cluster                                                                                                                                                                                                    |
| masterNodeSize        | String                           | Instance specifications of a Master node                                                                                                                                                                                                      |
| coreNodeSize          | String                           | Instance specifications of a Core node                                                                                                                                                                                                        |

| Parameter             | Type    | Description                                                                                                                                                                                                                                     |
|-----------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| masterNodeProductId   | String  | Product ID of a Master node                                                                                                                                                                                                                     |
| masterNodeSpecId      | String  | Specification ID of a Master node                                                                                                                                                                                                               |
| coreNodeProductId     | String  | Product ID of a Core node                                                                                                                                                                                                                       |
| coreNodeSpecId        | String  | Specification ID of a Core node                                                                                                                                                                                                                 |
| masterDataVolumeType  | String  | Data disk storage type of the Master node. Currently, SATA, SAS, and SSD are supported.                                                                                                                                                         |
| masterDataVolumeSize  | Integer | Data disk storage space of the Master node. To increase data storage capacity, you can add disks at the same time when creating a cluster. Value range: 100 GB to 32,000 GB                                                                     |
| masterDataVolumeCount | Integer | Number of data disks of the Master node.<br>The value can be set to <b>1</b> only.                                                                                                                                                              |
| coreDataVolumeType    | String  | Data disk storage type of the Core node. Currently, SATA, SAS, and SSD are supported.                                                                                                                                                           |
| coreDataVolumeSize    | Integer | Data disk storage space of the Core node. To increase data storage capacity, you can add disks at the same time when creating a cluster. You only need to enter a number without the unit GB. Value range: 100 GB to 32,000 GB                  |
| coreDataVolumeCount   | Integer | Number of data disks of the Core node.<br>Value range: 1 to 10                                                                                                                                                                                  |
| enterpriseProjectId   | String  | Enterprise project ID                                                                                                                                                                                                                           |
| isMrsManagerFinish    | Boolean | Whether MRS Manager installation is finished during cluster creation. <ul style="list-style-type: none"> <li>• <b>true</b>: MRS Manager installation is finished.</li> <li>• <b>false</b>: MRS Manager installation is not finished.</li> </ul> |

| Parameter     | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| clusterType   | Integer | Cluster type. The options are as follows: <ul style="list-style-type: none"> <li>• <b>0</b>: analysis cluster</li> <li>• <b>1</b>: streaming cluster</li> <li>• <b>2</b>: hybrid cluster</li> <li>• <b>3</b>: custom cluster</li> <li>• <b>4</b>: Offline cluster</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| logCollection | Integer | Whether to collect logs when cluster installation fails <ul style="list-style-type: none"> <li>• <b>0</b>: Do not collect.</li> <li>• <b>1</b>: Collect.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| periodType    | Integer | Whether the subscription type is yearly or monthly <ul style="list-style-type: none"> <li>• <b>0</b>: monthly subscription</li> <li>• <b>1</b>: yearly subscription</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| scale         | String  | Node change status. If this parameter is left blank, no change operation is performed on a cluster node. (The status is in English by default. When the X-Language request header is set to <b>zh-cn</b> , this field returns Chinese text.) The options are as follows: <ul style="list-style-type: none"> <li>• <b>scaling-out</b>: The cluster is being scaled out.</li> <li>• <b>scaling-in</b>: The cluster is being scaled in.</li> <li>• <b>scaling-error</b>: The cluster is in the running state and fails to be scaled in or out or the specifications fail to be scaled up for the last time.</li> <li>• <b>scaling-up</b>: The Master node specifications are being scaled up.</li> <li>• <b>scaling_up_first</b>: The standby Master node specifications are being scaled up.</li> <li>• <b>scaled_up_first</b>: The standby Master node specifications have been scaled up successfully.</li> <li>• <b>scaled-up-success</b>: The master node specifications have been scaled up.</li> </ul> |

| Parameter      | Type            | Description                                                                                           |
|----------------|-----------------|-------------------------------------------------------------------------------------------------------|
| nodeGroups     | List<NodeGroup> | List of Master, Core and Task nodes. For more parameter description, see <a href="#">Table 7-35</a> . |
| taskNodeGroups | List<NodeGroup> | List of Task nodes. For more parameter description, see <a href="#">Table 7-35</a> .                  |
| eipId          | String          | Unique ID of the cluster EIP                                                                          |
| eipAddress     | String          | IPv4 address of the cluster EIP                                                                       |
| eipv6Address   | String          | IPv6 address of the cluster EIP. This parameter is not returned when an IPv4 address is used.         |

**Table 7-34** componentAmb

| Parameter        | Type   | Description                                                                                                                    |
|------------------|--------|--------------------------------------------------------------------------------------------------------------------------------|
| componentId      | String | Component ID<br>For example, the <b>component_id</b> of Hadoop is MRS 3.0.2_001, MRS 2.1.0_001, MRS 1.9.2_001, MRS 1.8.10_001. |
| componentName    | String | Component name                                                                                                                 |
| componentVersion | String | Component version                                                                                                              |
| componentDesc    | String | Component description                                                                                                          |

**Table 7-35** NodeGroupV10

| Parameter     | Type    | Description                                                                                                                                                      |
|---------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GroupName     | String  | Node group name.                                                                                                                                                 |
| NodeNum       | Integer | Number of nodes. The value ranges from 0 to 500. The minimum number of Master and Core nodes is 1 and the total number of Core and Task nodes cannot exceed 500. |
| NodeSize      | String  | Instance specifications of a node.                                                                                                                               |
| NodeSpecId    | String  | Instance specification ID of a node                                                                                                                              |
| NodeProductId | String  | Instance product ID of a node                                                                                                                                    |



| Parameter                  | Type    | Description                                                                                                                                                                                                                      |
|----------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VmProductId                | String  | VM product ID of a node                                                                                                                                                                                                          |
| VmSpecCode                 | String  | VM specifications of a node                                                                                                                                                                                                      |
| RootVolumeSize             | Integer | System disk size of a node. This parameter is not configurable and its default value is <b>40 GB</b> .                                                                                                                           |
| RootVolumeProductId        | String  | System disk product ID of a node                                                                                                                                                                                                 |
| RootVolumeType             | String  | System disk type of a node                                                                                                                                                                                                       |
| RootVolumeResourceSpecCode | String  | System disk product specifications of a node                                                                                                                                                                                     |
| RootVolumeResourceType     | String  | System disk product type of a node                                                                                                                                                                                               |
| DataVolumeType             | String  | Data disk storage type of a node. Currently, SATA, SAS, and SSD are supported. <ul style="list-style-type: none"> <li>• <b>SATA</b>: common I/O</li> <li>• <b>SAS</b>: high I/O</li> <li>• <b>SSD</b>: ultra-high I/O</li> </ul> |
| DataVolumeCount            | Integer | Number of data disks of a node.                                                                                                                                                                                                  |
| DataVolumeSize             | Integer | Data disk storage space of a node.                                                                                                                                                                                               |
| DataVolumeProductId        | String  | Data disk product ID of a node                                                                                                                                                                                                   |
| DataVolumeResourceSpecCode | String  | Data disk product specifications of a node                                                                                                                                                                                       |
| DataVolumeResourceType     | String  | Data disk product type of a node                                                                                                                                                                                                 |

**Table 7-36** BootstrapScripts

| Parameter | Type   | Description                                                                                                                                                                                                                            |
|-----------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | String | Name of a bootstrap action script. It must be unique in a cluster.<br>The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and cannot start with a space.<br>The value can contain 1 to 64 characters. |

| Parameter              | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri                    | String       | <p>Path of the shell script. Set this parameter to an OBS bucket path or a local VM path.</p> <ul style="list-style-type: none"> <li>OBS bucket path: Enter a script path manually. For example, enter the path of the public sample script provided by MRS. Example: <b>s3a://bootstrap/presto/presto-install.sh</b>. If <b>dualroles</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>dualroles</b>. If <b>worker</b> is installed, the parameter of the <b>presto-install.sh</b> script is <b>worker</b>. Based on the Presto usage habit, you are advised to install <b>dualroles</b> on the active Master nodes and <b>worker</b> on the Core nodes.</li> <li>Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul> |
| parameters             | String       | Bootstrap action script parameters                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| nodes                  | Array String | Type of a node where the bootstrap action script is executed. The value can be Master, Core, or Task.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| active_master          | Boolean      | <p>Whether the bootstrap action script runs only on active Master nodes.</p> <p>The default value is <b>false</b>, indicating that the bootstrap action script can run on all Master nodes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| before_component_start | Boolean      | <p>Time when the bootstrap action script is executed. Currently, the following two options are available: <b>Before component start</b> and <b>After component start</b></p> <p>The default value is <b>false</b>, indicating that the bootstrap action script is executed after the component is started.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Parameter     | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fail_action   | String           | <p>Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to be executed.</p> <ul style="list-style-type: none"> <li>• <b>continue</b>: Continue to execute subsequent scripts.</li> <li>• <b>errorout</b>: Stop the action.</li> </ul> <p>The default value is <b>errorout</b>, indicating that the action is stopped.</p> <p><b>NOTE</b><br/>You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p> |
| start_time    | Long             | Execution time of one boot operation script.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| state         | String           | <p>Running state of one bootstrap action script</p> <ul style="list-style-type: none"> <li>• <b>PENDING</b></li> <li>• <b>IN_PROGRESS</b></li> <li>• <b>SUCCESS</b></li> <li>• <b>FAILURE</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                           |
| action_stages | Array of strings | <p>Select the time when the bootstrap action script is executed.</p> <ul style="list-style-type: none"> <li>• <b>BEFORE_COMPONENT_FIRST_START</b>: before initial component starts</li> <li>• <b>AFTER_COMPONENT_FIRST_START</b>: after initial component starts</li> <li>• <b>BEFORE_SCALE_IN</b>: before scale-in</li> <li>• <b>AFTER_SCALE_IN</b>: after scale-in</li> <li>• <b>BEFORE_SCALE_OUT</b>: before scale-out</li> <li>• <b>AFTER_SCALE_OUT</b>: after scale-out</li> </ul>                                                                                                                            |

## Example Request

```
GET /v1.1/{project_id}/cluster_infos/{cluster_id}
```

## Example response

**Status code: 200**

Cluster details are queried successfully.

```
{
  "cluster" : {
    "clusterId" : "bdb064ff-2855-4624-90d5-e9a6376abd6e",
    "clusterName" : "c17022001",
    "masterNodeNum" : "2",
    "coreNodeNum" : "3",
    "clusterState" : "scaling-in",
    "stageDesc" : "Installing MRS Manager",
    "createAt" : "1487570757",
    "updateAt" : "1487668974",
    "billingType" : "Metered",
    "dataCenter" : "cn-hk1",
    "vpc" : "vpc-autotest",
    "vpclId" : "e2978efd-ca12-4058-9332-1ca0bfbab592",
    "duration" : "0",
    "fee" : "0",
    "hadoopVersion" : "",
    "masterNodeSize" : "s3.2xlarge.2.linux.bigdata",
    "coreNodeSize" : "s3.2xlarge.2.linux.bigdata",
    "componentList" : [ {
      "componentId" : "MRS 3.2.0-LTS.1_001",
      "componentName" : "Hadoop",
      "componentVersion" : "3.1.1",
      "componentDesc" : "A framework that allows for the distributed processing of large data sets across
clusters."
    }, {
      "componentId" : "MRS 3.2.0-LTS.1_002",
      "componentName" : "Spark",
      "componentVersion" : "2.3.2",
      "componentDesc" : "A fast and general engine for large-scale data processing."
    }, {
      "componentId" : "MRS 3.2.0-LTS.1_004",
      "componentName" : "Hive",
      "componentVersion" : "3.1.0",
      "componentDesc" : "A data warehouse infrastructure that provides data summarization and ad hoc
querying."
    }, {
      "componentId" : "MRS 3.2.0-LTS.1_003",
      "componentName" : "HBase",
      "componentVersion" : "2.1.1",
      "componentDesc" : "A scalable, distributed database that supports structured data storage for large
tables."
    }
  ],
  "externalIp" : "100.XXX.XXX.XXX",
  "externalAlternateIp" : "100.XXX.XXX.XXX",
  "internalIp" : "192.XXX.XXX.XXX",
  "eipId" : "b16dd5eb-5e5b-486a-906a-2e8f6e814a7a",
  "eipAddress" : "100.XXX.XXX.XXX",
  "eipv6Address" : "2403:XXXX:XXXX:XXXX:XXXX",
  "deploymentId" : "4ac46ca7-a488-4b91-82c2-e4d7aa9c40c2",
  "remark" : "",
  "orderId" : "null",
  "azId" : "null",
  "masterNodeProductId" : "b35cf2d2348a445ca74b32289a160882",
  "masterNodeSpecId" : "8ab05e503b4c42abb304e2489560063b",
  "coreNodeProductId" : "dc970349d128460e960a0c2b826c427c",
  "coreNodeSpecId" : "cdc6035a249a40249312f5ef72a23cd7",
  "azName" : "AZ1",
  "instanceId" : "4ac46ca7-a488-4b91-82c2-e4d7aa9c40c2",
  "vnc" : null,
  "tenantId" : "3f99e3319a8943ceb15c584f3325d064",
  "volumeSize" : "600",
  "volumeType" : "SATA",
  "subnetId" : "6b96eec3-4f8d-4c83-93e2-6ec625001d7c",
  "subnetName" : "subnet-ftest",
  "securityGroupId" : "930e34e2-195d-401f-af07-0b64ea6603f8",
  "slaveSecurityGroupId" : "2ef3343e-3477-4a0d-80fe-4d874e4f81b8",
  "bootstrapScripts" : [ {
    "name" : "test1-success",
  }
  ]
}
```

```
"uri" : "s3a://bootscrip/script/simple/basic_success.sh",
"parameters" : "",
"nodes" : [ "master", "core" ],
"active_master" : true,
"fail_action" : "errorout",
"before_component_start" : true,
"state" : "SUCCESS",
"start_time" : 1527681083,
"action_stages" : [ "AFTER_SCALE_IN", "AFTER_SCALE_OUT" ]
}],
"isMrsManagerFinish" : false,
"safeMode" : "1",
"clusterVersion" : "MRS 3.2.0-LTS.1",
"nodePublicCertName" : "myp",
"masterNodeIp" : "192.XXX.XXX.XXX",
"privateIpFirst" : "192.XXX.XXX.XXX",
"errorInfo" : null,
"tags" : "k1=v1,k2=v2,k3=v3",
"clusterType" : "",
"logCollection" : "1",
"nodeGroups" : [ {
  "GroupName" : "master_node_default_group",
  "NodeNum" : "1",
  "NodeSize" : "s3.xlarge.2.linux.bigdata",
  "NodeSpecId" : "cdc6035a249a40249312f5ef72a23cd7",
  "VmProductId" : "",
  "VmSpecCode" : null,
  "NodeProductId" : "dc970349d128460e960a0c2b826c427c",
  "RootVolumeSize" : "480",
  "RootVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
  "RootVolumeType" : "SATA",
  "RootVolumeResourceSpecCode" : "",
  "RootVolumeResourceType" : "",
  "DataVolumeType" : "SATA",
  "DataVolumeCount" : "1",
  "DataVolumeSize" : "600",
  "DataVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
  "DataVolumeResourceSpecCode" : "",
  "DataVolumeResourceType" : ""
}, {
  "GroupName" : "core_node_analysis_group",
  "NodeNum" : "1",
  "NodeSize" : "s3.xlarge.2.linux.bigdata",
  "NodeSpecId" : "cdc6035a249a40249312f5ef72a23cd7",
  "VmProductId" : "",
  "VmSpecCode" : null,
  "NodeProductId" : "dc970349d128460e960a0c2b826c427c",
  "RootVolumeSize" : "480",
  "RootVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
  "RootVolumeType" : "SATA",
  "RootVolumeResourceSpecCode" : "",
  "RootVolumeResourceType" : "",
  "DataVolumeType" : "SATA",
  "DataVolumeCount" : "1",
  "DataVolumeSize" : "600",
  "DataVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
  "DataVolumeResourceSpecCode" : "",
  "DataVolumeResourceType" : ""
}, {
  "GroupName" : "task_node_analysis_group",
  "NodeNum" : "1",
  "NodeSize" : "s3.xlarge.2.linux.bigdata",
  "NodeSpecId" : "cdc6035a249a40249312f5ef72a23cd7",
  "VmProductId" : "",
  "VmSpecCode" : null,
  "NodeProductId" : "dc970349d128460e960a0c2b826c427c",
  "RootVolumeSize" : "480",
  "RootVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
  "RootVolumeType" : "SATA",
```

```

"RootVolumeResourceSpecCode" : "",
"RootVolumeResourceType" : "",
"DataVolumeType" : "SATA",
"DataVolumeCount" : "1",
"DataVolumeSize" : "600",
"DataVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
"DataVolumeResourceSpecCode" : "",
"DataVolumeResourceType" : ""
}],
"taskNodeGroups" : [ {
"GroupName" : "task_node_default_group",
"NodeNum" : "1",
"NodeSize" : "s3.xlarge.2.linux.bigdata",
"NodeSpecId" : "cdc6035a249a40249312f5ef72a23cd7",
"VmProductId" : "",
"VmSpecCode" : null,
"NodeProductId" : "dc970349d128460e960a0c2b826c427c",
"RootVolumeSize" : "480",
"RootVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
"RootVolumeType" : "SATA",
"RootVolumeResourceSpecCode" : "",
"RootVolumeResourceType" : "",
"DataVolumeType" : "SATA",
"DataVolumeCount" : "1",
"DataVolumeSize" : "600",
"DataVolumeProductId" : "16c1dcf0897249758b1ec276d06e0572",
"DataVolumeResourceSpecCode" : "",
"DataVolumeResourceType" : ""
}],
"masterDataVolumeType" : "SATA",
"masterDataVolumeSize" : "600",
"masterDataVolumeCount" : "1",
"coreDataVolumeType" : "SATA",
"coreDataVolumeSize" : "600",
"coreDataVolumeCount" : "1",
"periodType" : "0"
}
}

```

## Status Codes

[Table 7-37](#) describes the status code.

**Table 7-37** Status code

| Status Code | Description                        |
|-------------|------------------------------------|
| 200         | Cluster details have been queried. |

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

### 7.1.5 Querying a Host List

#### Function

This API is used to query a host list of a specified cluster.

## Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- Format  
GET /v1.1/{project\_id}/clusters/{cluster\_id}/hosts
- Parameter description

**Table 7-38** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                       |
|------------|-----------|--------|-------------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .          |
| cluster_id | Yes       | String | Cluster ID For details about how to obtain the value, see <a href="#">Obtaining the MRS Cluster Information</a> . |

**Table 7-39** Query parameters

| Parameter   | Mandatory | Type    | Description                                                                                                        |
|-------------|-----------|---------|--------------------------------------------------------------------------------------------------------------------|
| pageSize    | No        | Integer | Maximum number of clusters displayed on a page<br>Value range: [1-2147483646].<br>The default value is <b>10</b> . |
| currentPage | No        | Integer | Current page number The default value is <b>1</b> .                                                                |

## Request Parameters

None

## Response Parameters

**Table 7-40** Response body parameters

| Parameter | Type                       | Description                                                      |
|-----------|----------------------------|------------------------------------------------------------------|
| total     | Integer                    | Total number of hosts in a list                                  |
| hosts     | Array of HostModel objects | Host parameters<br>For details, see <a href="#">Table 7-41</a> . |

**Table 7-41** HostModel

| Parameter            | Type                      | Description                                                             |
|----------------------|---------------------------|-------------------------------------------------------------------------|
| id                   | String                    | VM ID                                                                   |
| ip                   | String                    | VM IP address                                                           |
| availability_zone_id | String                    | The AZ.                                                                 |
| tags                 | Array of TagPlain objects | The tag list. For details, see <a href="#">Table 7-42</a> .             |
| flavor               | String                    | VM flavor ID                                                            |
| type                 | String                    | VM type<br>Currently, MasterNode, CoreNode, and TaskNode are supported. |
| name                 | String                    | VM name                                                                 |
| status               | String                    | Current VM state                                                        |
| resource_id          | String                    | Node resource ID                                                        |
| mem                  | String                    | Memory                                                                  |
| cpu                  | String                    | Number of CPU cores                                                     |
| root_volume_size     | String                    | OS disk capacity                                                        |
| data_volume_type     | String                    | Data disk type                                                          |
| data_volume_size     | Integer                   | Data disk capacity                                                      |
| data_volume_count    | Integer                   | Number of data disks                                                    |
| node_group_name      | String                    | Node group                                                              |



**Table 7-42** TagPlain parameters

| Parameter | Type   | Description                                                                                      |
|-----------|--------|--------------------------------------------------------------------------------------------------|
| key       | String | The tag key. A tag key cannot contain special characters (=*<>\, /) or start or end with spaces. |
| value     | String | The value. A tag value cannot contain special characters (=*<>\, /) or start or end with spaces. |

## Example Request

```
GET /v1.1/{project_id}/clusters/{cluster_id}/hosts
```

## Example Response

**Status code: 200**

Querying the host list is successful.

```
{
  "total": "5",
  "hosts": [ {
    "id": "063d1d47-ae91-4a48-840c-b3cfe4efbcf0",
    "name": "a78e161c-d14f-4b68-8c2d-0219920ce844_node_core_IQhiC",
    "ip": "192.168.0.169",
    "availability_zone_id": null,
    "tags": null,
    "status": "ACTIVE",
    "resource_id": "95c1eabc-ed1d-4037-97d1-62f0587790c7",
    "flavor": "c2.2xlarge.linux.mrs",
    "type": "Core",
    "mem": "16384",
    "cpu": "8",
    "root_volume_size": "480",
    "data_volume_type": "SATA",
    "data_volume_size": "600",
    "data_volume_count": "1"
    "node_group_name": "node_group_1"
  }, {
    "id": "dc5c6208-faa2-4727-a65a-2b1ce235d350",
    "name": "a78e161c-d14f-4b68-8c2d-0219920ce844_node_master1_ASzkl",
    "ip": "192.168.0.156",
    "availability_zone_id": null,
    "tags": null,
    "status": "ACTIVE",
    "resource_id": "95c1eabc-ed1d-4037-97d1-62f0587790c7",
    "flavor": "c2.4xlarge.linux.mrs",
    "type": "Master",
    "mem": "32768",
    "cpu": "16",
    "root_volume_size": "480",
    "data_volume_type": "SATA",
    "data_volume_size": "600",
    "data_volume_count": "1"
    "node_group_name": "master_node_default_group"
  }, {
    "id": "c0ce793d-848b-448a-835b-ea0cac534b09",
    "name": "a78e161c-d14f-4b68-8c2d-0219920ce844_node_core_ANnRN",
```

```

"ip" : "192.168.0.243",
"availability_zone_id" : null,
"tags" : null,
"status" : "ACTIVE",
"resource_id" : "95c1eabc-ed1d-4037-97d1-62f0587790c7",
"flavor" : "c2.2xlarge.linux.mrs",
"type" : "Core",
"mem" : "16384",
"cpu" : "8",
"root_volume_size" : "480",
"data_volume_type" : "SATA",
"data_volume_size" : "600",
"data_volume_count" : "1"
"node_group_name" : "node_group_1"
}, {
"id" : "95c23e43-ef6e-4732-b6ed-a5f1c7779fae",
"name" : "a78e161c-d14f-4b68-8c2d-0219920ce844_node_core_uRRiA",
"ip" : "192.168.0.126",
"availability_zone_id" : null,
"tags" : null,
"status" : "ACTIVE",
"resource_id" : "95c1eabc-ed1d-4037-97d1-62f0587790c7",
"flavor" : "c2.2xlarge.linux.mrs",
"type" : "Core",
"mem" : "16384",
"cpu" : "8",
"root_volume_size" : "480",
"data_volume_type" : "SATA",
"data_volume_size" : "600",
"data_volume_count" : "1"
"node_group_name" : "node_group_1"
}, {
"id" : "63bdbf75-1133-4a94-8c27-1fa12c8b9e70",
"name" : "a78e161c-d14f-4b68-8c2d-0219920ce844_node_master2_StqFu",
"ip" : "192.168.0.22",
"availability_zone_id" : null,
"tags" : null,
"status" : "ACTIVE",
"resource_id" : "95c1eabc-ed1d-4037-97d1-62f0587790c7",
"flavor" : "c2.4xlarge.linux.mrs",
"type" : "Master",
"mem" : "32768",
"cpu" : "16",
"root_volume_size" : "480",
"data_volume_type" : "SATA",
"data_volume_size" : "600",
"data_volume_count" : "1"
"node_group_name" : "master_node_default_group"
} ]
}

```

## Status Codes

[Table 7-43](#) describes the status code.

