



MapReduce Service

API Reference

Date 2024-11-30

Contents

1 Before You Start.....	1
2 API Overview.....	3
3 Selecting an API Type.....	8
4 Calling APIs.....	9
4.1 Making an API Request.....	9
4.2 Authentication.....	11
4.3 Response.....	13
5 Application Cases.....	15
5.1 Creating an MRS Cluster.....	15
5.2 Scaling Out a Cluster.....	17
5.3 Scaling in a Cluster.....	18
5.4 Creating a Job.....	19
5.5 Terminating a Job.....	20
5.6 Terminating a Cluster.....	21
6 API V2.....	23
6.1 Cluster Management APIs.....	23
6.1.1 Creating a Cluster.....	23
6.1.2 Changing a Cluster Name.....	78
6.1.3 Creating a Cluster and Submitting a Job.....	81
6.1.4 Scaling Out a Cluster.....	128
6.1.5 Scaling In a Cluster.....	131
6.1.6 Adding Components to a Cluster.....	134
6.1.7 Querying the Cluster Node List.....	140
6.2 Job Management APIs.....	152
6.2.1 Adding and Executing a Job.....	152
6.2.2 Querying Information About a Job.....	160
6.2.3 Querying a List of Jobs.....	166
6.2.4 Terminating a Job.....	178
6.2.5 Obtaining SQL Results.....	180
6.2.6 Deleting Jobs in Batches.....	183
6.3 Auto Scaling APIs.....	185

6.3.1 Viewing Auto Scaling Policies.....	185
6.3.2 Updating an Auto Scaling Policy.....	196
6.3.3 Deleting an AS policy.....	208
6.3.4 Creating an AS policy.....	210
6.4 Cluster HDFS File API.....	221
6.4.1 Obtaining the List of Files from a Specified Directory.....	221
6.5 SQL APIs.....	227
6.5.1 Submitting a SQL Statement.....	227
6.5.2 Querying SQL Results.....	232
6.5.3 Canceling a SQL Execution Task.....	236
6.6 Agency Management.....	239
6.6.1 Querying the Mapping Between a User (Group) and an IAM Agency.....	239
6.6.2 Updating the Mapping Between a User (Group) and an IAM Agency.....	243
6.7 Data Connection Management.....	247
6.7.1 Creating a Data Connection.....	248
6.7.2 Querying the Data Connection List.....	250
6.7.3 Updating a Data Connection.....	255
6.7.4 Deleting a Data Connection.....	258
6.8 Querying Version Metadata.....	259
6.8.1 Obtaining MRS Version List.....	259
6.8.2 Querying Available Specifications of an MRS Cluster Version.....	261
6.9 IAM Synchronization.....	265
6.9.1 Obtaining Synchronized IAM Users and User Groups.....	265
6.9.2 Synchronizing an IAM User and User Group.....	267
6.9.3 Cancelling Synchronization of Specified Users and User Groups.....	270
6.10 Tag Management APIs.....	273
6.10.1 Enabling or Disabling the Default Tag of a Cluster.....	273
6.10.2 Querying the Status of Default Cluster Tags.....	275
6.10.3 Querying Tag Quotas.....	278
7 API V1.1.....	281
7.1 Cluster Management APIs.....	281
7.1.1 Creating a Cluster and Executing a Job.....	281
7.1.2 Resizing a Cluster.....	338
7.1.3 Querying a Cluster List.....	349
7.1.4 Querying Cluster Details.....	374
7.1.5 Querying a Host List.....	396
7.1.6 Terminating a Cluster.....	403
7.2 Auto Scaling APIs.....	405
7.2.1 Configuring an Auto Scaling Rule.....	405
7.3 Tag Management APIs.....	420
7.3.1 Adding Tags to a Specified Cluster.....	420
7.3.2 Querying Tags of a Specified Cluster.....	423