**Table 7-43** Status code

| Status Code | Description                     |
|-------------|---------------------------------|
| 200         | The host list has been queried. |

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 7.1.6 Terminating a Cluster

### Function

This API is used to terminate a cluster after data processing and analysis are complete or the cluster is abnormal. This API is compatible with Sahara.

Clusters in any of the following states cannot be terminated:

- scaling-out
- scaling-in
- starting
- terminating
- terminated
- failed

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
DELETE /v1.1/{project\_id}/clusters/{cluster\_id}
- Parameter description

**Table 7-44** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                             |
|------------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .                |
| cluster_id | Yes       | String | Cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining the MRS Cluster Information</a> . |

## Request Parameters

None

## Response Parameters

None

## Example Request

```
DELETE /v1.1/{project_id}/clusters/{cluster_id}
```

## Example Response

None

## Status Codes

[Table 7-45](#) describes the status codes.

**Table 7-45** Status codes

| Status Code | Description                      |
|-------------|----------------------------------|
| 200         | The request is executed.         |
| 204         | The cluster has been terminated. |

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

# 7.2 Auto Scaling APIs

## 7.2.1 Configuring an Auto Scaling Rule

### Function

This API is used to configure auto scaling rules.

The API used for cluster creation and job execution can also be used to create an auto scaling rule.

### Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- Format  
POST /v1.1/{project\_id}/autoscaling-policy/{cluster\_id}
- Parameter description

**Table 7-46** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                            |
|------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .               |
| cluster_id | Yes       | String | Cluster ID For details about how to obtain the cluster ID, see <a href="#">Obtaining the MRS Cluster Information</a> . |

## Request Parameters

**Table 7-47** Request body parameters

| Parameter           | Mandatory | Type                     | Description                                                                                                                                                                   |
|---------------------|-----------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_group          | Yes       | String                   | Type of the node to which an auto scaling rule applies. Currently, only Task nodes support auto scaling rules, that is, the request value is <b>task_node_default_group</b> . |
| auto_scaling_policy | Yes       | AutoScalingPolicy object | The auto scaling policy. <a href="#">Table 7-48</a> describes these parameters.                                                                                               |

**Table 7-48** AutoScalingPolicy

| Parameter           | Mandatory | Type    | Description                              |
|---------------------|-----------|---------|------------------------------------------|
| auto_scaling_enable | Yes       | Boolean | Whether to enable the auto scaling rule. |

| Parameter       | Mandatory | Type    | Description                                                                                                                                                                                                                               |
|-----------------|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| min_capacity    | Yes       | Integer | Minimum number of nodes left in the node group.<br>Value range: 0 to 500                                                                                                                                                                  |
| max_capacity    | Yes       | Integer | Maximum number of nodes in the node group.<br>Value range: 0 to 500                                                                                                                                                                       |
| resources_plans | No        | List    | Resource plan list. For details, see <a href="#">Table 7-49</a> . If this parameter is left blank, the resource plan is disabled.<br><br>When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured. |
| exec_scripts    | No        | List    | List of custom scaling automation scripts. For details, see <a href="#">Table 7-50</a> . If this parameter is left blank, a hook script is disabled.                                                                                      |
| rules           | No        | List    | List of auto scaling rules. For details, see <a href="#">Table 7-51</a> .<br><br>When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured.                                                         |

**Table 7-49** ResourcesPlan

| Parameter   | Mandatory | Type   | Description                                                                                                                                                                                         |
|-------------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| period_type | Yes       | String | Cycle type of a resource plan. Currently, only the following cycle type is supported:<br>daily                                                                                                      |
| start_time  | Yes       | String | The start time of a resource plan. The value is in the format of <b>hour:minute</b> , indicating that the time ranges from 00:00 to 23:59.                                                          |
| end_time    | Yes       | String | End time of a resource plan. The value is in the same format as that of <b>start_time</b> . The interval between <b>end_time</b> and <b>start_time</b> must be greater than or equal to 30 minutes. |

| Parameter    | Mandatory | Type    | Description                                                                                        |
|--------------|-----------|---------|----------------------------------------------------------------------------------------------------|
| min_capacity | Yes       | Integer | Minimum number of the preserved nodes in a node group in a resource plan.<br>Value range: 0 to 500 |
| max_capacity | Yes       | Integer | Maximum number of the preserved nodes in a node group in a resource plan.<br>Value range: 0 to 500 |

**Table 7-50** ScaleScript

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                 |
|-----------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | Yes       | String | Name of a custom automation script. It must be unique in a same cluster.<br>The value can contain only digits, letters, spaces, hyphens (-), and underscores (_) and must not start with a space.<br>The value can contain 1 to 64 characters.                                                                                                              |
| uri       | Yes       | String | Path of a custom automation script. Set this parameter to an OBS bucket path or a local VM path. <ul style="list-style-type: none"> <li>OBS bucket path: Enter a script path manually. for example, <b>s3a://XXX/scale.sh</b>.</li> <li>Local VM path: Enter a script path. The script path must start with a slash (/) and end with <b>.sh</b>.</li> </ul> |

| Parameter     | Mandatory | Type             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------|-----------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| parameters    | No        | String           | <p>Parameters of a custom automation script.</p> <ul style="list-style-type: none"> <li>Multiple parameters are separated by space.</li> <li>The following predefined system parameters can be transferred: <ul style="list-style-type: none"> <li><i>\${mrs_scale_node_num}</i>: Number of the nodes to be added or removed</li> <li><i>\${mrs_scale_type}</i>: Scaling type. The value can be <b>scale_out</b> or <b>scale_in</b>.</li> <li><i>\${mrs_scale_node_hostnames}</i>: Host names of the nodes to be added or removed</li> <li><i>\${mrs_scale_node_ips}</i>: IP addresses of the nodes to be added or removed</li> <li><i>\${mrs_scale_rule_name}</i>: Name of the rule that triggers auto scaling</li> </ul> </li> <li>Other user-defined parameters are used in the same way as those of common shell scripts. Parameters are separated by space.</li> </ul> |
| nodes         | Yes       | Array of strings | Type of a node where the custom automation script is executed. The node type can be Master, Core, or Task.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| active_master | No        | Boolean          | <p>Whether the custom automation script runs only on the active Master node.</p> <p>The default value is <b>false</b>, indicating that the custom automation script can run on all Master nodes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



| Parameter    | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| action_stage | Yes       | String | Time when a script is executed.<br>The following four options are supported: <ul style="list-style-type: none"> <li>• <b>before_scale_out</b>: before scale-out</li> <li>• <b>before_scale_in</b>: before scale-in</li> <li>• <b>after_scale_out</b>: after scale-out</li> <li>• <b>after_scale_in</b>: after scale-in</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                             |
| fail_action  | Yes       | String | Whether to continue to execute subsequent scripts and create a cluster after the custom automation script fails to be executed. <ul style="list-style-type: none"> <li>• <b>continue</b>: Continue to execute subsequent scripts.</li> <li>• <b>errorout</b>: Stop the action.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>• You are advised to set this parameter to <b>continue</b> in the commissioning phase so that the cluster can continue to be installed and started no matter whether the custom automation script is executed successfully.</li> <li>• The scale-in operation cannot be undone. Therefore, <b>fail_action</b> must be set to <b>continue</b> for the scripts that are executed after scale-in.</li> </ul> |

**Table 7-51** Rule

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                            |
|-----------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | Yes       | String | Name of an auto scaling rule.<br>A cluster name can contain only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.<br>Rule names must be unique in a node group. |

| Parameter          | Mandatory | Type           | Description                                                                                                                                                                                      |
|--------------------|-----------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| description        | No        | String         | Description about an auto scaling rule.<br>It contains a maximum of 1,024 characters.                                                                                                            |
| adjustment_type    | Yes       | String         | Auto scaling rule adjustment type.<br>Possible values: <ul style="list-style-type: none"> <li>• <b>scale_out</b>: cluster scale-out</li> <li>• <b>scale_in</b>: cluster scale-in</li> </ul>      |
| cool_down_minutes  | Yes       | Integer        | Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.<br>Value range: 0 to 10,080. One week is equal to 10,080 minutes. |
| scaling_adjustment | Yes       | Integer        | Number of nodes that can be adjusted once.<br>Value range: 1 to 100                                                                                                                              |
| trigger            | Yes       | Trigger object | Condition for triggering a rule. For details, see <a href="#">Table 7-52</a> .                                                                                                                   |

**Table 7-52** Trigger

| Parameter    | Mandatory | Type   | Description                                                                                                                                                                                                                                                 |
|--------------|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| metric_name  | Yes       | String | Metric name.<br>This triggering condition makes a judgment according to the value of the metric.<br>A metric name contains a maximum of 64 characters.<br>For details about metric names, see <a href="#">Configuring Auto Scaling for an MRS Cluster</a> . |
| metric_value | Yes       | String | Metric threshold to trigger a rule<br>The parameter value can only be an integer or number with two decimal places. The value type and range must correspond to the <b>metric_name</b> .                                                                    |

| Parameter           | Mandatory | Type    | Description                                                                                                                                                                                                                                                    |
|---------------------|-----------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| comparison_operator | No        | String  | Metric judgment logic operator.<br>Possible values: <ul style="list-style-type: none"> <li>• <b>LT</b>: less than</li> <li>• <b>GT</b>: greater than</li> <li>• <b>LTOE</b>: less than or equal to</li> <li>• <b>GTOE</b>: greater than or equal to</li> </ul> |
| evaluation_periods  | Yes       | Integer | Number of consecutive five-minute periods, during which a metric threshold is reached<br>Value range: 1 to 288                                                                                                                                                 |

## Response Parameters

None

## Example Request

Configure an auto scaling policy for the MRS cluster.

POST https://{endpoint}/v1.1/{project\_id}/autoscaling-policy/{cluster\_id}

```
{
  "node_group": "task_node_analysis_group",
  "auto_scaling_policy": {
    "auto_scaling_enable": "true",
    "min_capacity": "1",
    "max_capacity": "3",
    "resources_plans": [ {
      "period_type": "daily",
      "start_time": "9:50",
      "end_time": "10:20",
      "min_capacity": "2",
      "max_capacity": "3"
    }, {
      "period_type": "daily",
      "start_time": "10:20",
      "end_time": "12:30",
      "min_capacity": "0",
      "max_capacity": "2"
    }
  ],
  "exec_scripts": [ {
    "name": "before_scale_out",
    "uri": "s3a://XXX/zeppelin_install.sh",
    "parameters": "${mrs_scale_node_num} ${mrs_scale_type} xxx",
    "nodes": [ "master_node_default_group", "core_node_analysis_group", "task_node_analysis_group" ],
    "active_master": "true",
    "action_stage": "before_scale_out",
    "fail_action": "continue"
  }, {
    "name": "after_scale_out",
    "uri": "s3a://XXX/storm_rebalance.sh",
    "parameters": "${mrs_scale_node_hostnames} ${mrs_scale_node_ips}",
    "nodes": [ "master_node_default_group", "core_node_analysis_group", "task_node_analysis_group" ],
    "active_master": "true",
    "action_stage": "after_scale_out",
  }
  ]
}
```

```

    "fail_action": "continue"
  } ],
  "rules": [ {
    "name": "default-expand-1",
    "adjustment_type": "scale_out",
    "cool_down_minutes": "5",
    "scaling_adjustment": "1",
    "trigger": {
      "metric_name": "YARNMemoryAvailablePercentage",
      "metric_value": "25",
      "comparison_operator": "LT",
      "evaluation_periods": "10"
    }
  }, {
    "name": "default-shrink-1",
    "adjustment_type": "scale_in",
    "cool_down_minutes": "5",
    "scaling_adjustment": "1",
    "trigger": {
      "metric_name": "YARNMemoryAvailablePercentage",
      "metric_value": "70",
      "comparison_operator": "GT",
      "evaluation_periods": "10"
    }
  } ]
}

```

 **NOTE**

A new auto scaling rule will overwrite the auto scaling rule saved in the original database. If you want to modify the original rule, query the original rule first, modify the rule, and submit a modification task. For details, see [Querying Cluster Details](#).

## Example Response

**Status code: 200**

The operation is successful.

```

{
  "result": "succeeded"
}

```

## Status Codes

[Table 7-53](#) describes the status code.

**Table 7-53** Status code

| Status Code | Description                   |
|-------------|-------------------------------|
| 200         | The cluster has been created. |

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 7.3 Tag Management APIs

### 7.3.1 Adding Tags to a Specified Cluster

#### Function

This API is used to add tags to a specified cluster.

A cluster has a maximum of 20 tags. This API is idempotent. If a tag to be created has the same key as an existing tag, the tag will overwrite the existing one.

#### Constraints

None

#### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

#### URI

- Format  
POST /v1.1/{project\_id}/clusters/{cluster\_id}/tags
- Parameter description

**Table 7-54** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                            |
|------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .               |
| cluster_id | Yes       | String | Cluster ID For details about how to obtain the cluster ID, see <a href="#">Obtaining the MRS Cluster Information</a> . |

#### Request Parameters

**Table 7-55** Request body parameter

| Parameter | Mandatory | Type                       | Description      |
|-----------|-----------|----------------------------|------------------|
| tag       | Yes       | <a href="#">Tag</a> object | The cluster tag. |

**Table 7-56** tags parameters

| Parameter | Mandatory | Type   | Description                                                                                       |
|-----------|-----------|--------|---------------------------------------------------------------------------------------------------|
| key       | Yes       | String | Key. A tag key cannot contain special characters (=*<>\\ /) or start or end with a space.         |
| value     | Yes       | String | Tag value. A tag value cannot contain special characters (=*<>\\ /) or start or end with a space. |

## Response Parameters

None

## Example Request

Add a tag to a specified cluster.

```
{
  "tag":
  {
    "key": "DEV",
    "value": "DEV1"
  }
}
```

## Example response

None

## Status Codes

[Table 7-57](#) describes the status code.

**Table 7-57** Status code

| Status Code | Description                  |
|-------------|------------------------------|
| 204         | The operation is successful. |

## Error Codes

See [Error Codes](#).

## 7.3.2 Querying Tags of a Specified Cluster

### Function

This API is used to query tags of a specified cluster.

## Constraints

None

## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

GET /v1.1/{project\_id}/clusters/{cluster\_id}/tags

**Table 7-58** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                            |
|------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .               |
| cluster_id | Yes       | String | Cluster ID For details about how to obtain the cluster ID, see <a href="#">Obtaining the MRS Cluster Information</a> . |

## Request Parameters

None

## Response Parameters

**Status code: 200**

**Table 7-59** Response parameter

| Parameter | Type                      | Description                                             |
|-----------|---------------------------|---------------------------------------------------------|
| tags      | Array of TagPlain objects | Tag list. For details, see <a href="#">Table 7-60</a> . |

**Table 7-60** TagPlain parameters

| Parameter | Type   | Description                                                                                |
|-----------|--------|--------------------------------------------------------------------------------------------|
| key       | String | Key. A tag key cannot contain special characters (=*<>\\, /) or start or end with a space. |

| Parameter | Type   | Description                                                                                       |
|-----------|--------|---------------------------------------------------------------------------------------------------|
| value     | String | Tag value. A tag value cannot contain special characters (=*<>\, /) or start or end with a space. |

## Example Request

```
GET /v1.1/{project_id}/clusters/{cluster_id}/tags
```

## Example Response

**Status code: 200**

The operation is successful.

```
{
  "tags": [ {
    "key": "key1",
    "value": "value1"
  }, {
    "key": "key2",
    "value": "value3"
  } ]
}
```

## Status Codes

[Table 7-61](#) describes the status code.

**Table 7-61** Status code

| Status Code | Description                  |
|-------------|------------------------------|
| 200         | The operation is successful. |

## Error Codes

See [Error Codes](#).

## 7.3.3 Deleting Tags from a Specified Cluster

### Function

This API is used to delete tags from a specified cluster.

### Constraints

None



## Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

## URI

- Format  
DELETE /v1.1/{project\_id}/clusters/{cluster\_id}/tags/{key}
- Parameter description

**Table 7-62** URI parameters

| Parameter  | Mandator<br>y | Type   | Description                                                                                              |
|------------|---------------|--------|----------------------------------------------------------------------------------------------------------|
| project_id | Yes           | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |
| cluster_id | Yes           | String | Cluster ID For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .  |
| key        | Yes           | String | The tag key.                                                                                             |

## Request Parameters

None

## Response Parameters

None

## Example Request

```
DELETE /v1.1/{project_id}/clusters/{cluster_id}/tags/{key}
```

## Example Response

None

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 7.3.4 Adding Tags to a Cluster in Batches

### Function

This API is used to add tags to a specified cluster in batches.

A cluster can have a maximum of 20 tags.

This API is idempotent. If a tag to be created has the same key as an existing tag in a cluster, the tag will overwrite the existing one.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

POST /v1.1/{project\_id}/clusters/{cluster\_id}/tags/action

**Table 7-63** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                              |
|------------|-----------|--------|----------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |
| cluster_id | Yes       | String | Cluster ID For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> .  |

### Request Parameters

**Table 7-64** Request parameters

| Parameter | Mandatory | Type                 | Description                                                                             |
|-----------|-----------|----------------------|-----------------------------------------------------------------------------------------|
| action    | Yes       | String               | Operation to be performed. The value can be set to <b>create</b> or <b>delete</b> only. |
| tags      | Yes       | Array of tag objects | Tag list. For details about the parameter, see <a href="#">Table 7-65</a> .             |

**Table 7-65** tags parameters

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                |
|-----------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| key       | Yes       | String | <p>The tag key.</p> <ul style="list-style-type: none"> <li>The tag key cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\\, /).</li> <li>It can contain a maximum of 128 Unicode characters and cannot be an empty string.</li> <li>The tag key of a resource must be unique.</li> </ul> |
| value     | Yes       | String | <p>The value.</p> <ul style="list-style-type: none"> <li>The tag value cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\\, /).</li> <li>The value can contain a maximum of 255 Unicode characters and can be an empty string.</li> </ul>                                                |

## Response Parameters

None

## Example Request

Add tags to a cluster in batches.

```
POST /v1.1/{project_id}/clusters/{cluster_id}/tags/action
{
  "action": "create",
  "tags": [{
    "key": "DEV1",
    "value": "DEV1"
  }, {
    "key": "DEV2",
    "value": "DEV2"
  }]
}
```

## Example Response

None

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 7.3.5 Deleting Tags from a Cluster in Batches

### Function

This API is used to delete tags from a specified cluster in batches.

A cluster can have a maximum of 20 tags.

This API is idempotent.

- When tags are being deleted and some tags do not exist, the operation is considered to be successful by default. The character set of the tags will not be checked. A key and a value can respectively contain up to 128 and 255 Unicode characters. The tags structure cannot be missing. The key cannot be empty or an empty string.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

POST /v1.1/{project\_id}/clusters/{cluster\_id}/tags/action

**Table 7-66** URI parameters

| Parameter  | Mandatory | Type   | Description                                                                                                  |
|------------|-----------|--------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |
| cluster_id | Yes       | String | The cluster ID. For details about how to obtain the cluster ID, see <a href="#">Obtaining a Cluster ID</a> . |

## Request Parameters

**Table 7-67** Request body parameters

| Parameter | Mandatory | Type                 | Description                                                                      |
|-----------|-----------|----------------------|----------------------------------------------------------------------------------|
| action    | Yes       | String               | Operation to be performed. The value can be <b>create</b> or <b>delete</b> only. |
| tags      | Yes       | Array of tag objects | The tag list. For details about the parameter, see <a href="#">Table 7-68</a> .  |

**Table 7-68** Tag

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                               |
|-----------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| key       | Yes       | String | <p>The tag key.</p> <ul style="list-style-type: none"> <li>The tag key cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\, /).</li> <li>It can contain a maximum of 128 Unicode characters and cannot be an empty string.</li> <li>The tag key of a resource must be unique.</li> </ul> |
| value     | Yes       | String | <p>The value.</p> <ul style="list-style-type: none"> <li>The tag value cannot start or end with spaces or contain non-printable ASCII characters (0–31) and special characters (=*&lt;&gt;\, /).</li> <li>The value can contain a maximum of 255 Unicode characters and can be an empty string.</li> </ul>                                                |

## Response Parameters

None

## Example Request

Delete tags from a cluster in batches.

```
POST /v1.1/{project_id}/clusters/{cluster_id}/tags/action
{
  "action": "delete",
  "tags": [ {
    "key": "DEV1",
    "value": "DEV1"
  }, {
    "key": "DEV2",
    "value": "DEV2"
  } ]
}
```

### Example Response

None

### Status Codes

See [Status Codes](#).

### Error Codes

See [Error Codes](#).

## 7.3.6 Querying All Tags

### Function

This API is used to query all tags of a specified region.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

GET /v1.1/{project\_id}/clusters/tags

**Table 7-69** URI parameter

| Parameter  | Mandatory | Type   | Description                                                                                              |
|------------|-----------|--------|----------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

## Request Parameters

None

## Response Parameters

Status code: 200

**Table 7-70** Response parameter

| Parameter | Type                               | Description                                             |
|-----------|------------------------------------|---------------------------------------------------------|
| tags      | Array of TagWithMultiValue objects | Tag list. For details, see <a href="#">Table 7-71</a> . |

**Table 7-71** tags parameters

| Parameter | Type             | Description                                                                                       |
|-----------|------------------|---------------------------------------------------------------------------------------------------|
| key       | String           | Tag key. A tag key cannot contain special characters (=*<>\,/ ) or start or end with a space.     |
| values    | Array of strings | Tag value. A tag value cannot contain special characters (=*<>\,/ ) or start or end with a space. |

## Example Request

Query all tags.

```
GET https://{endpoint}/v1.1/{project_id}/clusters/tags
```

## Example Response

Status code: 200

The operation is successful.

```
{
  "tags": [ {
    "key": "key1",
    "values": [ "value1", "value2" ]
  }, {
    "key": "key2",
    "values": [ "value1", "value2" ]
  } ]
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

## 7.3.7 Querying a List of Clusters with Specified Tags

### Function

This API is used to filter clusters by tag.

By default, clusters and tags are sorted in descending order of creation time.

### Constraints

None

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and provide the sample SDK code debugging.

### URI

- Format  
POST /v1.1/{project\_id}/clusters/resource\_instances/action
- Parameter description

**Table 7-72** URI parameter

| Parameter  | Mandatory | Type   | Description                                                                                                  |
|------------|-----------|--------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | String | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |



## Request Parameters

**Table 7-73** Request body parameters

| Parameter | Mandatory | Type                               | Description                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------|-----------|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tags      | No        | Array of TagWithMultiValue objects | The return result contains resources corresponding to all tags in this parameter. This parameter contains a maximum of 10 keys, and each key contains a maximum of 10 values. The structure body cannot be missing, and the key cannot be left blank or set to an empty string. For details about the parameters, see <a href="#">Table 7-74</a> .                                                                 |
| tags_any  | No        | Array of TagWithMultiValue objects | The return result contains resources corresponding to any tag in this parameter. This parameter contains a maximum of 10 keys, and each key contains a maximum of 10 values. The structure body cannot be missing, and the key cannot be left blank or set to an empty string. Keys must be unique and values of a key must be unique. For details about the parameters, see <a href="#">Table 7-74</a> .          |
| not_tags  | No        | Array of TagWithMultiValue objects | The return result does not contain resources corresponding to all tags in this parameter. This parameter contains a maximum of 10 keys, and each key contains a maximum of 10 values. The structure body cannot be missing, and the key cannot be left blank or set to an empty string. Keys must be unique and values of a key must be unique. For details about the parameters, see <a href="#">Table 7-74</a> . |

| Parameter    | Mandatory | Type                               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------|-----------|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| not_tags_any | No        | Array of TagWithMultiValue objects | The return result does not contain resources corresponding to any tag in this parameter. This parameter contains a maximum of 10 keys, and each key contains a maximum of 10 values. The structure body cannot be missing, and the key cannot be left blank or set to an empty string. Keys must be unique and values of a key must be unique. For details about the parameters, see <a href="#">Table 7-74</a> .                                                                                                                                                   |
| limit        | No        | Integer                            | Number of records. This parameter is not available when <b>action</b> is <b>count</b> . The default value is <b>1000</b> when <b>action</b> is <b>filter</b> . The maximum value is <b>1000</b> , and the minimum value is <b>1</b> . The value cannot be a negative number.                                                                                                                                                                                                                                                                                        |
| offset       | No        | Integer                            | Index position. The query starts from the next piece of data specified by the <b>offset</b> parameter. This parameter is not required when you query data on the first page. The value in the response body returned for querying data on the previous page will be included in this parameter for querying data on subsequent pages. This parameter is not available when <b>action</b> is set to <b>count</b> . If <b>action</b> is set to <b>filter</b> , the value must be a number, and the default value is <b>0</b> . The value cannot be a negative number. |
| action       | Yes       | String                             | Operation to be performed. The value can be <b>filter</b> or <b>count</b> .<br>The value <b>filter</b> indicates pagination query. The value <b>count</b> indicates that the total number of query results meeting the search criteria will be returned.                                                                                                                                                                                                                                                                                                            |

| Parameter | Mandatory | Type                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------|-----------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| matches   | No        | Array of Match objects | <p>Search field. <b>key</b> indicates the field to be matched, for example, <b>resource_name</b>. <b>value</b> indicates the matched value. This field is a fixed dictionary value.</p> <p>Determine whether fuzzy match is required based on different fields. For example, if <b>key</b> is <b>resource_name</b>, fuzzy search is used by default. If <b>value</b> is an empty string, exact match is used. For details about the parameters, see <a href="#">Table 7-75</a>.</p> |

**Table 7-74** TagWithMultiValue

| Parameter | Mandatory | Type             | Description                                                                                                               |
|-----------|-----------|------------------|---------------------------------------------------------------------------------------------------------------------------|
| key       | Yes       | String           | Key. A tag key cannot contain special characters (=*<>\, /) or start or end with spaces.                                  |
| values    | No        | Array of strings | <p>List of values.</p> <p>A tag value cannot contain special characters (=*&lt;&gt;\, /) or start or end with spaces.</p> |

**Table 7-75** Match

| Parameter | Mandatory | Type   | Description                                                                                                              |
|-----------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------|
| key       | No        | String | Key. Currently, only <b>resource_name</b> is available, indicating the cluster name. Other keys will be supported later. |
| value     | No        | String | Value. A value contains a maximum of 64 Unicode characters.                                                              |

## Response Parameters

**Table 7-76** Response parameters

| Parameter   | Type                         | Description                                                     |
|-------------|------------------------------|-----------------------------------------------------------------|
| resources   | Array of MRSResource objects | Resource details. For details, see <a href="#">Table 7-77</a> . |
| total_count | Integer                      | Total number of resources.                                      |

**Table 7-77** MRSResource

| Parameter       | Type                      | Description                                             |
|-----------------|---------------------------|---------------------------------------------------------|
| resource_detail | String                    | Resource details.                                       |
| resource_id     | String                    | Resource ID.                                            |
| resource_name   | String                    | Resource name.                                          |
| tags            | Array of TagPlain objects | Tag list. For details, see <a href="#">Table 7-78</a> . |

**Table 7-78** TagPlain

| Parameter | Type   | Description                                                                                       |
|-----------|--------|---------------------------------------------------------------------------------------------------|
| key       | String | Key. A tag key cannot contain special characters (=*<>\, /) or start or end with a space.         |
| value     | String | Tag value. A tag value cannot contain special characters (=*<>\, /) or start or end with a space. |

## Example Request

- Query the cluster list when **action** is set to **filter**.  
POST `https://{endpoint}/v1.1/{project_id}/{resource_type}/resource_instances/action`

```
{
  "offset": "100",
  "limit": "100",
  "action": "filter",
  "matches": [ {
    "key": "resource_name",
    "value": "clusterA"
  } ],
  "not_tags": [ {
    "key": "key1",
    "values": [ "value1", "value2" ]
  } ]
}
```

```
    },
    "tags" : [ {
      "key" : "key1",
      "values" : [ "value1", "value2" ]
    } ],
    "tags_any" : [ {
      "key" : "key1",
      "values" : [ "value1", "value2" ]
    } ],
    "not_tags_any" : [ {
      "key" : "key1",
      "values" : [ "value1", "value2" ]
    } ]
  }
}
```

- Query the cluster list when **action** is set to **count**.

POST [https://{endpoint}/v1.1/{project\\_id}/{resource\\_type}/resource\\_instances/action](https://{endpoint}/v1.1/{project_id}/{resource_type}/resource_instances/action)

```
{
  "action" : "count",
  "not_tags" : [ {
    "key" : "key1",
    "values" : [ "value1", "value2" ]
  } ],
  "tags" : [ {
    "key" : "key1",
    "values" : [ "value1", "value2" ]
  }, {
    "key" : "key2",
    "values" : [ "value1", "value2" ]
  } ],
  "tags_any" : [ {
    "key" : "key1",
    "values" : [ "value1", "value2" ]
  } ],
  "not_tags_any" : [ {
    "key" : "key1",
    "values" : [ "value1", "value2" ]
  } ],
  "matches" : [ {
    "key" : "resource_name",
    "value" : "clusterA"
  } ]
}
```

## Example Response

**Status code: 200**

The operation is successful.

```
{
  "resources" : [ {
    "resource_detail" : null,
    "resource_id" : "cdfs_cefs_wesas_12_dsad",
    "resource_name" : "clusterA",
    "tags" : [ {
      "key" : "key1",
      "value" : "value1"
    }, {
      "key" : "key2",
      "value" : "value1"
    } ]
  } ],
  "total_count" : "1000"
}
```

## Status Codes

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

# 7.4 Availability Zones

## 7.4.1 Querying AZ Information

### Function

This API is used to query the availability zone (AZ) ID used when you create a cluster.

### URI

GET /v1.1/{region\_id}/available-zones

**Table 7-79** URI parameters

| Parameter | Mandatory | Type   | Description |
|-----------|-----------|--------|-------------|
| region_id | Yes       | String | Region ID   |

**Table 7-80** Query parameters

| Parameter | Mandatory | Type   | Description                                                                                                   |
|-----------|-----------|--------|---------------------------------------------------------------------------------------------------------------|
| scope     | No        | String | AZ range. The options are as follows: <ul style="list-style-type: none"><li>• Center</li><li>• Edge</li></ul> |

### Request Parameters

None

### Response Parameters

Status code: 200

**Table 7-81** Response body parameter

| Parameter                 | Type                             | Description                                                                          |
|---------------------------|----------------------------------|--------------------------------------------------------------------------------------|
| available_zones           | Array of AvailableZoneV2 objects | AZ list. For details about <b>AvailableZoneV2</b> , see <a href="#">Table 7-82</a> . |
| default_az_code           | String                           | Default AZ code                                                                      |
| support_physical_az_group | Boolean                          | Supported physical AZ group                                                          |

**Table 7-82** AvailableZoneV2

| Parameter   | Type                | Description                                                                                                                                                                                                                     |
|-------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id          | String              | AZ ID                                                                                                                                                                                                                           |
| az_code     | String              | AZ code                                                                                                                                                                                                                         |
| az_name     | String              | AZ name                                                                                                                                                                                                                         |
| az_id       | String              | AZ ID                                                                                                                                                                                                                           |
| status      | String              | AZ status                                                                                                                                                                                                                       |
| region_id   | String              | Region ID                                                                                                                                                                                                                       |
| az_group_id | String              | AZ group ID                                                                                                                                                                                                                     |
| az_type     | String              | AZ type <ul style="list-style-type: none"> <li>• <b>Core</b></li> <li>• <b>Satellite</b></li> <li>• <b>Dedicated</b></li> <li>• <b>Virtual</b></li> <li>• <b>Edge</b></li> <li>• <b>EdgeCental</b>: central and edge</li> </ul> |
| az_tags     | AvailableTag object | AZ tag. For details about <b>AvailableTag</b> , see <a href="#">Table 7-83</a> .                                                                                                                                                |

**Table 7-83** AvailableTag

| Parameter | Type   | Description                                                   |
|-----------|--------|---------------------------------------------------------------|
| mode      | String | AZ mode. The value can be <b>dedicated</b> or <b>shared</b> . |
| alias     | String | Alias of an AZ                                                |

| Parameter           | Type   | Description                                                         |
|---------------------|--------|---------------------------------------------------------------------|
| public_border_group | String | Group to which the AZ belongs. The default value is <b>center</b> . |

## Example Request

None

## Example Response

**Status code: 200**

Query AZ information.

```
{
  "available_zones" : [ {
    "id" : "xxx",
    "az_code" : "xxx",
    "az_name" : "AZ 1",
    "az_id" : "xxx",
    "status" : "Running",
    "region_id" : "xxx",
    "az_type" : "Dedicated",
    "az_group_id" : "",
    "az_tags" : {
      "mode" : null,
      "alias" : null,
      "public_border_group" : "center"
    }
  }, {
    "id" : "xxx",
    "az_code" : "xxx",
    "az_name" : "AZ 2",
    "az_id" : "xxx",
    "status" : "Running",
    "region_id" : "xxx",
    "az_type" : "Dedicated",
    "az_tags" : {
      "mode" : null,
      "alias" : null,
      "public_border_group" : "center"
    }
  }, {
    "id" : "xxx",
    "az_code" : "xxx",
    "az_name" : "AZ 3",
    "az_id" : "xxx",
    "status" : "Running",
    "region_id" : "xxx",
    "az_type" : "Dedicated",
    "az_tags" : {
      "mode" : null,
      "alias" : null,
      "public_border_group" : "center"
    }
  }
],
  "default_az_code" : "xxx",
  "support_physical_az_group" : true
}
```



## Status Code

See [Table 7-84](#).

**Table 7-84** Status Code

| Status Code | Description    |
|-------------|----------------|
| 200         | AZ information |

See [Status Codes](#).

## Error Codes

See [Error Codes](#).

# 7.5 Version Metadata

## 7.5.1 Querying the Metadata of a Cluster Version

### Function

This API is used to query the metadata of a cluster version. If the cluster ID is specified, the latest metadata of the cluster that has been patched can be queried.

### Debugging

You can debug this API in API Explorer. Automatic authentication is supported. API Explorer can automatically generate sample SDK code and support sample SDK code debugging.

### URI

GET /v1.1/{project\_id}/metadata/versions/{version\_name}

**Table 7-85** URI parameters

| Parameter    | Mandatory | Type   | Description                                                                                                                                               |
|--------------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| project_id   | Yes       | String | Project ID                                                                                                                                                |
| version_name | Yes       | String | The cluster version, for example: MRS 3.1.0. If the request client does not support automatic escape, escape the space to %20, for example, MRS %203.1.0. |

**Table 7-86** Query parameters

| Parameter  | Mandatory | Type   | Description                                                                                                            |
|------------|-----------|--------|------------------------------------------------------------------------------------------------------------------------|
| cluster_id | No        | String | Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. |

## Request Parameters

None

## Response Parameters

**Status code: 200**

**Table 7-87** Response body parameter

| Parameter            | Type                              | Description                                                                                  |
|----------------------|-----------------------------------|----------------------------------------------------------------------------------------------|
| other                | Map<String, Object>               | Others                                                                                       |
| name                 | String                            | Image version                                                                                |
| template_type        | String                            | Template type                                                                                |
| image_id             | String                            | Image ID                                                                                     |
| status               | String                            | Version status                                                                               |
| features             | Array of strings                  | Features                                                                                     |
| cluster_types        | Array of strings                  | Cluster types                                                                                |
| version_type         | String                            | Version type                                                                                 |
| components           | Array of VersionComponent objects | Components. For details, see <a href="#">Table 7-88</a> .                                    |
| resource_requirement | Array of strings                  | Description of required resources, such as IP addresses                                      |
| constraints          | VersionConstraint object          | Version restrictions. For details about the parameters, see <a href="#">Table 7-90</a> .     |
| flavors              | FlavorLists object                | Specifications. For details, see <a href="#">Table 7-93</a> .                                |
| role_deploy_meta     | Array of RoleDeployMeta objects   | Role deployment policy of component instances. For details, see <a href="#">Table 7-94</a> . |

**Table 7-88** VersionComponent

| Parameter               | Type                                         | Description                                                         |
|-------------------------|----------------------------------------------|---------------------------------------------------------------------|
| other                   | Map<String, Object>                          | Others                                                              |
| name                    | String                                       | Component name                                                      |
| version                 | String                                       | Supported versions                                                  |
| depend_on               | Array of strings                             | Component dependencies                                              |
| description             | String                                       | Component description                                               |
| available_cluster_types | Array of strings                             | Cluster types that support the component                            |
| external_data_sources   | Array of ComponentExternalDatasource objects | External data source. For details, see <a href="#">Table 7-89</a> . |
| resource_requirement    | Array of strings                             | Description of required resources, such as IP addresses             |
| valid_roles             | Array of strings                             | Valid roles                                                         |
| visible                 | Boolean                                      | Whether the component is visible                                    |
| children_components     | Array of strings                             | Sub-components                                                      |
| multi_az_support_status | String                                       | Whether multi-AZ is supported                                       |

**Table 7-89** ComponentExternalDatasource

| Parameter | Type             | Description                      |
|-----------|------------------|----------------------------------|
| name      | String           | Name of the external data source |
| types     | Array of strings | Type of the external data source |

**Table 7-90** VersionConstraint

| Parameter       | Type                   | Description                                                      |
|-----------------|------------------------|------------------------------------------------------------------|
| other           | Map<String, Object>    | Other restrictions                                               |
| node_constraint | NodeConstraints object | Node restrictions. For details, see <a href="#">Table 7-91</a> . |

| Parameter                             | Type             | Description                                                            |
|---------------------------------------|------------------|------------------------------------------------------------------------|
| safe_mode_kerberos_exclude_components | Array of strings | Excluded components in security mode (Kerberos authentication enabled) |

**Table 7-91** NodeConstraints

| Parameter       | Type                  | Description                                                                                                                 |
|-----------------|-----------------------|-----------------------------------------------------------------------------------------------------------------------------|
| other           | Map<String, Object>   | Restrictions on other nodes, including the number of nodes and disks                                                        |
| master          | NodeConstraint object | Master node restrictions. For details, see <a href="#">Table 7-92</a> .                                                     |
| core            | NodeConstraint object | Core node restrictions. For details, see <a href="#">Table 7-92</a> .                                                       |
| task            | NodeConstraint object | Task node restrictions. For details, see <a href="#">Table 7-92</a> .                                                       |
| core_separate   | NodeConstraint object | Core node restrictions when management and control nodes are different nodes. For details, see <a href="#">Table 7-92</a> . |
| core_combine    | NodeConstraint object | Core node restrictions when management and control nodes are different nodes. For details, see <a href="#">Table 7-92</a> . |
| task_separate   | NodeConstraint object | Task node restrictions when management and control nodes are different nodes. For details, see <a href="#">Table 7-92</a> . |
| task_combine    | NodeConstraint object | Task node restrictions when management and control nodes are different nodes. For details, see <a href="#">Table 7-92</a> . |
| node_group_task | NodeConstraint object | Task node group restrictions. For details, see <a href="#">Table 7-92</a> .                                                 |

**Table 7-92** NodeConstraint

| Parameter | Type                | Description        |
|-----------|---------------------|--------------------|
| other     | Map<String, Object> | Other restrictions |

| Parameter                  | Type                | Description                  |
|----------------------------|---------------------|------------------------------|
| min_node_num               | Integer             | Minimum number of nodes      |
| max_node_num               | Integer             | Minimum number of nodes      |
| min_core_num               | Map<String,Integer> | Minimum number of cores      |
| min_mem_size               | Map<String,Integer> | Minimum memory               |
| min_disk_size              | Integer             | Minimum disk capacity        |
| max_node_group_num         | Integer             | Maximum node groups          |
| min_data_volume_total_size | Map<String,Integer> | Minimum data volume capacity |
| disk_type_constraint       | Map<String,String>  | Disk type limitation         |
| min_root_disk_size         | Integer             | Minimum root disk size       |

**Table 7-93** FlavorLists

| Parameter | Type             | Description                              |
|-----------|------------------|------------------------------------------|
| master    | Array of strings | Specifications supported by Master nodes |
| core      | Array of strings | Specifications supported by Core nodes   |
| task      | Array of strings | Specifications supported by Task nodes   |

**Table 7-94** RoleDeployMeta

| Parameter       | Type               | Description                         |
|-----------------|--------------------|-------------------------------------|
| other           | Map<String,Object> | Others                              |
| name            | String             | Role name                           |
| code_name       | String             | Role name                           |
| component       | String             | Component to which the role belongs |
| node_preference | String             | Node deployment preference          |
| count           | String             | Maximum number of roles             |

| Parameter                   | Type             | Description                                 |
|-----------------------------|------------------|---------------------------------------------|
| affinity                    | String           | Affinity                                    |
| affinity_target             | String           | Objectives of affinity                      |
| multi_instance              | Integer          | Multi-instance                              |
| role_kind                   | String           | Role type                                   |
| constraints                 | Array of strings | Role scope                                  |
| multi_az_placement          | String           | Multi-AZ deployment                         |
| arbitration_deployment      | Boolean          | Arbitration deployment                      |
| support_elb                 | Boolean          | Whether ELB is supported                    |
| multi_affinity_group_enable | Boolean          | Whether the multi-affinity group is enabled |
| local_disks_anti_affinity   | Boolean          | Local disk anti-affinity                    |
| multi_instance_name_pattern | String           | Multi-instance name pattern                 |
| private_ip                  | String           | Private IP address                          |
| weight                      | String           | Weight                                      |

## Example Request

Obtain the metadata of MRS 3.1.0

```
GET /v1.1/{project_id}/metadata/versions/MRS%203.1.0
```

## Example Response

**Status code: 200**

Version metadata details. Some contents are omitted in this example. For details, call the API.