7.3.3 Deleting Tags from a Specified Cluster.....	426
7.3.4 Adding Tags to a Cluster in Batches.....	428
7.3.5 Deleting Tags from a Cluster in Batches.....	431
7.3.6 Querying All Tags.....	435
7.3.7 Querying a List of Clusters with Specified Tags.....	437
7.4 Availability Zones.....	446
7.4.1 Querying AZ Information.....	446
7.5 Version Metadata.....	451
7.5.1 Querying the Metadata of a Cluster Version.....	451
8 Out-of-Date APIs.....	466
8.1 Job API Management (Deprecated).....	466
8.1.1 Adding and Executing a Job (Deprecated).....	466
8.1.2 Querying the exe Object List of Jobs (Deprecated).....	477
8.1.3 Querying exe Object Details (Deprecated).....	482
8.1.4 Deleting a Job Execution Object (Deprecated).....	485
9 Permissions Policies and Supported Actions.....	487
9.1 Introduction.....	487
10 Appendix.....	491
10.1 Status Codes.....	491
10.2 Error Codes.....	495
10.3 Obtaining a Project ID.....	532
10.4 Obtaining Tenant ID.....	533
10.5 Obtaining the MRS Cluster Information.....	533
10.6 Roles and components supported by MRS.....	534
10.7 Change History.....	536

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** in the *MapReduce Service User Guide*.

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. To obtain the regions and endpoints, contact technical support.

Concepts

- Account

An account is created upon successful registration with the cloud platform. 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. For security purposes, do not directly use the account to perform routine management but create IAM users and grant them permissions for routine management.

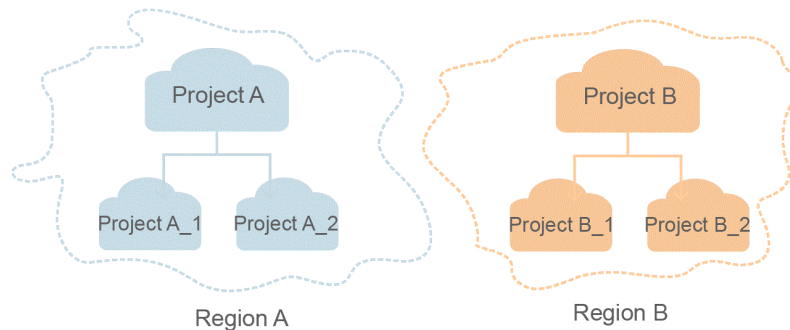
- User

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

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.
- **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 cloud service 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



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

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

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

Table 2-2 V1.1 APIs

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

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

Table 2-3 Historical APIs

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

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](#).
- **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.

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://{{endpoint}}/v3/auth/tokens
```

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.

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.

```
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.

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": "xxxxxxx"
      }
    }
  }
}
```

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:

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 4-1 Response header

```
"X-Frame-Options": "SAMEORIGIN",
"X-IAM-ETag-id": "2562365939-d8f6f12921974cb097338ac11fcec8a",
"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 [VPC > Querying VPCs](#) and [VPC > Creating a VPC](#). For details about how to create a subnet, see [Subnet > Querying Subnets](#) and [Subnet > 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.
- In this example, an analysis cluster is created.

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 3.3.1-LTS",
  "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,Loader,FTP-
Server,Solr,Flink,Oozie,GraphBase,HetuEngine,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 VPC > Querying VPCs and VPC > Creating a VPC in the <i>VPC API Reference</i> .
Subnet	See operation instructions in Subnet > Querying Subnets and Subnet > Creating a Subnet in the <i>VPC API Reference</i> .
Key Pair	See operation instructions in ECS SSH Key Management > Querying SSH Key Pairs and ECS SSH Key Management > Creating and Importing an SSH Key Pair in the <i>ECS API Reference</i> .
Zone	See Endpoints for details about regions and AZs.
Version	Currently, MRS 3.3.1-LTS is supported.

Resource	How to Obtain
Component	<p>MRS 3.3.1-LTS supports the following components:</p> <ul style="list-style-type: none"> • An analysis cluster contains the following components: Hadoop, Spark, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, JobGateway, Guardian, and Doris. • A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger. • A hybrid cluster contains the following components: Hadoop, Spark, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, Flume, JobGateway, Guardian, and Doris. • A custom cluster contains the following components: CDL, Hadoop, Spark, HBase, Hive, Hue, IoTDB, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse, Guardian, JobGateway, Doris and MemArtsCC.