```
{
  "other" : { },
  "name" : "MRS 3.1.0",
  "template_type" : "clusterFusion",
  "image_id" : "ECS:X86:deda7c89-6069-46d3-af7c-4afc31eec0d9,ECS:ARM:d5f0087b-9fd7-4977-b9c8-d07262f452d6,BMS:ks1ne:bf071a00-d1d6-472f-943b-22f733248f36,BMS:d2:a9a1a279-85c9-4c05-bd0a-4cb786937a09",
  "status" : "GA",
  "features" : [ "anti_affinity", "auto_scaling", "bootstrap_scripts", "bootstrap_support_obs", "custom_config", "custom_topology", "datasource_manager", "decouple_install", "elastic_ip", "fi_cluster", "fi_cluster_v2", "iam_user_sync", "log_collect", "log_collect_v2", "master_scale_up", "mrs_ecs_agency", "multi_disks", "multi_login_mode", "obs_integration", "obs_user_policy", "opensource_port_matrix", "ops_channel", "password_policy_v2", "random_passwd_for_install", "register_dns_server", "resize", "resources_plans",
```

```

"safe_mode", "scale_scripts", "security_job_submit", "smn_alarm", "spark_sql_validation_in_executor",
"support_obs_protocol", "task_node", "unified_management_ui", "sql_execution", "multi_az",
"weak_multi_az_deployment_constraint", "multi_master_scale_up", "ommm_random_pwd", "metric_monitor",
"auth_manager", "cluster_patch", "hw_domain_name", "fi_custom_config", "force_scale_up",
"skip_fi_evs_expand", "detach_node" ],
"cluster_types": [ "analysis", "streaming", "mixed", "custom" ],
"version_type": "basic",
"components": [ { {
  "other": { },
  "name": "Hadoop",
  "version": "3.1.1",
  "depend_on": [ "ZooKeeper", "Ranger" ],
  "description": "A framework that allows for the distributed processing of large data sets across clusters.",
  "available_cluster_types": [ "analysis", "mixed", "custom" ],
  "external_datasources": null,
  "resource_requirement": [ "privatelp:1" ],
  "valid_roles": [ "NameNode", "Zkfc", "JournalNode", "DataNode", "ResourceManager", "NodeManager",
"JobHistoryServer", "TimelineServer", "HttpFS" ],
  "visible": true,
  "children_components": [ "HDFS", "Mapreduce", "Yarn" ],
  "multi_az_support_status": "multi_az_ha"
} }, {
  "other": { },
  "name": "Spark2x",
  "version": "2.4.5",
  "depend_on": [ "HDFS", "Yarn", "Hive", "KrbClient", "KrbServer", "ZooKeeper" ],
  "description": "Apache Spark2x is a fast and general engine based on open source Spark2.x for large-
scale data processing.",
  "available_cluster_types": [ "analysis", "mixed", "custom" ],
  "external_datasources": null,
  "resource_requirement": null,
  "valid_roles": [ "JobHistory2x", "JDBCServer2x", "SparkResource2x", "IndexServer2x" ],
  "visible": true,
  "children_components": [ "Spark2x" ],
  "multi_az_support_status": "multi_az_ha"
} }, {
  "other": { },
  "name": "HBase",
  "version": "2.2.3",
  "depend_on": [ "HDFS", "Yarn", "KrbServer", "ZooKeeper" ],
  "description": "A scalable, distributed database that supports structured data storage for large tables.",
  "available_cluster_types": [ "analysis", "mixed", "custom" ],
  "external_datasources": null,
  "resource_requirement": null,
  "valid_roles": [ "HMaster", "RegionServer", "ThriftServer", "Thrift1Server", "RETServer",
"RegionServer_1" ],
  "visible": true,
  "children_components": [ "HBase" ],
  "multi_az_support_status": "multi_az_ha"
} }, {
  "other": { },
  "name": "Hive",
  "version": "3.1.0",
  "depend_on": [ "HDFS", "Mapreduce", "Yarn", "DBService", "ZooKeeper" ],
  "description": "A data warehouse infrastructure that provides data summarization and ad hoc querying.",
  "available_cluster_types": [ "analysis", "mixed", "custom" ],
  "external_datasources": [ { {
    "name": "hive_metastore",
    "types": [ "RDS_MYSQL", "DLCATALOG" ]
  } },
  "resource_requirement": null,
  "valid_roles": [ "MetaStore", "WebHCat", "HiveServer" ],
  "visible": true,
  "children_components": [ "Hive" ],
  "multi_az_support_status": "multi_az_ha"
} }, {
  "other": { },
  "name": "ZooKeeper",
  "version": null,

```

```
"depend_on" : [ "KrbClient", "KrbServer" ],
"description" : null,
"available_cluster_types" : [ "analysis", "streaming", "mixed", "custom" ],
"external_datasources" : null,
"resource_requirement" : null,
"valid_roles" : [ "quorumpeer" ],
"visible" : true,
"children_components" : [ "ZooKeeper" ],
"multi_az_support_status" : "multi_az_ha"
}],
"resource_requirement" : [ "privateIp:2" ],
"constraints" : {
  "other" : { },
  "node_constraint" : {
    "other" : { },
    "master" : {
      "other" : { },
      "min_node_num" : 2,
      "max_node_num" : 9,
      "min_core_num" : null,
      "min_mem_size" : null,
      "min_disk_size" : 600,
      "max_node_group_num" : 1,
      "min_data_volume_total_size" : null,
      "disk_type_constraint" : null,
      "min_root_disk_size" : 480
    },
    "core" : {
      "other" : { },
      "min_node_num" : 3,
      "max_node_num" : 500,
      "min_core_num" : null,
      "min_mem_size" : null,
      "min_disk_size" : 600,
      "max_node_group_num" : null,
      "min_data_volume_total_size" : null,
      "disk_type_constraint" : null,
      "min_root_disk_size" : 480
    },
    "task" : null,
    "core_separate" : null,
    "core_combine" : null,
    "task_separate" : null,
    "task_combine" : null,
    "node_group_task" : {
      "other" : { },
      "min_node_num" : 2,
      "max_node_num" : 10000,
      "min_core_num" : null,
      "min_mem_size" : null,
      "min_disk_size" : 600,
      "max_node_group_num" : 9,
      "min_data_volume_total_size" : null,
      "disk_type_constraint" : null,
      "min_root_disk_size" : 480
    }
  },
  "safe_mode_kerberos_exclude_components" : [ "Presto", "Pulsar" ]
},
"flavors" : {
  "master" : [ "ac7.4xlarge.4", "ac7.8xlarge.4", "ac7.16xlarge.4", "ac7.32xlarge.4", "ac7.8xlarge.2",
"ac7.16xlarge.2", "ac7.32xlarge.2", "am7.2xlarge.8", "am7.4xlarge.8", "am7.8xlarge.8", "c3ne.4xlarge.4",
"c3ne.8xlarge.4", "c3ne.15xlarge.4" ],
  "core" : [ "ac7.4xlarge.4", "ac7.8xlarge.4", "ac7.16xlarge.4", "ac7.32xlarge.4", "ac7.8xlarge.2",
"ac7.16xlarge.2", "ac7.32xlarge.2", "am7.2xlarge.8", "am7.4xlarge.8", "am7.8xlarge.8", "c3ne.4xlarge.4",
"c3ne.8xlarge.4", "c3ne.15xlarge.4", "c6.4xlarge.4", "c6.8xlarge.2", "c6.8xlarge.4", "c6.16xlarge.2" ],
  "task" : [ "ac7.4xlarge.4", "ac7.8xlarge.4", "ac7.16xlarge.4", "ac7.32xlarge.4", "ac7.8xlarge.2",
"ac7.16xlarge.2", "ac7.32xlarge.2", "am7.2xlarge.8", "am7.4xlarge.8", "am7.8xlarge.8", "c3ne.4xlarge.4",
"c3ne.8xlarge.4", "c3ne.15xlarge.4" ]
}
```



```
},
"role_deploy_meta" : [ {
  "other" : { },
  "name" : "OMSServer",
  "code_name" : "OMS",
  "component" : "OMSServer",
  "node_preference" : "MASTER",
  "count" : "2",
  "affinity" : null,
  "affinity_target" : null,
  "multi_instance" : null,
  "role_kind" : null,
  "constraints" : null,
  "multi_az_placement" : "AT_LEAST_2",
  "arbitration_deployment" : false,
  "support_elb" : false,
  "multi_affinity_group_enable" : false,
  "local_disks_anti_affinity" : false,
  "multi_instance_name_pattern" : null,
  "private_ip" : null,
  "weight" : null
}, {
  "other" : { },
  "name" : "NameNode",
  "code_name" : "NN",
  "component" : "HDFS",
  "node_preference" : "MASTER",
  "count" : "2",
  "affinity" : null,
  "affinity_target" : null,
  "multi_instance" : null,
  "role_kind" : null,
  "constraints" : null,
  "multi_az_placement" : "AT_LEAST_2",
  "arbitration_deployment" : false,
  "support_elb" : false,
  "multi_affinity_group_enable" : false,
  "local_disks_anti_affinity" : false,
  "multi_instance_name_pattern" : null,
  "private_ip" : null,
  "weight" : null
}, {
  "other" : { },
  "name" : "Zkfc",
  "code_name" : "ZKFC",
  "component" : "HDFS",
  "node_preference" : null,
  "count" : null,
  "affinity" : "JUST_COLOCATE",
  "affinity_target" : "NameNode",
  "multi_instance" : null,
  "role_kind" : null,
  "constraints" : null,
  "multi_az_placement" : "AT_LEAST_2",
  "arbitration_deployment" : false,
  "support_elb" : false,
  "multi_affinity_group_enable" : false,
  "local_disks_anti_affinity" : false,
  "multi_instance_name_pattern" : null,
  "private_ip" : null,
  "weight" : null
}, {
  "other" : { },
  "name" : "HttpFS",
  "code_name" : "HFS",
  "component" : "HDFS",
  "node_preference" : "MASTER",
  "count" : "[0-10]",
  "affinity" : null,
```

```
"affinity_target" : null,
"multi_instance" : null,
"role_kind" : null,
"constraints" : null,
"multi_az_placement" : "AT_LEAST_2",
"arbitration_deployment" : false,
"support_elb" : false,
"multi_affinity_group_enable" : false,
"local_disks_anti_affinity" : false,
"multi_instance_name_pattern" : null,
"private_ip" : null,
"weight" : null
}, {
  "other" : { },
  "name" : "JournalNode",
  "code_name" : "JN",
  "component" : "HDFS",
  "node_preference" : "MASTER",
  "count" : "[3-60],step=2",
  "affinity" : null,
  "affinity_target" : null,
  "multi_instance" : null,
  "role_kind" : null,
  "constraints" : null,
  "multi_az_placement" : "QUORUM_LIKE",
  "arbitration_deployment" : true,
  "support_elb" : false,
  "multi_affinity_group_enable" : false,
  "local_disks_anti_affinity" : false,
  "multi_instance_name_pattern" : null,
  "private_ip" : null,
  "weight" : null
}, {
  "other" : { },
  "name" : "DataNode",
  "code_name" : "DN",
  "component" : "HDFS",
  "node_preference" : "NO_LIMIT",
  "count" : "[3-10000]",
  "affinity" : null,
  "affinity_target" : null,
  "multi_instance" : null,
  "role_kind" : "stateful",
  "constraints" : null,
  "multi_az_placement" : "AT_LEAST_2",
  "arbitration_deployment" : false,
  "support_elb" : false,
  "multi_affinity_group_enable" : false,
  "local_disks_anti_affinity" : false,
  "multi_instance_name_pattern" : null,
  "private_ip" : null,
  "weight" : null
}, {
  "other" : { },
  "name" : "ResourceManager",
  "code_name" : "RM",
  "component" : "Yarn",
  "node_preference" : "MASTER",
  "count" : "2",
  "affinity" : null,
  "affinity_target" : null,
  "multi_instance" : null,
  "role_kind" : null,
  "constraints" : null,
  "multi_az_placement" : "AT_LEAST_2",
  "arbitration_deployment" : false,
  "support_elb" : false,
  "multi_affinity_group_enable" : false,
  "local_disks_anti_affinity" : false,
```

```
"multi_instance_name_pattern" : null,  
"private_ip" : null,  
"weight" : null  
}, {  
  "other" : { },  
  "name" : "NodeManager",  
  "code_name" : "NM",  
  "component" : "Yarn",  
  "node_preference" : "NO_LIMIT",  
  "count" : "[3-10000]",  
  "affinity" : null,  
  "affinity_target" : null,  
  "multi_instance" : null,  
  "role_kind" : "stateless",  
  "constraints" : null,  
  "multi_az_placement" : "AT_LEAST_2",  
  "arbitration_deployment" : false,  
  "support_elb" : false,  
  "multi_affinity_group_enable" : false,  
  "local_disks_anti_affinity" : false,  
  "multi_instance_name_pattern" : null,  
  "private_ip" : null,  
  "weight" : null  
} ]  
}
```

## Status Codes

For details, see [Status Codes](#).

## Error Codes

For details, see [Error Codes](#).

# 8 Out-of-Date APIs

---

## 8.1 Job API Management (Deprecated)

### 8.1.1 Adding and Executing a Job (Deprecated)

#### Function

This API is used to add a job to an MRS cluster and execute the job. This API is incompatible with Sahara.

#### URI

- Format  
POST /v1.1/{project\_id}/jobs/submit-job
- Parameter description

**Table 8-1** URI parameter

| Parameter  | Mandatory | Description                                                                                                  |
|------------|-----------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

## Request Parameters

**Table 8-2** Request parameters

| Parameter  | Mandatory | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_type   | Yes       | Integer | <p>Job type code</p> <ul style="list-style-type: none"> <li>• 1: MapReduce</li> <li>• 2: Spark</li> <li>• 3: Hive Script</li> <li>• 4: HiveQL (not supported currently)</li> <li>• 5: DistCp, importing and exporting data. For details, see <a href="#">Table 8-3</a>.</li> <li>• 6: Spark Script</li> <li>• 7: Spark SQL, submitting Spark SQL statements. For details, see <a href="#">Table 8-4</a>. (Not supported in this API currently.)</li> </ul> <p><b>NOTE</b><br/>Spark and Hive jobs can be added to only clusters that include Spark and Hive components.</p> |
| job_name   | Yes       | String  | <p>Job name</p> <p>Contains only 1 to 64 letters, digits, hyphens (-), and underscores (_).</p> <p><b>NOTE</b><br/>Identical job names are allowed but not recommended.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| cluster_id | Yes       | String  | Cluster ID                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| jar_path  | Yes       | String | <p>Path of the JAR or SQL file for program execution</p> <p>The parameter must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• Contains a maximum of 1,023 characters, excluding special characters such as ; &amp;&gt;&lt;'\$. The address cannot be empty or full of spaces.</li> <li>• Starts with / or s3a://. The OBS path does not support files or programs encrypted by KMS.</li> <li>• Spark Script must end with .sql while MapReduce and Spark Jar must end with .jar.sql and jar are case-insensitive.</li> </ul>                                                                                                                                                                                                     |
| arguments | No        | String | <p>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.</p> <p>The parameter contains a maximum of 2,047 characters, excluding special characters such as ; &amp;&gt;&lt;'\$, and can be left blank.</p> <p><b>NOTE</b></p> <p>When entering a parameter containing sensitive information (for example, login password), you can add an at sign (@) before the parameter name to encrypt the parameter value. This prevents the sensitive information from being persisted in plaintext. Therefore, when you view job information on the MRS, sensitive information will be displayed as asterisks (*).</p> <p>Example: username=xxx<br/>@password=yyy</p> |

| Parameter        | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| input            | No        | String | <p>Path for inputting data, which must start with / or <b>s3a://</b>. Set this parameter to a correct OBS path. The OBS path does not support files or programs encrypted by KMS.</p> <p>The parameter contains a maximum of 1,023 characters, excluding special characters such as ; &amp;&gt;'&lt;\$, and can be left blank.</p>                                                                                                                                                                                                                |
| output           | No        | String | <p>Path for outputting data, which must start with / or <b>s3a://</b>. A correct OBS path is required. If the path does not exist, the system automatically creates it.</p> <p>The parameter contains a maximum of 1,023 characters, excluding special characters such as ; &amp;&gt;'&lt;\$, and can be left blank.</p>                                                                                                                                                                                                                          |
| job_log          | No        | String | <p>Path for storing job logs that record job running status. The path must start with / or <b>s3a://</b>. A correct OBS path is required.</p> <p>The parameter contains a maximum of 1,023 characters, excluding special characters such as ; &amp;&gt;'&lt;\$, and can be left blank.</p>                                                                                                                                                                                                                                                        |
| hive_script_path | Yes       | String | <p>SQL program path</p> <p>This parameter is needed by Spark Script and Hive Script jobs only, and must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• Contains a maximum of 1,023 characters, excluding special characters such as ; &amp;&gt;'&lt;\$. The address cannot be empty or full of spaces.</li> <li>• The path must start with / or <b>s3a://</b>. The OBS path does not support files or programs encrypted by KMS.</li> <li>• The path must end with <b>.sql.sql</b> is case-insensitive.</li> </ul> |

**Table 8-3** DistCp parameters

| Parameter   | Mandatory | Type   | Description                                                                                                                                                                                                                                                     |
|-------------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_name    | Yes       | String | <p>Job name</p> <p>Contains only 1 to 64 letters, digits, hyphens (-), and underscores (_).</p> <p><b>NOTE</b><br/>Identical job names are allowed but not recommended.</p>                                                                                     |
| input       | No        | String | <p>Data source path</p> <ul style="list-style-type: none"> <li>When you import data, the parameter is set to an OBS path. Files or programs encrypted by KMS are not supported.</li> <li>When you export data, the parameter is set to an HDFS path.</li> </ul> |
| output      | No        | String | <p>Data receiving path</p> <ul style="list-style-type: none"> <li>When you import data, the parameter is set to an HDFS path.</li> <li>When you export data, the parameter is set to an OBS path.</li> </ul>                                                    |
| file_action | Yes       | String | <p>Types of file operations, including:</p> <ul style="list-style-type: none"> <li>export: Export data from HDFS to OBS.</li> <li>import: Import data from OBS to HDFS.</li> </ul>                                                                              |



**Table 8-4** Spark SQL parameters

| Parameter | Mandatory | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| hql       | Yes       | String | <p>Spark SQL statement, which needs Base64 encoding and decoding. <b>ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789+/-</b> is a standard encoding table. MRS uses <b>ABCDEFGHIJKLMN OPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789+/-</b> for Base64 encoding. The value of the <b>hql</b> parameter is generated by adding any letter to the beginning of the encoded character string. The Spark SQL statement is generated by decoding the value in the background.</p> <p>Example:</p> <ol style="list-style-type: none"> <li>Obtain the Base64 encoding tool.</li> <li>Enter the <b>show tables;</b> Spark SQL statement in the encoding tool to perform Base64 encoding.</li> <li>Obtain the encoded character string <b>c2hvdYB0YWwLsZXM7</b>.</li> <li>At the beginning of <b>c2hvdYB0YWwLsZXM7</b>, add any letter, for example, <b>g</b>. Then, the character string becomes <b>gc2hvdYB0YWwLsZXM7</b>, that is, the value of the <b>hql</b> parameter.</li> </ol> |
| job_name  | Yes       | String | <p>Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> <p><b>NOTE</b><br/>Identical job names are allowed but not recommended.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## Response Parameters

**Table 8-5** Response parameter

| Parameter     | Type   | Description                                               |
|---------------|--------|-----------------------------------------------------------|
| job_execution | Object | Job details. For details, see <a href="#">Table 8-6</a> . |

**Table 8-6** job\_execution parameters

| Parameter     | Type    | Description                                                                                                                                       |
|---------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| templated     | Bool    | Whether job execution objects are generated by job templates.                                                                                     |
| created_at    | Integer | Creation time, which is a 10-bit timestamp.                                                                                                       |
| updated_at    | Integer | Update time, which is a 10-bit timestamp.                                                                                                         |
| id            | String  | Job ID                                                                                                                                            |
| tenant_id     | String  | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .                                          |
| job_id        | String  | Job application ID                                                                                                                                |
| job_name      | String  | Job name                                                                                                                                          |
| input_id      | String  | Data input ID                                                                                                                                     |
| output_id     | String  | Data output ID                                                                                                                                    |
| start_time    | Integer | Start time of job execution, which is a 10-bit timestamp.                                                                                         |
| end_time      | Integer | End time of job execution, which is a 10-bit timestamp.                                                                                           |
| cluster_id    | String  | Cluster ID                                                                                                                                        |
| engine_job_id | String  | Workflow ID of Oozie                                                                                                                              |
| return_code   | Integer | Returned code for an execution result                                                                                                             |
| is_public     | Bool    | Whether a job is public <ul style="list-style-type: none"> <li>true</li> <li>false</li> </ul> The current version does not support this function. |

| Parameter    | Type    | Description                                                                                                                                                                                                                                                                               |
|--------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| is_protected | Bool    | Whether a job is protected <ul style="list-style-type: none"> <li>• true</li> <li>• false</li> </ul> <p>The current version does not support this function.</p>                                                                                                                           |
| group_id     | String  | Group ID of a job                                                                                                                                                                                                                                                                         |
| jar_path     | String  | Path of the <b>.jar</b> file for program execution                                                                                                                                                                                                                                        |
| input        | String  | Address for inputting data                                                                                                                                                                                                                                                                |
| output       | String  | Address for outputting data                                                                                                                                                                                                                                                               |
| job_log      | String  | Address for storing job logs                                                                                                                                                                                                                                                              |
| job_type     | Integer | Job type code <ul style="list-style-type: none"> <li>• 1: MapReduce</li> <li>• 2: Spark</li> <li>• 3: Hive Script</li> <li>• 4: HiveQL (not supported currently)</li> <li>• 5: DistCp</li> <li>• 6: Spark Script</li> <li>• 7: Spark SQL (not supported in this API currently)</li> </ul> |
| file_action  | String  | Data import and export                                                                                                                                                                                                                                                                    |
| arguments    | String  | Key parameter for program execution. The parameter is specified by the function of the user's internal program. MRS is only responsible for loading the parameter. This parameter can be empty.                                                                                           |
| job_state    | Integer | Job status code <ul style="list-style-type: none"> <li>• -1: Terminated</li> <li>• 1: Starting</li> <li>• 2: Running</li> <li>• 3: Completed</li> <li>• 4: Abnormal</li> <li>• 5: Error</li> </ul>                                                                                        |

| Parameter        | Type    | Description                                                                                                                                                                                      |
|------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_final_status | Integer | Final job status <ul style="list-style-type: none"> <li>• 0: unfinished</li> <li>• 1: terminated due to an execution error</li> <li>• 2: executed successfully</li> <li>• 3: canceled</li> </ul> |
| hive_script_path | String  | Address of the Hive script                                                                                                                                                                       |
| create_by        | String  | User ID for creating jobs<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                                         |
| finished_step    | Integer | Number of completed steps<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                                         |
| job_main_id      | String  | Main ID of a job<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                                                  |
| job_step_id      | String  | Step ID of a job<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                                                  |
| postpone_at      | Integer | Delay time, which is a 10-bit timestamp.<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                          |
| step_name        | String  | Step name of a job<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                                                |
| step_num         | Integer | Number of steps<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                                                   |
| task_num         | Integer | Number of tasks<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions.                                                                   |

| Parameter        | Type   | Description                                                                                                                              |
|------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------|
| update_by        | String | User ID for updating jobs                                                                                                                |
| credentials      | String | Token<br>The current version does not support this function.                                                                             |
| user_id          | String | User ID for creating jobs<br>This parameter is not used in the current version, but is retained for compatibility with earlier versions. |
| job_configs      | String | Key-value pair set for saving job running configurations                                                                                 |
| extra            | String | Authentication information<br>The current version does not support this function.                                                        |
| data_source_urls | String | Data source URL                                                                                                                          |
| info             | String | Key-value pair set, containing job running information returned by Oozie                                                                 |

## Example

- Example request

The following is an example of a MapReduce job request:

```
{
  "job_type": 1,
  "job_name": "mrs_test_jobone_20170602_141106",
  "cluster_id": "e955a7a3-d334-4943-a39a-994976900d56",
  "jar_path": "s3a://mrs-opsadm/jarpath/hadoop-mapreduce-examples-2.7.2.jar",
  "arguments": "wordcount",
  "input": "s3a://mrs-opsadm/input/",
  "output": "s3a://mrs-opsadm/output/",
  "job_log": "s3a://mrs-opsadm/log/",
  "file_action": "",
  "hql": "",
  "hive_script_path": ""
}
```

Example request of a Spark job

```
{
  "job_type": 2,
  "job_name": "mrs_test_sparkjob_20170602_141106",
  "cluster_id": "e955a7a3-d334-4943-a39a-994976900d56",
  "jar_path": "s3a://mrs-opsadm/jarpath/spark-test.jar",
  "arguments": "org.apache.spark.examples.SparkPi 10",
  "input": "",
  "output": "s3a://mrs-opsadm/output/",
  "job_log": "s3a://mrs-opsadm/log/",
  "file_action": "",
  "hql": "",
  "hive_script_path": ""
}
```

Example request of a Hive Script job

```
{
  "job_type": 3,
  "job_name": "mrs_test_SparkScriptJob_20170602_141106",
  "cluster_id": "e955a7a3-d334-4943-a39a-994976900d56",
  "jar_path": "s3a://mrs-opsadm/jarpath/Hivescript.sql",
  "arguments": "",
  "input": "s3a://mrs-opsadm/input/",
  "output": "s3a://mrs-opsadm/output/",
  "job_log": "s3a://mrs-opsadm/log/",
  "file_action": "",
  "hql": "",
  "hive_script_path": "s3a://mrs-opsadm/jarpath/Hivescript.sql"
}
```

#### Example request of a DistCp job for import

```
{
  "job_type": 5,
  "job_name": "mrs_test_importjob_20170602_141106",
  "cluster_id": "e955a7a3-d334-4943-a39a-994976900d56",
  "input": "s3a://mrs-opsadm/jarpath/hadoop-mapreduce-examples-2.7.2.jar",
  "output": "/user",
  "file_action": "import"
}
```

#### Example request for exporting a DistCp job

```
{
  "job_type": 5,
  "job_name": "mrs_test_exportjob_20170602_141106",
  "cluster_id": "e955a7a3-d334-4943-a39a-994976900d56",
  "input": "/user/hadoop-mapreduce-examples-2.7.2.jar",
  "output": "s3a://mrs-opsadm/jarpath/",
  "file_action": "export"
}
```

#### Example request of a Spark Script job

```
{
  "job_type": 6,
  "job_name": "mrs_test_sparkscriptjob_20170602_141106",
  "cluster_id": "e955a7a3-d334-4943-a39a-994976900d56",
  "jar_path": "s3a://mrs-opsadm/jarpath/sparkscript.sql",
  "arguments": "",
  "input": "s3a://mrs-opsadm/input/",
  "output": "s3a://mrs-opsadm/output/",
  "job_log": "s3a://mrs-opsadm/log/",
  "file_action": "",
  "hql": "",
  "hive_script_path": "s3a://mrs-opsadm/jarpath/sparkscript.sql"
}
```

- Example response

```
{
  "job_execution": {
    "templated": false,
    "created_at": 1496387588,
    "updated_at": 1496387588,
    "id": "12ee9ae4-6ee1-48c6-bb84-fb0b4f76cf03",
    "tenant_id": "c71ad83a66c5470496c2ed6e982621cc",
    "job_id": "",
    "job_name": "mrs_test_jobone_20170602_141106",
    "input_id": null,
    "output_id": null,
    "start_time": 1496387588,
    "end_time": null,
    "cluster_id": "e955a7a3-d334-4943-a39a-994976900d56",
    "engine_job_id": null,
    "return_code": null,
    "is_public": null,
    "is_protected": false,
    "group_id": "12ee9ae4-6ee1-48c6-bb84-fb0b4f76cf03",
    "jar_path": "s3a://mrs-opsadm/jarpath/hadoop-mapreduce-examples-2.7.2.jar",
  }
}
```

```

"input": "s3a://mrs-opsadm/input/",
"output": "s3a://mrs-opsadm/output/",
"job_log": "s3a://mrs-opsadm/log/",
"job_type": 1,
"file_action": "",
"arguments": "wordcount",
"hql": "",
"job_state": 2,
"job_final_status": 0,
"hive_script_path": "",
"create_by": "b67132be2f054a45b247365647e05af0",
"finished_step": 0,
"job_main_id": "",
"job_step_id": "",
"postpone_at": 1496387588,
"step_name": "",
"step_num": 0,
"task_num": 0,
"update_by": "b67132be2f054a45b247365647e05af0",
"credentials": "",
"user_id": "b67132be2f054a45b247365647e05af0",
"job_configs": null,
"extra": null,
"data_source_urls": null,
"info": null
}
}

```

## Status Codes

[Table 8-7](#) describes the status code.

**Table 8-7** Status code

| Status Code | Description             |
|-------------|-------------------------|
| 200         | The job has been added. |

See [Status Codes](#).

## 8.1.2 Querying the exe Object List of Jobs (Deprecated)

### Function

This API is used to query the exe object list of all jobs. This API is incompatible with Sahara.

### URI

- Format  
GET /v1.1/{project\_id}/job-exes
- Parameter description

**Table 8-8** URI parameter

| Parameter  | Mandatory | Description                                                                                                  |
|------------|-----------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

## Request Parameters

**Table 8-9** Request parameters

| Parameter    | Mandatory | Type    | Description                                                                                                                                             |
|--------------|-----------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| cluster_id   | Yes       | String  | Cluster ID                                                                                                                                              |
| id           | No        | String  | Job execution object ID                                                                                                                                 |
| page_size    | No        | Integer | Maximum number of jobs displayed on a page<br>Value range: 1 to 100                                                                                     |
| current_page | No        | Integer | Current page number                                                                                                                                     |
| job_name     | No        | String  | Job name                                                                                                                                                |
| state        | No        | Integer | Job status code <ul style="list-style-type: none"> <li>• -1: Terminated</li> <li>• 2: Running</li> <li>• 3: Completed</li> <li>• 4: Abnormal</li> </ul> |

## Response Parameters

**Table 8-10** Response parameters

| Parameter      | Type    | Description                                                         |
|----------------|---------|---------------------------------------------------------------------|
| totalRecord    | Integer | The total number of jobs in the job list.                           |
| job_executions | Array   | Job list parameter<br>For details, see <a href="#">Table 8-11</a> . |



**Table 8-11** job\_executions parameters

| Parameter   | Type    | Description                                                                                                                                                                                                                                                                       |
|-------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id          | String  | Job ID                                                                                                                                                                                                                                                                            |
| create_at   | Integer | Creation time, which is a 13-bit timestamp.                                                                                                                                                                                                                                       |
| update_at   | Integer | Update time, which is a 13-bit timestamp.                                                                                                                                                                                                                                         |
| tenant_id   | String  | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> .                                                                                                                                                                          |
| job_id      | String  | Job ID of the YARN                                                                                                                                                                                                                                                                |
| job_name    | String  | Job name                                                                                                                                                                                                                                                                          |
| start_time  | Integer | Start time of job execution, which is a 13-bit timestamp.                                                                                                                                                                                                                         |
| end_time    | Integer | End time of job execution, which is a 13-bit timestamp.                                                                                                                                                                                                                           |
| cluster_id  | String  | Cluster ID of a job                                                                                                                                                                                                                                                               |
| group_id    | String  | Group ID of a job                                                                                                                                                                                                                                                                 |
| jar_path    | String  | Path of the <b>.jar</b> file or <b>.sql</b> file for program execution                                                                                                                                                                                                            |
| input       | String  | Address for inputting data                                                                                                                                                                                                                                                        |
| output      | String  | Address for outputting data                                                                                                                                                                                                                                                       |
| job_log     | String  | Address for storing job logs                                                                                                                                                                                                                                                      |
| job_type    | Integer | Job type code <ul style="list-style-type: none"><li>• 1: MapReduce</li><li>• 2: Spark</li><li>• 3: Hive Script</li><li>• 4: HiveQL (not supported currently)</li><li>• 5: DistCp</li><li>• 6: Spark Script</li><li>• 7: Spark SQL (not supported in this API currently)</li></ul> |
| file_action | String  | Data import and export                                                                                                                                                                                                                                                            |

| Parameter        | Type    | Description                                                                                                                                                                                       |
|------------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| arguments        | String  | Key parameter for program execution. The parameter is specified by the function of the user's internal program. MRS is only responsible for loading the parameter. This parameter can be empty.   |
| hql              | String  | HiveQL statement                                                                                                                                                                                  |
| job_state        | Integer | Job status code <ul style="list-style-type: none"> <li>• -1: Terminated</li> <li>• 2: Running</li> <li>• 3: Completed</li> <li>• 4: Abnormal</li> </ul>                                           |
| job_final_status | Integer | Final job status. <ul style="list-style-type: none"> <li>• 0: unfinished</li> <li>• 1: terminated due to an execution error</li> <li>• 2: executed successfully</li> <li>• 3: canceled</li> </ul> |
| hive_script_path | String  | Address of the Hive script                                                                                                                                                                        |
| create_by        | String  | User ID for creating jobs                                                                                                                                                                         |
| finished_step    | Integer | Number of completed steps                                                                                                                                                                         |
| job_main_id      | String  | Main ID of a job                                                                                                                                                                                  |
| job_step_id      | String  | Step ID of a job                                                                                                                                                                                  |
| postpone_at      | Integer | Delay time, which is a 13-bit timestamp.                                                                                                                                                          |
| step_name        | String  | Step name of a job                                                                                                                                                                                |
| step_num         | Integer | Number of steps                                                                                                                                                                                   |
| task_num         | Integer | Number of tasks                                                                                                                                                                                   |
| update_by        | String  | User ID for updating jobs                                                                                                                                                                         |
| spend_time       | Integer | Duration of job execution (unit: s)                                                                                                                                                               |
| step_seq         | Integer | Step sequence of a job                                                                                                                                                                            |
| progress         | String  | Job execution progress                                                                                                                                                                            |

## Example

- **Example request**  
GET/v1.1/{project\_id}/job-exes?  
page\_size=10&current\_page=1&state=3&job\_name=myfirstjob&clusterId=20ca8601-72a2-4570-  
b788-2a20fec81a95

- **Example response**

```
{
  "totalRecord": 14,
  "job_executions": [
    {
      "id": "669476bd-89d2-45aa-8f1a-872d16de377e",
      "create_at": 1484641003707,
      "update_at": 1484641003707,
      "tenant_id": "3f99e3319a8943ceb15c584f3325d064",
      "job_id": "",
      "job_name": "myfirstjob",
      "start_time": 1484641003707,
      "end_time": null,
      "cluster_id": "2b460e01-3351-4170-b0a7-57b9dd5ffef3",
      "group_id": "669476bd-89d2-45aa-8f1a-872d16de377e",
      "jar_path": "s3a://jp-test1/program/hadoop-mapreduce-examples-2.4.1.jar",
      "input": "s3a://jp-test1/input/",
      "output": "s3a://jp-test1/output/",
      "job_log": "s3a://jp-test1/joblogs/",
      "job_type": 1,
      "file_action": "",
      "arguments": "wordcount",
      "hql": "",
      "job_state": 2,
      "job_final_status": 1,
      "hive_script_path": null,
      "create_by": "3f99e3319a8943ceb15c584f3325d064",
      "finished_step": 0,
      "job_main_id": "",
      "job_step_id": "",
      "postpone_at": 1484641003174,
      "step_name": "",
      "step_num": 0,
      "task_num": 0,
      "update_by": "3f99e3319a8943ceb15c584f3325d064",
      "spend_time": null,
      "step_seq": 222,
      "progress": "first progress"
    }
  ]
}
```

## Status Codes

[Table 8-12](#) describes the status code.

**Table 8-12** Status code

| Status Code | Description                             |
|-------------|-----------------------------------------|
| 200         | The exe object list of jobs is queried. |

See [Status Codes](#).

## 8.1.3 Querying exe Object Details (Deprecated)

### Function

This API is used to query detailed information about the exe object of a job. This API is incompatible with Sahara.

### URI

- Format  
GET /v1.1/{project\_id}/job-exes/{job\_exe\_id}
- Parameter description

**Table 8-13** URI parameters

| Parameter  | Mandatory | Description                                                                                                  |
|------------|-----------|--------------------------------------------------------------------------------------------------------------|
| project_id | Yes       | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |
| job_exe_id | Yes       | The job ID.                                                                                                  |

### Request Parameters

None

### Response Parameters

**Table 8-14** Response parameter

| Parameter     | Type   | Description                                                |
|---------------|--------|------------------------------------------------------------|
| job_execution | Object | Job details. For details, see <a href="#">Table 8-15</a> . |

**Table 8-15** job\_execution parameters

| Parameter | Type    | Description                                                                                              |
|-----------|---------|----------------------------------------------------------------------------------------------------------|
| id        | String  | Job ID                                                                                                   |
| create_at | Integer | Creation time, which is a 13-bit timestamp.                                                              |
| update_at | Integer | Update time, which is a 13-bit timestamp.                                                                |
| tenant_id | String  | Project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |

| Parameter   | Type    | Description                                                                                                                                                                                                                                                                               |
|-------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_id      | String  | Job ID                                                                                                                                                                                                                                                                                    |
| job_name    | String  | Job name                                                                                                                                                                                                                                                                                  |
| start_time  | Integer | Start time of job execution, which is a 13-bit timestamp.                                                                                                                                                                                                                                 |
| end_time    | Integer | End time of job execution, which is a 13-bit timestamp.                                                                                                                                                                                                                                   |
| cluster_id  | String  | Cluster ID of a job                                                                                                                                                                                                                                                                       |
| group_id    | String  | Group ID of a job                                                                                                                                                                                                                                                                         |
| jar_path    | String  | Path of the <b>.jar</b> file or <b>.sql</b> file for program execution                                                                                                                                                                                                                    |
| input       | String  | Address for inputting data                                                                                                                                                                                                                                                                |
| output      | String  | Address for outputting data                                                                                                                                                                                                                                                               |
| job_log     | String  | Address for storing job logs                                                                                                                                                                                                                                                              |
| job_type    | Integer | Job type code <ul style="list-style-type: none"> <li>• 1: MapReduce</li> <li>• 2: Spark</li> <li>• 3: Hive Script</li> <li>• 4: HiveQL (not supported currently)</li> <li>• 5: DistCp</li> <li>• 6: Spark Script</li> <li>• 7: Spark SQL (not supported in this API currently)</li> </ul> |
| file_action | String  | Data import and export                                                                                                                                                                                                                                                                    |
| arguments   | String  | 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. This parameter can be empty.                                                                                                    |
| hql         | String  | HiveQL statement                                                                                                                                                                                                                                                                          |
| job_state   | Integer | Job status code <ul style="list-style-type: none"> <li>• -1: Terminated</li> <li>• 1: Starting</li> <li>• 2: Running</li> <li>• 3: Completed</li> <li>• 4: Abnormal</li> <li>• 5: Error</li> </ul>                                                                                        |

| Parameter        | Type    | Description                                                                                                                                                                              |
|------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_final_status | Integer | Final job status <ul style="list-style-type: none"> <li>0: unfinished</li> <li>1: terminated due to an execution error</li> <li>2: executed successfully</li> <li>3: canceled</li> </ul> |
| hive_script_path | String  | Address of the Hive script                                                                                                                                                               |
| create_by        | String  | User ID for creating jobs                                                                                                                                                                |
| finished_step    | Integer | Number of completed steps                                                                                                                                                                |
| job_main_id      | String  | Main ID of a job                                                                                                                                                                         |
| job_step_id      | String  | Step ID of a job                                                                                                                                                                         |
| postpone_at      | Integer | Delay time, which is a 13-bit timestamp.                                                                                                                                                 |
| step_name        | String  | Step name of a job                                                                                                                                                                       |
| step_num         | Integer | Number of steps                                                                                                                                                                          |
| task_num         | Integer | Number of tasks                                                                                                                                                                          |
| update_by        | String  | User ID for updating jobs                                                                                                                                                                |
| spend_time       | Integer | Duration of job execution (unit: s)                                                                                                                                                      |
| step_seq         | Integer | Step sequence of a job                                                                                                                                                                   |
| progress         | String  | Job execution progress                                                                                                                                                                   |

## Example

- Example request  
None.
- Example response

```
{
  "job_execution": {
    "id": "632863d5-15d4-4691-9dc1-1a72340cb097",
    "create_at": 1484240559176,
    "update_at": 1484240559176,
    "tenant_id": "3f99e3319a8943ceb15c584f3325d064",
    "job_id": "632863d5-15d4-4691-9dc1-1a72340cb097",
    "job_name": "hive_script",
    "start_time": 1484240559176,
    "end_time": null,
    "cluster_id": "8b1d55f6-150e-45e2-8347-b2ca608d366b",
    "group_id": "632863d5-15d4-4691-9dc1-1a72340cb097",
    "jar_path": "s3a://jp-test1/program/Hivescript.sql",
    "input": "s3a://jp-test1/input/",
    "output": "s3a://jp-test1/output/",
    "job_log": "s3a://jp-test1/joblogs/",
    "job_type": 3,
  }
}
```