URI

POST /v2/{project_id}/clusters

Table 6-2 URI parameter

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value None</p>

Request Parameters

Table 6-3 Request body parameters

Parameter	Mandatory	Type	Description
cluster_version	Yes	String	<p>Explanation Cluster version.</p> <p>Constraints None</p> <p>Value range MRS 3.3.1-LTS</p> <p>Default value N/A</p>
cluster_name	Yes	String	<p>Explanation Cluster name</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The cluster name must globally unique. • The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). <p>Default value N/A</p>
cluster_type	Yes	String	<p>Explanation Cluster type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • ANALYSIS: analysis cluster • STREAMING: streaming cluster • MIXED: hybrid cluster • CUSTOM: custom cluster, which is supported only by MRS 3.x. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
region	Yes	String	<p>Explanation Information about the region where the cluster is located. For details, see Endpoints.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
is_dec_project	No	Boolean	<p>Explanation Whether the cluster is specific for the DeC.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The resource is a DeC resource. • false: The resource is not a DeC resource. <p>Default value false</p>
vpc_name	Yes	String	<p>Explanation Name of the VPC where the subnet locates. Perform the following operations to obtain the VPC name from the VPC management console:</p> <ol style="list-style-type: none"> 1. Log in to the management console. 2. Choose Virtual Private Cloud > My VPCs. On the Virtual Private Cloud page, obtain the VPC name from the list. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
subnet_id	No	String	<p>Explanation</p> <p>Subnet ID, which can be obtained by performing the following operations on the VPC management console:</p> <ol style="list-style-type: none"> 1. Log in to the VPC management console. 2. Choose Virtual Private Cloud > My VPCs. 3. Locate the row that contains the target VPC and click the number in the Subnets column to view the subnet information. 4. Click the subnet name to obtain the network ID. <p>Constraints</p> <p>At least one of subnet_id and subnet_name must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. subnet_id is recommended.</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
subnet_name	Yes	String	<p>Subnet name. Perform the following operations to obtain the subnet name from the VPC management console:</p> <ol style="list-style-type: none"> 1. Log in to the management console. 2. Choose Virtual Private Cloud > My VPCs. 3. Locate the row that contains the target VPC and click the number in the Subnets column to obtain the subnet name. <p>Constraints At least one of subnet_id and subnet_name must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. If only subnet_name 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. subnet_id is recommended.</p> <p>Value range N/A</p> <p>Default value N/A</p>
components	Yes	String	<p>Explanation List of component names, which are separated by commas (.). For details about the component names, see the component list of each version in Table 6-1.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
external_datasources	No	Array of ClusterDataConnect or Map objects	<p>Explanation</p> <p>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 Table 6-4.</p> <p>Constraints</p> <p>N/A</p>
availability_zone	Yes	String	<p>Explanation</p> <p>AZ name. Multi-AZ clusters are not supported.</p> <p>See Endpoints for details about AZs.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>
security_groups_id	No	String	<p>Explanation</p> <p>Security group ID of the cluster.</p> <ul style="list-style-type: none"> • If this parameter is left blank, MRS automatically creates a security group, whose name starts with mrs_{cluster_name}. • If this parameter is specified, a fixed security group is used to create a cluster. The passed ID must be the security group ID in the current tenant. • Multiple security group IDs are supported and separated by commas (,). <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
auto_create_default_security_group	No	Boolean	<p>Explanation Whether to create the default security group for the MRS cluster.</p> <p>Constraints If this parameter is set to true, the default security group will be created for the cluster regardless of whether security_groups_id is specified.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The default security group is created for the MRS cluster. • false: The default security group is not created. <p>Default value false</p>
safe_mode	Yes	String	<p>Explanation Run mode of an MRS cluster.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SIMPLE: normal cluster. In a normal cluster, Kerberos authentication is disabled, and users can use all functions provided by the cluster. • KERBEROS: 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. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