```

"file_action": "",
"arguments": "wordcount",
"hql": null,
"job_state": 3,
"job_final_status": 1,
"hive_script_path": "s3a://jp-test1/program/Hivescript.sql",
"create_by": "3f99e3319a8943ceb15c584f3325d064",
"finished_step": 0,
"job_main_id": "",
"job_step_id": "",
"postpone_at": 1484240558705,
"step_name": "",
"step_num": 0,
"task_num": 0,
"update_by": "3f99e3319a8943ceb15c584f3325d064",
"spend_time": null,
"step_seq": 222,
"progress": "first progress"
}
}

```

## Status Code

[Table 8-16](#) describes the status code of this API.

**Table 8-16** Status code

| Status code | Description                                      |
|-------------|--------------------------------------------------|
| 200         | The exe object details are queried successfully. |

For the description about error status codes, see [Status Codes](#).

## 8.1.4 Deleting a Job Execution Object (Deprecated)

### Function

This API is used to delete a job execution object. This API is compatible with Sahara.

### URI

- Format  
DELETE /v1.1/{project\_id}/job-executions/{job\_execution\_id}
- Parameter description

**Table 8-17** URI parameters

| Parameter        | Mandatory | Description                                                                                                  |
|------------------|-----------|--------------------------------------------------------------------------------------------------------------|
| project_id       | Yes       | The project ID. For details about how to obtain the project ID, see <a href="#">Obtaining a Project ID</a> . |
| job_execution_id | Yes       | The job ID.                                                                                                  |

## Request Parameters

None

## Response Parameters

None

## Example

- Example request  
None
- Example response  
None

## Status Codes

[Table 8-18](#) describes the status code.

**Table 8-18** Status code

| Status Code | Description                          |
|-------------|--------------------------------------|
| 204         | The job execution object is deleted. |

See [Status Codes](#).



# 9 Permissions Policies and Supported Actions

---

## 9.1 Introduction

Create IAM users and assign permissions to the users so you can perform fine-grained permissions management for your MRS resources. If your Huawei Cloud account meets service requirements, you do not need to create IAM users. In this case, skip this section.

By default, new IAM users do not have permissions assigned. You need to add them to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

You can grant users permissions by using [roles](#) and [policies](#). Roles are a type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. Policies define API-based permissions for operations on specific resources under certain conditions, allowing for more fine-grained, secure access control of cloud resources.

### NOTE

Policy-based authorization is useful if you want to allow or deny the access to an API.

Each account has all the permissions required to call all APIs, but IAM users must be assigned the required permissions. The permissions required for calling an API are determined by the actions supported by the API. Only users who have been granted permissions allowing the actions can call the API successfully. For example, if an IAM user queries MRS clusters using an API, the user must have been granted permissions that allow the **mrs:cluster:list** action.

## Supported Actions

MRS provides system-defined policies that can be directly used in IAM. MRS cluster administrators can also create custom policies and use them to supplement system-defined policies, implementing more refined access control. Operations supported by policies are specific to APIs. The following are common concepts related to policies:

- **Permission:** A statement in a policy that allows or denies certain operations.
- **API:** REST APIs that can be called by a user who has been granted specific permissions.
- **Action:** Specific operations that are allowed or denied.
- **Dependencies:** actions on which a specific action depends. When allowing an action for a user, you also need to allow any existing action dependencies for that user.
- **IAM or enterprise projects:** Type of projects for which an action will take effect. Policies that contain actions for both IAM and enterprise projects can be used and take effect for both IAM and Enterprise Management. Policies that only contain actions for IAM projects can be used and only take effect for IAM.

 **NOTE**

The check mark (√) indicates that an action takes effect. The cross mark (x) indicates that an action does not take effect.

**Table 9-1** Actions

| Permissions                                               | API                                                | Action             | IAM Project | Enterprise Project |
|-----------------------------------------------------------|----------------------------------------------------|--------------------|-------------|--------------------|
| Creating a Cluster and Executing a Job (V1)               | POST /v1.1/{project_id}/run-job-flow               | mrs:cluster:create | √           | √                  |
| Creating a Cluster (V2)                                   | POST/v2/{project_id}/clusters                      |                    | √           | √                  |
| Querying a Cluster List (V1)                              | GET /v1.1/{project_id}/cluster_infos               | mrs:cluster:list   | √           | √                  |
| Obtaining a Cluster List (V2) (Obtaining Cluster Details) | GET/v2/{project_id}/clusters                       |                    | √           | √                  |
| Deleting a Cluster                                        | DELETE /v1.1/{project_id}/clusters/{cluster_id}    | mrs:cluster:delete | √           | √                  |
| Querying a Host List (V1)                                 | GET /v1.1/{project_id}/clusters/{cluster_id}/hosts | mrs:host:list      | √           | √                  |
| Querying a File List (V2)                                 | GET/v2/{project_id}/clusters/{cluster_id}/files    | mrs:file:list      | √           | √                  |

| Permissions                               | API                                                                                     | Action         | IAM Project | Enterprise Project |
|-------------------------------------------|-----------------------------------------------------------------------------------------|----------------|-------------|--------------------|
| Adding and Executing a Job (V1)           | POST /v1.1/{project_id}/jobs/submit-job                                                 | mrs:job:submit | √           | √                  |
| Adding and Executing a Job (V2)           | POST /v2/{project_id}/clusters/{cluster_id}/job-executions                              |                | √           | √                  |
| Querying the exe Object List of Jobs (V1) | GET /v1.1/{project_id}/job-exes                                                         | mrs:job:list   | √           | √                  |
| Querying Information About a Job (V2)     | GET /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}            |                | √           | √                  |
| Querying a List of Jobs (V2)              | GET /v2/{project_id}/clusters/{cluster_id}/job-executions                               |                | √           | √                  |
| Obtaining the SQL Result (V2)             | GET /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/sql-result |                | √           | √                  |
| Querying exe Object Details (V1)          | GET /v1.1/{project_id}/job-exes/{job_exe_id}                                            | mrs:job:get    | √           | √                  |
| Querying User Agent Information           | GET/v2/{project_id}/clusters/{cluster_id}/agency-mapping                                |                | √           | √                  |
| Querying Job Log Details                  | GET/v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/log-detail  |                | √           | √                  |
| Querying Tags of a Specified Cluster      | GET /v1.1/{project_id}/clusters/{cluster_id}/tags                                       | mrs:tag:list   | √           | √                  |
| Querying All Tags                         | GET /v1.1/{project_id}/clusters/tags                                                    |                | √           | √                  |
| Creating a Tag for a Cluster              | POST/v1.1/{project_id}/clusters/{cluster_id}/tags                                       | mrs:tag:create | √           | √                  |

| Permissions                                     | API                                                                                | Action                | IAM Project | Enterprise Project |
|-------------------------------------------------|------------------------------------------------------------------------------------|-----------------------|-------------|--------------------|
| Adding or Deleting Cluster Tags in Batches      | POST /v1.1/{project_id}/clusters/{cluster_id}/tags/action                          | mrs:tag:batch Operate | √           | √                  |
| Querying a List of Clusters with Specified Tags | POST /v1.1/{project_id}/clusters/resource_instances/action                         | mrs:tag:listResource  | √           | ×                  |
| Terminating a Job (V2)                          | POST /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/kill | mrs:job:stop          | √           | √                  |
| Deleting Jobs in Batches (V2)                   | POST /v2/{project_id}/clusters/{cluster_id}/job-executions/batch-delete            | mrs:job:batch Delete  | √           | √                  |
| Canceling SQL Execution                         | POST/v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}/cancel           | mrs:sql:cancel        | √           | √                  |
| Submitting a SQL Statement                      | POST/v2/{project_id}/clusters/{cluster_id}/sql-execution                           | mrs:sql:execute       | √           | √                  |
| Obtaining All AS Policies                       | GET/v2/{project_id}/autoscaling-policy/{cluster_id}                                | mrs:cluster:policy    | √           | √                  |
| Configuring an Auto Scaling Rule                | POST /v1.1/{project_id}/autoscaling-policy/{cluster_id}                            |                       | √           | √                  |
| Updating User Agent Information                 | PUT/v2/{project_id}/clusters/{cluster_id}/agency-mapping                           | mrs:cluster:syncUser  | √           | √                  |
| Obtaining Job Execution Results                 | GET/v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}                   | mrs:sql:get           | √           | √                  |

## References

- [Permissions Management](#)
- [Creating a User and Granting Permissions](#)
- [MRS Custom Policies](#)

# 10 Appendix

---

## 10.1 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 10-1** 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 10-2** 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 10-3** 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                 |

## 10.2 BMS Specifications Used by MRS

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

- Disk-intensive (D2)
- I/O-optimized (IO2)

### 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 10-4** Specifications of I/O-optimized (IO2) BMSs

| Flavor/ID           | CPU                                           | Memory (GB)       | Local Disk                              | Extended Configuration |
|---------------------|-----------------------------------------------|-------------------|-----------------------------------------|------------------------|
| physical.io2.xlarge | 2 x 22 Core Intel Xeon Gold 6161 V5 (2.2 GHz) | 384 DDR4 RAM (GB) | 2 x 800 GB SSD RAID 1 + 10 x 800 GB SSD | 2 x 2 x 10GE           |

**Table 10-5** Specifications of disk-intensive (D2) BMSs

| Flavor/ID         | CPU                                           | Memory (GB)       | Local Disk                                          | Extended Configuration |
|-------------------|-----------------------------------------------|-------------------|-----------------------------------------------------|------------------------|
| physical.d2.large | 2 x 12 Core Intel Xeon Gold 5118 V5 (2.3 GHz) | 192 DDR4 RAM (GB) | 2 x 600 GB SAS System Disk RAID 1 + 12 x 10 TB SATA | 2 x 2 x 10GE           |

## 10.3 Status Codes

**Table 10-6** describes status codes.

**Table 10-6** Status codes

| Status Code | Message             | Description                                                                                                                                                                                            |
|-------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 100         | Continue            | The client should continue with its request.<br>This interim response is used to inform the client that the initial part of the request has been received and has not yet been rejected by the server. |
| 101         | Switching Protocols | The protocol should be switched. The protocol can only be switched to a newer protocol.<br>For example, the current HTTPS protocol is switched to a later version.                                     |
| 200         | OK                  | The request has been fulfilled.                                                                                                                                                                        |



| Status Code | Message                       | Description                                                                                                                                                                                                                       |
|-------------|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 201         | Created                       | The request has been fulfilled and a new resource has been created.                                                                                                                                                               |
| 202         | Accepted                      | The request has been accepted, but the processing has not been completed.                                                                                                                                                         |
| 203         | Non-Authoritative Information | The server has successfully processed the request, but is returning information that may be from another source.                                                                                                                  |
| 204         | NoContent                     | The request has been fulfilled, but the HTTPS response does not contain a response body.<br>The status code is returned in response to an HTTPS OPTIONS request.                                                                  |
| 205         | Reset Content                 | The server has fulfilled the request, but the requester is required to reset the content.                                                                                                                                         |
| 206         | Partial Content               | The server has successfully processed the partial GET request.                                                                                                                                                                    |
| 300         | Multiple Choices              | There are multiple options for the location of the requested resource. The response contains a list of resource characteristics and addresses from which a user terminal (such as a browser) can choose the most appropriate one. |
| 301         | Moved Permanently             | The requested resource has been assigned a new permanent URI, and the new URI is contained in the response.                                                                                                                       |
| 302         | Found                         | The requested resource resides temporarily under a different URI.                                                                                                                                                                 |
| 303         | See Other                     | The response to the request can be found under a different URI.<br>It can be retrieved by using a GET or POST method.                                                                                                             |
| 304         | Not Modified                  | The requested resource has not been modified. When the server returns this status code, it does not return any resources.                                                                                                         |
| 305         | Use Proxy                     | The requested resource is available only through a proxy.                                                                                                                                                                         |

| Status Code | Message                       | Description                                                                                                                                                    |
|-------------|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 306         | Unused                        | The HTTPS status code is no longer used.                                                                                                                       |
| 400         | BadRequest                    | The request is invalid.<br>The client should not repeat the request without modifications.                                                                     |
| 401         | Unauthorized                  | This status code is returned after the client provides the authentication information, indicating that the authentication information is incorrect or invalid. |
| 402         | Payment Required              | This status code is reserved for future use.                                                                                                                   |
| 403         | Forbidden                     | The server understood the request, but is refusing to fulfill it.<br>The client should not repeat the request without modifications.                           |
| 404         | NotFound                      | The requested resource cannot be found.<br>The client should not repeat the request without modifications.                                                     |
| 405         | MethodNotAllowed              | A request method is not supported for the requested resource.<br>The client should not repeat the request without modifications.                               |
| 406         | Not Acceptable                | The server cannot fulfill the request according to the content characteristics of the request.                                                                 |
| 407         | Proxy Authentication Required | This status code is similar to 401, but indicates that the client must first authenticate itself with the proxy.                                               |
| 408         | Request Time-out              | The server has timed out waiting for the request.<br>The client may repeat the request without modifications at a later time.                                  |

| Status Code | Message                         | Description                                                                                                                                                                                                                                                                                                                         |
|-------------|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 409         | Conflict                        | The request could not be processed due to a conflict with the current state of the resource.<br><br>This status code indicates that the resource that the client is attempting to create already exists, or that the request has failed to be processed because of the update of the conflict request.                              |
| 410         | Gone                            | The requested resource has been deleted permanently and is no longer available.                                                                                                                                                                                                                                                     |
| 411         | Length Required                 | The server is refusing to process the request without a defined <b>Content-Length</b> .                                                                                                                                                                                                                                             |
| 412         | Precondition Failed             | The server did not meet one of the preconditions that the requester put on the request.                                                                                                                                                                                                                                             |
| 413         | Request Entity Too Large        | The server is refusing to process a request because the request entity is too large for the server to process. The server may close the connection to prevent the client from continuing the request. If the server is only temporarily unable to process the request, the response will contain a <b>Retry-After</b> header field. |
| 414         | Request-URI Too Large           | The Request-URI is too long for the server to process.                                                                                                                                                                                                                                                                              |
| 415         | Unsupported Media Type          | The server is unable to process the media format in the request.                                                                                                                                                                                                                                                                    |
| 416         | Requested range not satisfiable | The requested range is invalid.                                                                                                                                                                                                                                                                                                     |
| 417         | Expectation Failed              | The server has failed to meet the requirements of the <b>Expect</b> request-header field.                                                                                                                                                                                                                                           |
| 422         | UnprocessableEntity             | The request is well-formed but is unable to be processed due to semantic errors.                                                                                                                                                                                                                                                    |

| Status Code | Message                     | Description                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 429         | TooManyRequests             | The client has sent excessive number of requests to the server within a given time (exceeding the limit on the access frequency of the client), or the server has received an excessive number of requests within a given time (beyond its processing capability). In this case, the client should resend the request after the time specified in the <b>Retry-After</b> header of the response has elapsed. |
| 500         | InternalServerError         | The server is able to receive the request but unable to understand it.                                                                                                                                                                                                                                                                                                                                       |
| 501         | Not Implemented             | The server does not support the function required to fulfill the request.                                                                                                                                                                                                                                                                                                                                    |
| 502         | Bad Gateway                 | The server was acting as a gateway or proxy and received an invalid request from the remote server.                                                                                                                                                                                                                                                                                                          |
| 503         | ServiceUnavailable          | The requested service is invalid. The client should not repeat the request without modifications.                                                                                                                                                                                                                                                                                                            |
| 504         | ServerTimeout               | The request cannot be fulfilled within a given time. This status code is returned to the client only if the <b>Timeout</b> parameter is specified in the request.                                                                                                                                                                                                                                            |
| 505         | HTTPS Version not supported | The server does not support the HTTPS protocol version used in the request.                                                                                                                                                                                                                                                                                                                                  |

## 10.4 Error Codes

If an error code starting with **APIGW** is returned after you call an API, rectify the fault by referring to [Error Codes](#).

| Status Code | Error Code | Error Message                                           | Description                                                       | Measure                                                                                                                                                                           |
|-------------|------------|---------------------------------------------------------|-------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 400         | 0023       | Failed to obtain cluster details.                       | Failed to obtain cluster details.                                 | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. Check whether the EIP exists and use a correct EIP. |
| 400         | 0056       | Jobs cannot be submitted to a cluster in the xxx state. | Failed to submit the job because the cluster is in the xxx state. | Wait until the cluster task is complete and the status changes to <b>Running</b> .                                                                                                |
| 400         | 0057       | Spark jobs cannot be submitted.                         | Spark jobs cannot be submitted.                                   | Check whether the Spark service exists and whether the Spark service is normal.                                                                                                   |
| 400         | 0093       | Metadata of version xxx not found.                      | Failed to find version metadata.<br>Version: xxx.                 | Check whether the cluster version is correct.                                                                                                                                     |
| 400         | 0160       | Failed to kill the job.                                 | Failed to terminate a job.                                        | Call the API again, contact technical support, or check whether the cluster is normal.                                                                                            |

| Status Code | Error Code | Error Message                       | Description                         | Measure                                                                                |
|-------------|------------|-------------------------------------|-------------------------------------|----------------------------------------------------------------------------------------|
| 400         | 0161       | Failed to delete jobs in batches.   | Failed to delete jobs in batches.   | Call the API again, contact technical support, or check whether the cluster is normal. |
| 400         | 0162       | Failed to query the job.            | Failed to query the job.            | Call the API again, contact technical support, or check whether the cluster is normal. |
| 400         | 0165       | Failed to verify the SQL statement. | Failed to verify the SQL statement. | Inspect the SQL statement.                                                             |
| 400         | 0166       | Failed to query a job list.         | Failed to query the jobs.           | Call the API again, contact technical support, or check whether the cluster is normal. |
| 400         | 0167       | The v2 job API cannot be accessed.  | The v2 job APIs cannot be accessed. | Use the v1 APIs.                                                                       |
| 400         | 0168       | Hive jobs cannot be submitted.      | Failed to submit the Hive job.      | Check whether the Hive service exists and whether the Hive service is normal.          |
| 400         | 0169       | Flink jobs cannot be submitted.     | Failed to submit the Flink job.     | Check whether the Flink service exists and whether the Flink service is normal.        |

| Status Code | Error Code | Error Message                                              | Description                                              | Measure                                                                                                                       |
|-------------|------------|------------------------------------------------------------|----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 400         | 0170       | Failed to collect job log directory information.           | Failed to collect job log directory information.         | Call the API again, contact technical support, or check whether the cluster is normal.                                        |
| 400         | 0171       | Failed to collect job log details.                         | Failed to collect job log details.                       | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 0172       | Failed to collect the SQL job result.                      | Failed to obtain the SQL statement execution result.     | Call the API again, contact technical support, or check whether the cluster is normal.                                        |
| 400         | 0173       | Failed to submit the job (cluster ID: xxx; job name: xxx). | Failed to submit the job. Cluster ID: xxx, job name: xxx | Call the API again, contact technical support, or check whether the cluster is normal.                                        |
| 400         | 0174       | Failed to query the job.                                   | Failed to query the job.                                 | Call the API again, contact technical support, or check whether the cluster is normal.                                        |

| Status Code | Error Code | Error Message                                                | Description                                                  | Measure                                                                                |
|-------------|------------|--------------------------------------------------------------|--------------------------------------------------------------|----------------------------------------------------------------------------------------|
| 400         | 0175       | Failed to kill the job.                                      | Failed to terminate a job.                                   | Initiate the call again or contact technical support.                                  |
| 400         | 0176       | The job does not exist.                                      | Job not exist.                                               | Modify the request parameters based on the error message.                              |
| 400         | 0177       | The number of jobs running in each cluster cannot exceed 10. | The number of running jobs in each cluster cannot exceed 10. | Modify the request parameters based on the error message.                              |
| 400         | 0178       | The job ID cannot be left blank.                             | The job ID cannot be left blank.                             | Modify the request parameters based on the error message.                              |
| 400         | 0179       | The job type must be SparkSql or SparkScript.                | The job type must be SparkSql or SparkScript.                | Modify the request parameters based on the error message.                              |
| 400         | 0180       | The job is being submitted.                                  | The job is being submitted.                                  | Perform the operation after the current operation is complete.                         |
| 400         | 0181       | The SQL job result collection is empty.                      | The SQL job result is empty.                                 | Call the API again, contact technical support, or check whether the cluster is normal. |



| Status Code | Error Code | Error Message                                              | Description                                                | Measure                                                                                |
|-------------|------------|------------------------------------------------------------|------------------------------------------------------------|----------------------------------------------------------------------------------------|
| 400         | 0182       | Failed to read the SQL job result.                         | Failed to read the result of the SQL job.                  | Call the API again, contact technical support, or check whether the cluster is normal. |
| 400         | 0183       | The job is running.                                        | The job is running.                                        | Perform the operation after the current operation is complete.                         |
| 400         | 0185       | The log type does not exist.                               | Log type not exist.                                        | Modify the request parameters based on the error message.                              |
| 400         | 0187       | The log aggregation path is empty. Logs cannot be queried. | The log aggregation path is empty. Logs cannot be queried. | Check whether the job log path is correct.                                             |
| 400         | 0188       | The job query result is empty. Logs cannot be obtained.    | The job query result is empty. Logs cannot be obtained.    | Check whether the job log path is correct.                                             |
| 400         | 0189       | Failed to delete the job list.                             | Failed to delete the jobs.                                 | Call the API again, contact technical support, or check whether the cluster is normal. |
| 400         | 0190       | The user who submits the job cannot be empty.              | The user who submits the job cannot be empty.              | Modify the request parameters based on the error message.                              |

| Status Code | Error Code | Error Message                                                                                                                                                     | Description                                                                                                                                                  | Measure                                                   |
|-------------|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| 400         | 0190       | The user who submits the job cannot be empty.                                                                                                                     | The user who submits the job cannot be empty.                                                                                                                | Modify the request parameters based on the error message. |
| 400         | 0191       | Failed to query the user who submits the job on MRS Manager.                                                                                                      | Failed to query the user who submits a job on MRS manager.                                                                                                   | Modify the request parameters based on the error message. |
| 400         | 0192       | The current user does not exist on MRS Manager. Grant the user sufficient permissions on IAM and then perform IAM user synchronization on the Dashboard tab page. | The current user does not exist on MRS Manager. Escalate privileges for the user and then perform IAM user synchronization on the <b>Dashboard</b> tab page. | Rectify the fault based on the error message.             |
| 400         | 0193       | Failed to operate the database job records.                                                                                                                       | Failed to operate the database job records.                                                                                                                  | Initiate the call again or contact technical support.     |
| 400         | 0194       | Failed to start the launcher and submit the job.                                                                                                                  | Failed to start the job and to submit the job.                                                                                                               | Initiate the call again or contact technical support.     |
| 400         | 0199       | Failed to delete the jobs.                                                                                                                                        | Failed to delete the job.                                                                                                                                    | Initiate the call again or contact technical support.     |

| Status Code | Error Code | Error Message                                                                                                                                                                     | Description                                                                                                                                                                         | Measure                                                                  |
|-------------|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| 400         | 0200       | Failed to kill the jobs on Yarn.                                                                                                                                                  | Failed to terminate a job on Yarn.                                                                                                                                                  | Initiate the call again or contact technical support.                    |
| 400         | 0201       | The job does not exist.                                                                                                                                                           | Job not exist.                                                                                                                                                                      | Modify the request parameters based on the error message.                |
| 400         | 0202       | Too many jobs are being submitted, please try again later.                                                                                                                        | Too many jobs are submitted. Please try again later.                                                                                                                                | Initiate the call again or contact technical support.                    |
| 400         | 0211       | The maximum number of bound security groups has been reached. A maximum of four security groups can be bound, excluding the default security group that is automatically created. | The number of bound security groups exceeded the upper limit. A maximum of four security groups can be bound, excluding the default security groups that are automatically created. | Modify the request parameters based on the error message.                |
| 400         | 12000002   | The parameter is invalid.                                                                                                                                                         | Invalid parameter.                                                                                                                                                                  | Check whether the parameters are correct according to the API reference. |
| 400         | 12000003   | The cluster does not exist.                                                                                                                                                       | Cluster not exist.                                                                                                                                                                  | Check whether the cluster exists.                                        |

| Status Code | Error Code | Error Message                                                                                                                                                                                                | Description                                                                                                                                                                                             | Measure                                                                            |
|-------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| 400         | 12000009   | The parameter is invalid.                                                                                                                                                                                    | Invalid parameter.                                                                                                                                                                                      | Check whether the parameters are correct according to the API document.            |
| 400         | 12000012   | Tasks are being executed in the cluster. Subcontract period cluster is not allowed for clusters that are in this state. Scale-out or scale-in is not allowed for clusters that are not in the running state. | Task being executed in the cluster. Switching to the yearly/ monthly cluster is not allowed in this status. Scale-out or scale-in is not allowed for clusters that are not in the <b>Running</b> state. | Wait until the cluster task is complete and the status changes to <b>Running</b> . |
| 400         | 12000013   | Failed to scale in the cluster (ID: xxx). The type and the quantity of nodes to be deleted are xxx and xxx, respectively. The task node does not exist.                                                      | The task node does not exist.                                                                                                                                                                           | Check whether the node exists.                                                     |
| 400         | 12000014   | Failed to scale out the cluster (ID: xxx). The type and the quantity of nodes to be added are xxx and xxx, respectively.                                                                                     | Failed to scale out the cluster. Cluster ID: xxx, node type: xxx, and number of nodes: xxx                                                                                                              | Initiate the call again or contact technical support.                              |

| Status Code | Error Code | Error Message                                                              | Description                                                          | Measure                                                                                                                       |
|-------------|------------|----------------------------------------------------------------------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 400         | 12000018   | Scale-out or scale-in cannot be performed again because it is in progress. | Scale-out or scale-in cannot be performed because it is in progress. | Wait until the cluster task is complete and the status changes to <b>Running</b> .                                            |
| 400         | 12000019   | Failed to obtain hosts of the cluster.                                     | Failed to obtain hosts of the cluster.                               | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 12000020   | Failed to terminate the cluster.                                           | Failed to terminate the cluster.                                     | Initiate the call again or contact technical support.                                                                         |
| 400         | 12000021   | Clusters in the xxx state cannot be terminated.                            | Clusters in the xxx state cannot be terminated.                      | Wait until the cluster task is complete and the status changes to <b>Running</b> .                                            |
| 400         | 120000212  | Failed to obtain the AZs that you have permission to access.               | Failed to obtain the authorized AZ.                                  | Try to call the API again.                                                                                                    |
| 400         | 120000213  | The region [xxx] does not exist.                                           | The specified region [xxx] does not exist.                           | Modify the request parameters based on the error message.                                                                     |

| Status Code | Error Code | Error Message                                                                   | Description                                                                 | Measure                                                                                                                       |
|-------------|------------|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 400         | 120000214  | No permission to access the AZs [xxx]. The following AZs [xxx] can be accessed. | You do not have the permission to access AZ [xxx]. You can access AZ [xxx]. | Modify the request parameters based on the error message.                                                                     |
| 400         | 12000023   | Failed to obtain cluster details.                                               | Failed to obtain cluster details.                                           | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 12000024   | Yearly/ Monthly clusters cannot be terminated.                                  | Yearly/ Monthly clusters cannot be terminated.                              | Do not perform this operation.                                                                                                |
| 400         | 12000027   | Failed to verify the subnet when creating the cluster xxx.                      | Failed to create cluster xxx. The subnet verification failed.               | Check whether the number of available subnets meets the requirements for creating a cluster.                                  |
| 400         | 12000028   | The cluster has a maximum of xxx Core and Task nodes.                           | The total number of cores and task nodes in a cluster cannot exceed xxx.    | Try to reduce the number of nodes required in the request.                                                                    |
| 400         | 12000029   | Failed to obtain the quota.                                                     | Failed to obtain the quota.                                                 | Try to call the API again.                                                                                                    |

| Status Code | Error Code | Error Message                                                             | Description                                                                 | Measure                                                                       |
|-------------|------------|---------------------------------------------------------------------------|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| 400         | 12000030   | The requested number of nodes in the cluster exceeds the available quota. | The total number of nodes to be added to the cluster exceeds the quota.     | Try to reduce the amount of resources requested or increase available quotas. |
| 400         | 12000031   | The requested number of vCPUs in the cluster exceeds the available quota. | The total number of CPU cores to be added to the cluster exceeds the quota. | Try to reduce the amount of resources requested or increase available quotas. |
| 400         | 12000032   | The requested memory of the cluster exceeds the available quota.          | The total memory to be added to the cluster exceeds the quota.              | Try to reduce the amount of resources requested or increase available quotas. |
| 400         | 12000033   | The requested number of disks in the cluster exceeds the available quota. | The total number of disks to be added to the cluster exceeds the quota.     | Try to reduce the amount of resources requested or increase available quotas. |
| 400         | 12000034   | The requested disk capacity of the cluster exceeds the available quota.   | The total disk capacity to be expanded exceeds the quota.                   | Try to reduce the amount of resources requested or increase available quotas. |
| 400         | 12000036   | Failed to obtain product information.                                     | Failed to obtain product information.                                       | Try to call the API again.                                                    |
| 400         | 12000038   | Failed to obtain the security group.                                      | Failed to obtain the security group.                                        | Try to call the API again.                                                    |

| Status Code | Error Code | Error Message                                  | Description                                                     | Measure                                                                                                                       |
|-------------|------------|------------------------------------------------|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 400         | 12000041   | Failed to obtain the cluster list.             | Failed to obtain the cluster list.                              | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 12000042   | Failed to create a cluster.                    | Failed to create the cluster.                                   | Initiate the call again, or check whether the parameters are correct according to the API reference.                          |
| 400         | 12000043   | Duplicate cluster name: xxx.                   | Cluster name xxx already exists.                                | Use another cluster name.                                                                                                     |
| 400         | 12000044   | The minimum memory of a Master node is xxx GB. | The minimum memory of the Master node in the cluster is xxx GB. | Increase the memory size of the Master node.                                                                                  |
| 400         | 12000045   | Insufficient quota of the security group.      | Insufficient security group quota.                              | Increase the security group quota.                                                                                            |
| 400         | 12000046   | Insufficient quota of the security group rule. | Insufficient security group rule quota.                         | Increase the security group rule quota or delete unnecessary security group rules.                                            |
| 400         | 12000048   | Product specification xxx does not exist.      | Product specification xxx does not exist.                       | Change the product specifications or AZ.                                                                                      |



| Status Code | Error Code | Error Message                                             | Description                                                   | Measure                                                                   |
|-------------|------------|-----------------------------------------------------------|---------------------------------------------------------------|---------------------------------------------------------------------------|
| 400         | 12000050   | Incorrect certificate.                                    | Incorrect certificate.                                        | Replace the certificate with a correct one.                               |
| 400         | 12000052   | No access rights.                                         | Unauthorized access.                                          | Check whether the permission meets the requirements.                      |
| 400         | 12000053   | Invalid billing type.                                     | Invalid order type                                            | Check whether the parameters are correct according to the API document.   |
| 400         | 12000054   | The operation is not supported.                           | Operation not supported.                                      | Do not perform this operation.                                            |
| 400         | 12000055   | Failed to open the file.                                  | Failed to open the file.                                      | Initiate the call again or contact technical support.                     |
| 400         | 12000059   | Key pair xxx does not exist.                              | User key pair xxx does not exist.                             | Check whether the key pair exists. If yes, replace it with a correct one. |
| 400         | 12000060   | The number of running jobs per cluster cannot exceed xxx. | The number of running jobs in each cluster cannot exceed xxx. | Modify the request parameters based on the error message.                 |

| Status Code | Error Code | Error Message                                                                             | Description                                                                      | Measure                                                                                                                       |
|-------------|------------|-------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 400         | 12000061   | Failed to submit the job (cluster ID: xxx; job name: xxx; job ID: xxx).                   | Failed to submit the job. Cluster ID: xxx, job name: xxx, and job ID: xxx        | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 12000062   | Jobs cannot be submitted to a cluster in the xxx state.                                   | Failed to submit the job because the cluster is in the xxx state.                | Wait until the cluster task is complete and the status changes to <b>Running</b> .                                            |
| 400         | 12000063   | Spark jobs cannot be submitted.                                                           | Spark jobs cannot be submitted.                                                  | Check whether the Spark service exists and whether the Spark service is normal.                                               |
| 400         | 12000064   | Jobs cannot be submitted or inquired to a security cluster by API.                        | This API is not supported for security clusters for submitting or querying jobs. | Submit the job using a non-security cluster or a client.                                                                      |
| 400         | 12000065   | Scale-out or scale-in is not available for Yearly/ Monthly clusters.                      | Scale-out or scale-in is not available for Yearly/ Monthly clusters.             | Do not perform this operation.                                                                                                |
| 400         | 12000068   | If the job type is Hive or Spark Script, the value of the mains should not be left blank. | When the job type is Hive or Spark Script, <b>mains</b> cannot be empty.         | Modify the request parameters based on the error message.                                                                     |

| Status Code | Error Code | Error Message                                                                          | Description                                                      | Measure                                                                                                                       |
|-------------|------------|----------------------------------------------------------------------------------------|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 400         | 12000069   | If the job type is MapReduce or Spark, the value of the libs should not be left blank. | When the job type is Hive or Spark, <b>libs</b> cannot be empty. | Modify the request parameters based on the error message.                                                                     |
| 400         | 12000070   | An error occurred while accessing Knox.                                                | An exception occurred when accessing KNOX.                       | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 12000071   | The Executor server has an internal error.                                             | Internal Executor error.                                         | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 12000073   | Failed to access the HDFS directory.                                                   | Failed to connect to an HDFS directory.                          | Call the API again, contact technical support, or check whether the HDFS service of the cluster is normal.                    |

| Status Code | Error Code | Error Message                                                                                                                                                                                                                                                                                                                          | Description                                                                                                                                                                                                                                                                           | Measure                                                   |
|-------------|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| 400         | 12000077   | The password of [xxx] must meet the following requirements:<br>1.must be xxx to xxx characters long. 2.at least contain [xxx] types of the following characters: uppercase letters, lowercase letters, digits, and special characters(`~!@#\$% ^&*()-_+=  [{}];',<.>?).<br>3.cannot be the username or the username spelled backwards. | The password xxx must:<br>- Contain xxx to xxx characters.<br>- Contain at least xxx of the following character types: uppercase letters, lowercase letters, digits, and special character `~!@#\$% ^&*()-_+= [{}];',<.>?<br>- Not be the username or the username spelled backwards. | Modify the request parameters based on the error message. |
| 400         | 12000080   | The status of some nodes is not running in the cluster. Try again later.                                                                                                                                                                                                                                                               | Some nodes are not running in the cluster. Try again later.                                                                                                                                                                                                                           | Try to call the API again.                                |
| 400         | 12000081   | The jar_path parameter cannot be left blank if the job type is MapReduce or Spark.                                                                                                                                                                                                                                                     | If the job type is MapReduce or Spark, the <b>jar_path</b> parameter cannot be empty.                                                                                                                                                                                                 | Added the <b>jar_path</b> parameter.                      |

| Status Code | Error Code | Error Message                                                                                                                                                        | Description                                                                                                                                                                | Measure                                                                            |
|-------------|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| 400         | 12000082   | Node groups cannot be deleted in the cluster that is being scaled in or scaled out.                                                                                  | Node groups cannot be deleted from a cluster that is being scaled in or out.                                                                                               | Wait until the cluster task is complete and the status changes to <b>Running</b> . |
| 400         | 12000085   | This interface does not support the cluster of this version. Please use the <code>/v2/{project_id}/clusters/{cluster_id}/job-executions</code> job submit interface. | This API does not support clusters of this version. Use the <code>/v2/{project_id}/clusters/{cluster_id}/job-executions</code> to submit jobs.                             | Use a new API that meets the cluster version requirements.                         |
| 400         | 12000086   | This cluster version not support s3[an]: schema. Please use the obs: schema.                                                                                         | The cluster does not support file paths starting with <code>s3a://</code> or <code>s3n://</code> . Use a file path supported by the cluster, such as <code>obs://</code> . | Use a protocol that matches the cluster version.                                   |
| 400         | 12000087   | Failed to get billing records.                                                                                                                                       | Failed to obtain CDR files.                                                                                                                                                | Try to call the API again.                                                         |
| 400         | 12000090   | Products of the xxx specifications are no longer available in selected AZ.                                                                                           | The product specification xxx has been removed from the selected AZ.                                                                                                       | Change the product specifications or AZ.                                           |
| 400         | 12000092   | Failed to get metadata of version xxx.                                                                                                                               | Failed to obtain version metadata. Version: xxx.                                                                                                                           | Try to call the API again.                                                         |

| Status Code | Error Code | Error Message                                                                 | Description                                                                                | Measure                                                                                               |
|-------------|------------|-------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| 400         | 12000093   | Metadata of version xxx not found.                                            | Failed to find version metadata.<br>Version: xxx.                                          | Check whether the cluster version is correct and use the correct version.                             |
| 400         | 12000094   | The xxx in xxx version does not support the xxx flavor.                       | The xxx node does not support the xxx flavor in the xxx version.                           | Use another flavor.                                                                                   |
| 400         | 12000095   | Patch xxx is unavailable.                                                     | Patch xxx is unavailable.                                                                  | Check whether the patch version is correct.                                                           |
| 400         | 12000099   | Topology template of version xxx not found.                                   | Information about the topology template cannot be found.<br>Version: xxx.                  | Check whether the cluster version is correct.                                                         |
| 400         | 12000100   | Failed to stop the cluster. Only running or abnormal clusters can be stopped. | Failed to stop the cluster. Only clusters in the running or abnormal state can be stopped. | Wait until the cluster task is complete and the status changes to <b>Running</b> or <b>Abnormal</b> . |
| 400         | 12000101   | Failed to start the cluster. Only stopped clusters can be started.            | Failed to start the cluster. Only clusters in the stopped state can be started.            | Wait until the cluster task is complete and the status changes to <b>Stopped</b> .                    |

| Status Code | Error Code | Error Message                                           | Description                                                | Measure                                                                                                                       |
|-------------|------------|---------------------------------------------------------|------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 400         | 12000102   | Failed to stop the cluster.                             | Failed to stop the cluster.                                | Call the API again, contact technical support, or check whether the cluster is normal.                                        |
| 400         | 12000103   | Failed to start the cluster.                            | Failed to start the cluster.                               | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 400         | 12000104   | VPC ECS DSS DCC OpenStack service error.                | VPC ECS DSS DCC OpenStack service error.                   | Try to call the API again.                                                                                                    |
| 400         | 12000105   | VPC xxx does not exist.                                 | VPC xxx not exist.                                         | Check whether the VPC exists.                                                                                                 |
| 400         | 12000106   | Key pair xxx does not exist.                            | User key pair xxx not exist.                               | Check whether the key pair exists and replace it with a correct one.                                                          |
| 400         | 12000107   | Invalid project ID: xxx.                                | Invalid project ID: xxx.                                   | Use a correct project ID.                                                                                                     |
| 400         | 12000108   | Failed to verify the EIP when creating the cluster xxx. | Failed to create cluster xxx. The EIP verification failed. | Check whether the EIP exists and use a correct EIP.                                                                           |

| Status Code | Error Code | Error Message                                                           | Description                                                             | Measure                                                                                                                                                                           |
|-------------|------------|-------------------------------------------------------------------------|-------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 400         | 12000109   | Failed to bind the EIP to cluster xxx.                                  | Failed to bind an EIP to cluster xxx.                                   | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. Check whether the EIP exists and use a correct EIP. |
| 400         | 12000110   | Failed to unbind the EIP from cluster xxx.                              | Failed to unbind the EIP from cluster xxx.                              | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal.                                                     |
| 400         | 12000111   | Failed to bind the EIP. The EIP xxx has been bound to another resource. | Failed to bind the EIP. The EIP xxx has been bound to another resource. | Use another EIP that is not in use.                                                                                                                                               |
| 400         | 12000112   | The EIP xxx does not exist.                                             | EIP xxx not exist.                                                      | Check whether the EIP exists and use a correct EIP.                                                                                                                               |
| 400         | 12000113   | Failed to update the EIP.                                               | Failed to update the EIP.                                               | Check whether the EIP exists and use a correct EIP.                                                                                                                               |



| Status Code | Error Code | Error Message                                                                               | Description                                                                                    | Measure                                                              |
|-------------|------------|---------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| 400         | 12000114   | The hive_script_path parameter cannot be left blank if the job type is Hive or SparkScript. | If the job type is Hive or SparkScript, the <b>hive_script_path</b> parameter cannot be empty. | Add the missing parameters as prompted.                              |
| 400         | 12000115   | The ECS group quota is insufficient.                                                        | Insufficient ECS group quota.                                                                  | Increase the quota.                                                  |
| 400         | 12000116   | The VPC quota is insufficient. Select the existing VPC or increase the quota.               | Insufficient VPC quota. Select an existing VPC or apply for more quotas.                       | Increase the quota.                                                  |
| 400         | 12000117   | The subnet quota is insufficient. Select the existing subnet or increase the quota.         | Insufficient subnet quota. Select an existing subnet or apply for more quotas.                 | Increase the quota.                                                  |
| 400         | 12000118   | Failed to create the security group rule.                                                   | Failed to add a security group rule.                                                           | Check whether the security group exists or try again later.          |
| 400         | 12000119   | The security group rule already exists.                                                     | The security group rule already exists.                                                        | Do not perform this operation or manually modify the security group. |

| Status Code | Error Code | Error Message                                                                                                                                                                               | Description                                                                                                                                                                                               | Measure                                                                            |
|-------------|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| 400         | 12000121   | Failed to submit a request to scale out the Yearly/ Monthly cluster. Scale-out is not allowed because the cluster has an unpaid order. Scale out the cluster again after you pay the order. | Failed to submit the request for scaling out a yearly/ monthly subscribed cluster. Scale-out is not allowed because the cluster has an unpaid order. Scale out the cluster again after you pay the order. | Complete the order before performing the operation.                                |
| 400         | 12000122   | EPS service error.                                                                                                                                                                          | EPS error.                                                                                                                                                                                                | Try to call the API again.                                                         |
| 400         | 12000123   | Failed to update Task node information because the number of Task nodes is not 0.                                                                                                           | Failed to update task node information because the number of task nodes is not 0.                                                                                                                         | Remove all nodes from the node group.                                              |
| 400         | 12000124   | In the cluster xxx, the number of Task nodes can be adjusted only using auto scaling.                                                                                                       | Cluster xxx supports only auto scaling to adjust the number of task nodes.                                                                                                                                | Rectify the fault based on the error message.                                      |
| 400         | 12000125   | Failed to update Task node information, because the cluster state is scaling out or scaling in.                                                                                             | Failed to update task node information because the cluster is in the scale-out or scale-in state.                                                                                                         | Wait until the cluster task is complete and the status changes to <b>Running</b> . |

| Status Code | Error Code | Error Message                                                                                  | Description                                                                                                      | Measure                                                                            |
|-------------|------------|------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| 400         | 12000126   | Failed to obtain authentication information.                                                   | Failed to obtain authentication information.                                                                     | Try to call the API again.                                                         |
| 400         | 12000127   | Failed to lock cluster operation.                                                              | Failed to lock cluster operations.                                                                               | Try to call the API again.                                                         |
| 400         | 12000128   | Failed to unlock cluster operation.                                                            | Failed to unlock cluster operations.                                                                             | Try to call the API again.                                                         |
| 400         | 12000129   | Master node specifications cannot be scaled up for a cluster that is not in the running state. | The specifications of the Master node cannot be upgraded because the cluster is not in the <b>Running</b> state. | Wait until the cluster task is complete and the status changes to <b>Running</b> . |
| 400         | 12000130   | Specifications available for scale-up not found.                                               | Failed to find the node specifications that can be upgraded.                                                     | Use another flavor.                                                                |
| 400         | 12000131   | Master node specifications cannot be scaled up for a non-HA cluster.                           | The Master node in a non-HA cluster cannot be upgraded.                                                          | Rectify the fault based on the error message.                                      |
| 400         | 12000132   | vCPUs and memory cannot be reduced in the specification scale-up.                              | The number of vCPUs or memory cannot be reduced during node specification upgrade.                               | Rectify the fault based on the error message.                                      |
| 400         | 12000133   | Specification scale-up is not available for this type of nodes.                                | This node type is not supported for specification upgrade.                                                       | Rectify the fault based on the error message.                                      |

| Status Code | Error Code | Error Message                                                                                         | Description                                                                                           | Measure                                                        |
|-------------|------------|-------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|----------------------------------------------------------------|
| 400         | 12000134   | Failed to scale up the Master node specifications.                                                    | Failed to upgrade the specifications of the Master node.                                              | Try to call the API again.                                     |
| 400         | 12000135   | Master nodes available for specification scale-up not found                                           | Failed to find the Master node that can be upgraded.                                                  | Check whether the node exists.                                 |
| 400         | 12000138   | Failed to get max server Group members.                                                               | Failed to query the maximum number of VMs in an ECS anti-affinity group.                              | Try to call the API again.                                     |
| 400         | 12000139   | All evs volume type is sellout, please try again later.                                               | All EVS disks have been sold out. Please try again later.                                             | Try to call the API again.                                     |
| 400         | 12000140   | Evs volume type:xxx is sellout, please try again later.                                               | EVS disk xxx has been sold out. Please try again later.                                               | Try to call the API again.                                     |
| 400         | 12000141   | The disk size of a node cannot be less than {value} GB.                                               | The node disk cannot be less than {value} GB.                                                         | Rectify the fault based on the error message.                  |
| 400         | 12000142   | The disk size cannot exceed 32,000 GB.                                                                | The disk cannot be greater than 32,000 GB.                                                            | Rectify the fault based on the error message.                  |
| 400         | 12000154   | IAM synchronization is in progress and cannot be triggered again in the same cluster. Cluster ID: xxx | IAM synchronization is in progress and cannot be triggered again in the same cluster. Cluster ID: xxx | Perform the operation after the current operation is complete. |

| Status Code | Error Code | Error Message                           | Description                                   | Measure                                                                                |
|-------------|------------|-----------------------------------------|-----------------------------------------------|----------------------------------------------------------------------------------------|
| 400         | 12000156   | Failed to query iam group               | Failed to query the IAM user group.           | Initiate the call again or contact technical support.                                  |
| 400         | 12000157   | Failed to query iam user or role        | Failed to query the IAM role and user.        | Initiate the call again or contact technical support.                                  |
| 400         | 12000163   | Failed to query Manager user.           | Failed to query the Manager user.             | Call the API again, contact technical support, or check whether the cluster is normal. |
| 400         | 12000164   | Failed to query Manager user group.     | Failed to query the Manager user group.       | Call the API again, contact technical support, or check whether the cluster is normal. |
| 400         | 12000209   | Either a VPC ID or name is required.    | Either the VPC ID or the name is required.    | Rectify the fault based on the error message.                                          |
| 400         | 12000210   | Either a subnet ID or name is required. | Either the subnet ID or the name is required. | Rectify the fault based on the error message.                                          |

| Status Code | Error Code | Error Message                                                                                                                                                                                        | Description                                                                                                                             | Measure                                                                                                                                           |
|-------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| 400         | 12000233   | Insufficient resources for flavor xxx. Reduce the purchase quantity and try again. Alternatively, select another instance type or flavor, or switch the AZ or region to select your desired product. | Insufficient resources for specification xxx. Reduce the quantity, change service types and specifications, or change the region or AZ. | Reduce the purchase quantity and try again. Alternatively, select another instance with different types and flavors, or change the region and AZ. |
| 400         | 12000234   | Insufficient resources for the flavor xxx of the node to be scaled out. Reduce the purchase quantity and try again.                                                                                  | Insufficient resources for specification xxx. Reduce the quantity and try again.                                                        | Reduce the purchase quantity and try again, or change the specifications.                                                                         |
| 400         | 12000360   | Cluster name cannot be updated for a cluster that is in the terminating or terminated state.                                                                                                         | The name of a cluster in the <b>Deleting</b> or <b>Deleted</b> state cannot be changed.                                                 | Wait until the cluster task is complete and the status changes to <b>Running</b> .                                                                |
| 400         | 12000361   | Only offline cluster is allowed to modify manager access ip.                                                                                                                                         | The IP address for accessing FusionInsight Manager cannot be changed for a non-managed cluster.                                         | Rectify the fault based on the error message.                                                                                                     |

| Status Code | Error Code | Error Message                                                                                                                           | Description                                                                                                                 | Measure                                                    |
|-------------|------------|-----------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------|
| 400         | 12000362   | Cluster ID already exists or is not in UUID format.                                                                                     | The cluster ID already exists or is not in UUID format.                                                                     | Use correct request parameters based on the error message. |
| 400         | 12000363   | The offline cluster with same management node IPs[xxx] already exists.                                                                  | A management cluster with the same management node IP address [xxx] already exists.                                         | Modify the request parameters based on the error message.  |
| 400         | 12000405   | The number of nodes where the xxx role is deployed must be greater than or equal to xxx and less than or equal to xxx. The step is xxx. | The number of nodes to be deployed with the xxx role must be no less than xxx and no more than xxx. The step length is xxx. | Modify the request parameters based on the error message.  |
| 400         | 12000406   | The number of nodes where the xxx role is deployed must be greater than or equal to xxx and less than or equal to xxx.                  | The number of nodes to be deployed with the xxx role must be no less than xxx and no more than xxx.                         | Modify the request parameters based on the error message.  |
| 400         | 12000407   | The xxx role must be deployed on xxx nodes.                                                                                             | Role xxx must be deployed on xxx node(s).                                                                                   | Modify the request parameters based on the error message.  |

| Status Code | Error Code | Error Message                                                          | Description                                                                       | Measure                                                                            |
|-------------|------------|------------------------------------------------------------------------|-----------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| 400         | 12003001   | The service xxx was installed.                                         | The xxx service has been installed.                                               | Modify the request parameters based on the error message.                          |
| 400         | 12003002   | The service xxx is not in version metadata.                            | The xxx service is not described in metadata.                                     | Initiate the call again or contact technical support.                              |
| 400         | 12003003   | The group %s is not exist                                              | Node group %s does not exist.                                                     | Use an existing node group in the cluster.                                         |
| 400         | 12003004   | The cluster does not support add components.                           | Failed to add the service.                                                        | Initiate the call again or contact technical support.                              |
| 400         | 12003005   | The role xxx is not belong to the service xxx.                         | The xxx role does not belong to the xxx service.                                  | Modify the request parameters based on the error message.                          |
| 400         | 12003006   | The cluster is not in the [ running ] state.                           | The cluster is not in Running state.                                              | Wait until the cluster task is complete and the status changes to <b>Running</b> . |
| 400         | 12003008   | Occur service xxx that depend on xxx, can not add or delete.           | Adding or deleting is not allowed because xxx depends on xxx.                     | Modify the request parameters based on the error message.                          |
| 400         | 12003009   | Can not add services for nodegroup xxx that doesn't have normal hosts. | Failed to add services to node group xxx because all hosts are in abnormal state. | Initiate the call again or contact technical support.                              |



| Status Code | Error Code | Error Message                                                                                                                                                                         | Description                                                                                                                                       | Measure                                                   |
|-------------|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| 400         | 12003012   | xxx can not deploy to abnormal host, please check.                                                                                                                                    | xxx cannot be deployed on hosts in abnormal state. Please check.                                                                                  | Modify the request parameters based on the error message. |
| 400         | 12003013   | Can not deploy xxx to xxx, the host state is not normal.                                                                                                                              | Failed to deploy xxx on the host xxx because the host is in abnormal state.                                                                       | Modify the request parameters based on the error message. |
| 400         | 12003021   | Can not add the stateful roles [xxx] to the group whose type is task                                                                                                                  | Roles in xxx state cannot be added to Task node groups.                                                                                           | Check the node group type or role status.                 |
| 400         | 12003022   | There are multiple roles [xxx] whose local_disks_anti_affinity is true. In strict anti-affinity mode, a node group can have only one role with local_disks_anti_affinity set to true. | Roles whose <b>local_disks_anti_affinity</b> is <b>true</b> : xxx. In strict anti-affinity mode, only one such role can be added to a node group. | Modify the request parameters based on the error message. |
| 400         | 12003023   | The following components [xxx] cannot be added to the security cluster where Kerberos authentication is enabled.                                                                      | Components cannot be added to Kerberos authenticated security clusters: xxx                                                                       | Check the component or cluster type.                      |

| Status Code | Error Code | Error Message                                                                                                                                                          | Description                                                                                                                    | Measure                                                   |
|-------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| 400         | 12003054   | The xxx role is missing. Check the request parameters.                                                                                                                 | Role xxx required. Check the request parameters.                                                                               | Modify the request parameters based on the error message. |
| 400         | 12003090   | The index must be greater than zero and an integer.                                                                                                                    | The role index must be an integer greater than 0.                                                                              | Modify the request parameters based on the error message. |
| 400         | 12005001   | The number of tags in one cluster must not exceed xxx.                                                                                                                 | The number of tags in a cluster cannot exceed xxx.                                                                             | Modify the request parameters based on the error message. |
| 400         | 12005002   | Tag xxx does not comply with the specifications. The tag key and value cannot start or end with a space and cannot contain any of the following characters: =*<>\\, /. | Invalid tag: xxx. The key and value of the tag cannot start or end with spaces or contain the following characters: =*<>\\, /. | Modify the request parameters based on the error message. |
| 400         | 130000002  | The token is invalid.                                                                                                                                                  | Invalid token.                                                                                                                 | Update the token or check whether the token is correct.   |
| 400         | 13000046   | Security group xxx does not exist.                                                                                                                                     | Security group xxx does not exist.                                                                                             | Check whether the security group exists.                  |

| Status Code | Error Code   | Error Message                                                             | Description                                                                      | Measure                                                   |
|-------------|--------------|---------------------------------------------------------------------------|----------------------------------------------------------------------------------|-----------------------------------------------------------|
| 400         | MRS.00005045 | Role xxx must be deployed on the master node.                             | Role xxx must be deployed on master nodes.                                       | Modify the request parameters based on the error message. |
| 400         | MRS.00005058 | Current cluster version: xxx does not support config clickhouse password. | The ClickHouse password cannot be configured in the current cluster version xxx. | Modify the request parameters based on the error message. |
| 400         | MRS.00005059 | Security Cluster does not support config clickhouse password.             | The ClickHouse password cannot be configured for a security cluster.             | Modify the request parameters based on the error message. |
| 400         | MRS.00005060 | Current Service: xxx does not support config password.                    | The password cannot be configured in the current service xxx.                    | Modify the request parameters based on the error message. |
| 400         | MRS.00005061 | Clickhouse user password or default user password can not be empty.       | The password of ClickHouse or the default user cannot be left blank.             | Modify the request parameters based on the error message. |
| 400         | MRS.00005094 | Role xxx must be deployed on the non-master node.                         | Role xxx must be deployed on non-master nodes.                                   | Modify the request parameters based on the error message. |
| 400         | MRS.00005095 | Role xxx must be deployed on all nodes.                                   | Role xxx must be deployed on all nodes.                                          | Modify the request parameters based on the error message. |

| Status Code | Error Code       | Error Message                                                                                      | Description                                      | Measure                                                   |
|-------------|------------------|----------------------------------------------------------------------------------------------------|--------------------------------------------------|-----------------------------------------------------------|
| 400         | MRS.0000509<br>6 | The role metadata is abnormal.                                                                     | Role metadata is abnormal.                       | Initiate the call again or contact technical support.     |
| 400         | MRS.0000509<br>7 | The number of xxx role instances must be greater than or equal to 0 and less than or equal to xxx. | Role xxx must be deployed on 0 to xxx instances. | Modify the request parameters based on the error message. |

| Status Code | Error Code | Error Message        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Measure                                                  |
|-------------|------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|
| 400         | MRS.0010   | Dataconnector error. | Invalid connector name. The name can contain only uppercase and lowercase letters, digits, hyphens (-), and underscores(_). The connector ID cannot be empty.<br>Unsupported data type: xxx.<br>Unsupported component type: xxx. The data connection (ID: xxx) is not found.<br>The mapping between the cluster (ID: xxx) and the data connection is not found.<br>The mapping between the cluster (connection ID: xxx) and the data connection mapping is not found.<br>The RDS instance ID cannot be left blank. The instance ID, database name, username, | Locate and rectify the fault based on the error message. |

| Status Code | Error Code | Error Message | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Measure |
|-------------|------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|
|             |            |               | and password of RDS cannot be left blank. The AccessKey, SecretKey, and directory of OBS cannot be left blank. The IP address, port, database, username, and password of JDBC cannot be left blank. The driver that can be used by xxx is not found. Data connection type mismatch. The input type is xxx, but the actual data connection type is xxx. The VPC ID xxx of the RDS instance is inconsistent with the VPC ID xxx of the cluster. The subnet ID xxx of the RDS instance is inconsistent with the subnet ID xxx of the cluster. The security group ID xxx of the RDS instance is inconsistent |         |

| Status Code | Error Code | Error Message | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Measure |
|-------------|------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|
|             |            |               | <p>with the security group ID xxx of the cluster. A data connection (xxx) with the same location already exists in the cluster. The data connection has been used by another cluster (xxx). Data connection xxx already exists. The cluster cannot connect to the database of the RDS instance. Check the network, database name, username, and password. RDS instance (ID: xxx) cannot found. RDS instance xxx is abnormal. The status is xxx. Failed to synchronize data connection parameters to the cluster.</p> |         |

| Status Code | Error Code | Error Message                                                                                                                                                                                                                                                                                    | Description                                                                                                                                                                                                                                                                                                                                                          | Measure                                                                                             |
|-------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| 400         | MRS.0011   | SQL typed xxx can not run on the cluster which not installed dependent components. Cannot execute SQL on a cluster in the xxx state. Not support sql execution in cluster version xxx. Request with multiple SQL is not support. Failed to submit the SQL request to Executor (cluster ID: xxx). | Statement (SQL type: xxx) cannot run in the clusters where related dependent components are not installed. The cluster is in the xxx state. The SQL statement cannot be executed. The cluster (version xxx) does not support the SQL API. Multiple SQL statements can be submitted at a time. Failed to submit the SQL statement to Executor. The cluster ID is xxx. | Modify the request parameters based on the error message.                                           |
| 400         | MRS.0015   | The scaling operation failed due to not meeting the scaling conditions.                                                                                                                                                                                                                          | The scaling request fails because the scaling conditions are not met.                                                                                                                                                                                                                                                                                                | Check the cause details in the response body or contact technical support.                          |
| 400         | MRS.0016   | The request failed due to restrictions related to the cluster node group.                                                                                                                                                                                                                        | The request fails because the cluster node group does not meet a specific condition.                                                                                                                                                                                                                                                                                 | Rectify the fault based on the error information in the response body or contact technical support. |



| Status Code | Error Code | Error Message                                             | Description                                                            | Measure                                                        |
|-------------|------------|-----------------------------------------------------------|------------------------------------------------------------------------|----------------------------------------------------------------|
| 400         | MRS.0020   | The cluster does not support add components.              | The cluster does not support the function of adding components.        | Check the cluster version or type.                             |
| 400         | MRS.0205   | Failed to sync agency mapping configuration to cluster.   | Failed to synchronize the agency mapping configuration to the cluster. | Check whether the agency is correct.                           |
| 400         | MRS.0206   | Updating agency mapping task is running.                  | Updating the mapping configuration task...                             | Perform the operation after the current operation is complete. |
| 400         | MRS.0207   | Parse Json format failed.                                 | Failed to parse the JSON file.                                         | Check whether the JSON file is correct.                        |
| 400         | MRS.0208   | Create or modify policy failed.                           | Failed to create or modify the policy.                                 | Check whether the policy is correct.                           |
| 400         | MRS.0209   | Assign policy to agency failed.                           | Failed to assign the policy to the agency.                             | Check whether the agency is correct.                           |
| 400         | MRS.0210   | No secu_admin policy.                                     | The secu_admin permission is required.                                 | Obtain related permissions based on the error message.         |
| 400         | MRS.0211   | Failed to obtain new agency or new agency does not exist. | Failed to obtain the new agency or the agency does not exist.          | Check whether the agency is correct.                           |

| Status Code | Error Code | Error Message                                                                                    | Description                                                                                     | Measure                                                                 |
|-------------|------------|--------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| 400         | MRS.0212   | Updating ECS agency task is running.                                                             | Updating the ECS agency...                                                                      | Initiate the call again or contact technical support.                   |
| 400         | MRS.0216   | Failed to update ECS metadata.                                                                   | Failed to update ECS metadata.                                                                  | Initiate the call again or contact technical support.                   |
| 400         | MRS.0217   | Failed to unbind policy.                                                                         | Failed to unbind the policy.                                                                    | Check whether the policy is correct.                                    |
| 400         | MRS.0218   | Failed to check whether policies exist.                                                          | Failed to verify that the policy exists.                                                        | Check whether the policy is correct.                                    |
| 400         | MRS.1010   | The RDS instance whose ID is xxx is not found. The status (xxx) of RDS instance xxx is abnormal. | The RDS instance whose ID is xxx is not found. RDS instance xxx is abnormal. The status is xxx. | Rectify the fault based on the error message.                           |
| 401         | 12000001   | Invalid authentication                                                                           | Authentication failed                                                                           | Check whether the account status and authentication content are normal. |

| Status Code | Error Code | Error Message                                                                            | Description                                                                                                                 | Measure                                                                                                                       |
|-------------|------------|------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| 401         | 12000136   | Permission denied. Error message: Policy doesn't allow bss:order:update to be performed. | You do not have permission to perform this operation. Error message: Policy doesn't allow bss:order:update to be performed. | Obtain related permissions based on the error message.                                                                        |
| 404         | 12000057   | Failed to obtain the file list.                                                          | Failed to obtain the file list.                                                                                             | Check whether the MRS cluster status, Master node status, cluster network communication, and security group rules are normal. |
| 404         | 12005003   | The tag key xxx does not exist in cluster xxx.                                           | Tag key xxx does not exist in cluster xxx.                                                                                  | Modify the request parameters based on the error message.                                                                     |
| 500         | 12000004   | Internal server error.                                                                   | Internal server error.                                                                                                      | Initiate the call again or contact technical support.                                                                         |

## 10.5 Obtaining a Project ID

### Obtaining a Project ID from the Management Console

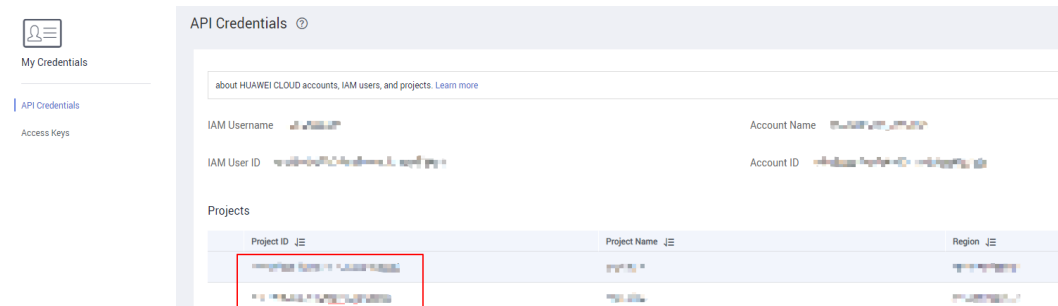
A project ID (**project\_id**) is required for some URLs when an API is called. To obtain a project ID, perform the following operations:

1. Log in to the management console.
2. Click the username and choose **My Credentials** from the drop-down list.

On the **API Credentials** page, view project IDs in the project list.

If there are multiple projects in one region, expand **Region** and view subproject IDs in the **Project ID** column.

**Figure 10-1** Viewing project IDs



## Obtaining a Project ID by Calling an API

You can obtain the project ID by calling the API used to [query project information based on the specified criteria](#).

The API for obtaining a project ID is **GET https://{Endpoint}/v3/projects/{Endpoint}**. **{Endpoint}** indicates the endpoint of IAM. For details, see [Endpoints](#).

The following is an example response. The value of **id** under **projects** is the project ID of the region specified by **name**.

```
{
  "projects": [
    {
      "domain_id": "65382450e8f64ac0870cd180d14e684b",
      "is_domain": false,
      "parent_id": "65382450e8f64ac0870cd180d14e684b",
      "name": "region_id",
      "description": "",
      "links": {
        "next": null,
        "previous": null,
        "self": "https://www.example.com/v3/projects/a4a5d4098fb4474fa22cd05f897d6b99"
      },
      "id": "a4a5d4098fb4474fa22cd05f897d6b99",
      "enabled": true
    }
  ],
  "links": {
    "next": null,
    "previous": null,
    "self": "https://www.example.com/v3/projects"
  }
}
```


## 10.6 Obtaining Account ID

An account ID (**domain-id**) is required for some URLs when an API is called. To obtain an account ID on the management console, perform the following steps:

1. Log in to the management console.
2. Click the username and choose **My Credentials** from the drop-down list.

On the displayed **API Credentials** page, view **Account ID**.

**Figure 10-2** Obtaining an account ID



## 10.7 Obtaining the MRS Cluster Information

### Components Supported by MRS

- MRS 3.2.0-LTS.1 supports the following components:
  - An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and Guardian.
  - A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.
  - A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, Flume, and Guardian.
  - A custom cluster contains the following components: CDL, Hadoop, Spark2x, HBase, Hive, Hue, IoTDB, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse, and Guardian.
- MRS 3.1.5 supports the following components:
  - An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Sqoop, and Guardian
  - A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.
  - A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Sqoop, Guardian, Kafka, and Flume.
  - A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Kafka, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, Kudu, Sqoop, and Guardian.
- MRS 3.1.2-LTS.3 supports the following components:
  - An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, and Tez.
  - A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.
  - A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, and Flume.
  - A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse.

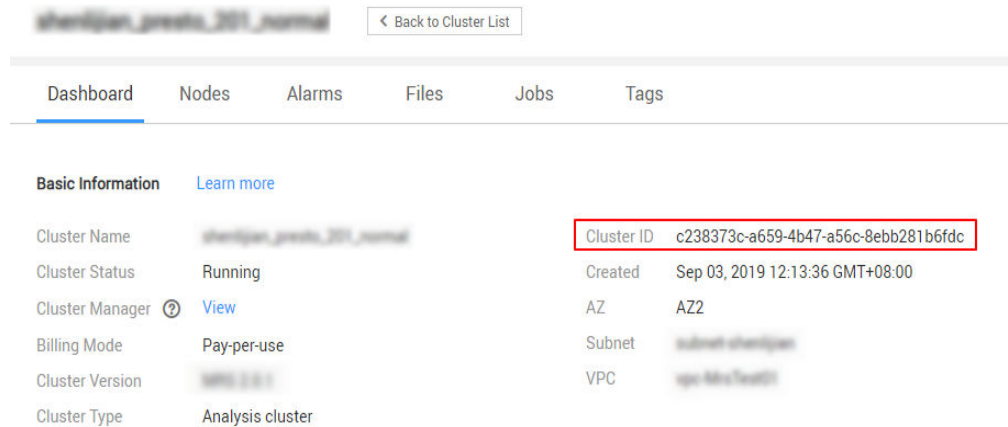
- MRS 3.1.0 supports the following components:
  - An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, and Kudu.
  - A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.
  - A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Kafka, and Flume.
  - A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Kafka, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, and Kudu.
- MRS 3.0.5 supports the following components:
  - An analysis cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, and Alluxio.
  - A streaming cluster contains the following components: Kafka, Storm, Flume, ZooKeeper, and Ranger.
  - A hybrid cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, Kudu, Alluxio, Kafka, Storm, and Flume.
  - A custom cluster contains the following components: Hadoop, Spark2x, HBase, Hive, Hue, Loader, Kafka, Storm, Flume, Flink, Oozie, ZooKeeper, Ranger, Tez, Impala, Presto, ClickHouse, Kudu, and Alluxio.
- MRS 2.1.0 supports the following components:
  - An analysis cluster contains the following components: Presto, Hadoop, Spark, HBase, Hive, Hue, Loader, Tez, Impala, Kudu, and Flink.
  - A streaming cluster contains the following components: Kafka, Storm, and Flume.
- MRS 1.9.2 supports the following components:
  - An analysis cluster contains the following components: Presto, Hadoop, Spark, HBase, OpenTSDB, Hive, Hue, Loader, Tez, Flink, Alluxio, and Ranger.
  - A streaming cluster contains the following components: Kafka, KafkaManager, Storm, and Flume.

## Obtaining a Cluster ID

A cluster ID (**cluster\_id**) is required for some URLs when an API is called. To obtain a cluster ID, perform the following operations:

1. Log in to the MRS management console.
2. On the **Active Clusters** page, and click the name of the cluster to be operated. The cluster details page is displayed.
3. Click the **Dashboard** tab and obtain the cluster ID in the **Basic Information** area.

**Figure 10-3** Obtaining a cluster ID

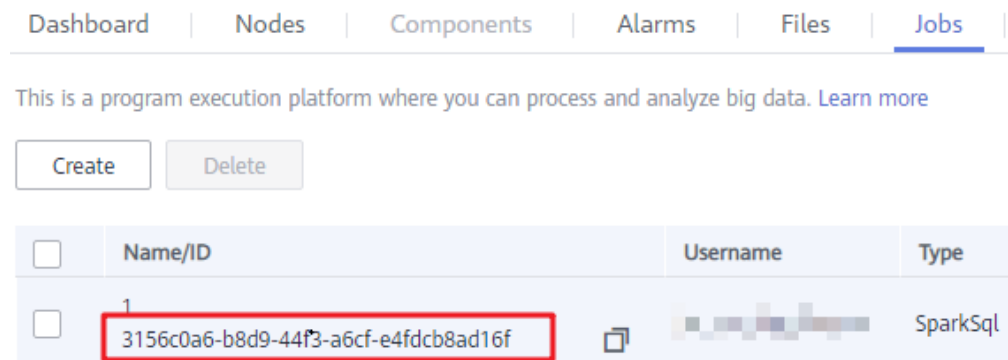


## Obtaining a Job ID

A job ID (**job\_execution\_id**) is required for some URLs when an API is called. To obtain a job ID, perform the following operations:

1. Log in to the MRS management console.
2. On the **Active Clusters** page, and click the name of the cluster to be operated. The cluster details page is displayed.
3. Click the **Jobs** tab and obtain the ID of the job to be operated from the job list.

**Figure 10-4** Obtaining a job ID



## 10.8 Roles and components supported by MRS

**Table 10-7** Roles and components supported by MRS

| Role Name | Component |
|-----------|-----------|
| OMSServer | OMSServer |
| NameNode  | HDFS      |
| Zkfc      | HDFS      |

| Role Name        | Component  |
|------------------|------------|
| JournalNode      | HDFS       |
| DataNode         | HDFS       |
| ResourceManager  | Yarn       |
| NodeManager      | Yarn       |
| JobHistoryServer | MapReduce  |
| quorumpeer       | ZooKeeper  |
| HMaster          | HBase      |
| ThriftServer     | HBase      |
| RegionServer     | HBase      |
| SlapdServer      | LdapServer |
| KerberosServer   | KrbServer  |
| KerberosAdmin    | KrbServer  |
| Hue              | Hue        |
| LoaderServer     | Loader     |
| JDBCServer       | Spark      |
| JobHistory       | Spark      |
| SparkResource    | Spark      |
| JDBCServer2x     | Spark2x    |
| JobHistory2x     | Spark2x    |
| SparkResource2x  | Spark2x    |
| MetaStore        | Hive       |
| WebHCat          | Hive       |
| HiveServer       | Hive       |
| MonitorServer    | Flume      |
| Flume            | Flume      |
| oozie            | Oozie      |
| KerberosClient   | KrbClient  |
| SlapdClient      | LdapClient |
| meta             | meta       |
| DBServer         | DBService  |



| <b>Role Name</b>   | <b>Component</b> |
|--------------------|------------------|
| Broker             | Kafka            |
| Supervisor         | Storm            |
| Logviewer          | Storm            |
| Nimbus             | Storm            |
| UI                 | Storm            |
| FlinkResource      | Flink            |
| ClickHouseServer   | ClickHouse       |
| ClickHouseBalancer | ClickHouse       |
| HSBroker           | HetuEngine       |
| HSConsole          | HetuEngine       |
| QAS                | HetuEngine       |
| CDLConnector       | CDL              |
| CDLService         | CDL              |
| IoTDBServer        | IoTDB            |
| ConfigNode         | IoTDB            |