manager_admin_password	Yes	String	<p>Explanation Password of the MRS Manager administrator.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • Must contain 8 to 26 characters. • Must contain at least four of the following: uppercase letters, lowercase letters, digits, and special characters (!@\$%^_-=+ [{}]:./?), but must not contain spaces. • Cannot be the username or the username spelled backwards. <p>Default value N/A</p>
login_mode	Yes	String	<p>Explanation Node login mode.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • PASSWORD: password-based login. If this value is selected, node_root_password cannot be left blank. • KEYPAIR: specifies the key pair used for login. If this value is selected, node_keypair_name cannot be left blank. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
node_root_password	No	String	<p>Explanation Password of user root for logging in to a cluster node.</p> <p>Constraints N/A</p> <p>Value range A password must meet the following requirements:</p> <ul style="list-style-type: none"> • Must be 8 to 26 characters long. • Must contain at least four of the following: uppercase letters, lowercase letters, digits, and special characters (!@#\$%^&*_+[]{};,:./?), but must not contain spaces. • Cannot be the username or the username spelled backwards. <p>Default value N/A</p>
node_keypair_name	No	String	<p>Explanation Name of a key pair You can use a key pair to log in to the Master node in the cluster.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
eip_address	No	String	<p>Explanation 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> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
eip_id	No	String	<p>Explanation ID of the bound EIP. This parameter is mandatory when eip_address is configured. To obtain the EIP ID, log in to the VPC console, choose Network > Elastic IP and Bandwidth > Elastic IP, click the EIP to be bound, and obtain the ID in the Basic Information area.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
mrs_ecs_default_agency	No	String	<p>Explanation</p> <p>Name of the agency bound to a cluster node by default. The value is fixed to MRS_ECS_DEFAULT_AGENCY.</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 MRS_ECS_DEFAULT_AGENCY agency has the OBS OperateAccess permission of OBS and the CES FullAccess (for users who have enabled fine-grained policies), CES Administrator, and KMS Administrator permissions in the region where the cluster is located.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
template_id	No	String	<p>Explanation</p> <p>Template used for node deployment when the cluster type is CUSTOM.</p> <ul style="list-style-type: none"> • mgmt_control_combined_v2: 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. • mgmt_control_separated_v2: 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. • mgmt_control_data_separated_v2: 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. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
tags	No	Array of tag objects	<p>Explanation Cluster tag For more parameter description, see Table 6-5.</p> <p>Constraints A cluster allows a maximum of 10 tags. A tag name (key) must be unique in a cluster.</p>
log_collection	No	Integer	<p>Explanation Whether to collect logs when cluster creation fails.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: Do not create an OBS bucket only for log collection when a cluster fails to be created. • 1: Create an OBS bucket only for collect logs when a cluster fails to be created. <p>Default value 1</p>
node_groups	Yes	Array of NodeGroupV2 objects	<p>Explanation Information about the node groups in the cluster. For details about the parameters, see Table 6-6.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
log_uri	No	String	<p>Explanation</p> <p>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 MRS_ECS_DEFAULT_AGENCY or customize an agency with the read and write permissions on the OBS path. This parameter is available only for cluster versions that support dumping cluster logs to OBS.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>
component_configs	No	Array of ComponentConfig objects	<p>Explanation</p> <p>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 ComponentConfig.</p> <p>Constraints</p> <p>The number of records cannot exceed 50.</p>
smn_notify	No	SmnNotify object	<p>Explanation</p> <p>SMN alarm notifications. For details about this parameter, see Table 6-17.</p> <p>Constraints</p> <p>N/A</p>

Table 6-4 ClusterDataConnectorMap

Parameter	Mandatory	Type	Description
map_id	No	Integer	<p>Explanation Data connection association ID</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
connector_id	No	String	<p>Explanation Data connection ID</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
component_name	No	String	<p>Explanation Component name</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
role_type	No	String	<p>Explanation Component role type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • hive_metastore: Hive Metastore role • hive_data: Hive role • hbase_data: HBase role • ranger_data: Ranger role <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
source_type	No	String	<p>Explanation Data connection type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • LOCAL_DB: local metadata • RDS_POSTGRES: RDS PostgreSQL database • RDS_MYSQL: RDS MySQL database • gaussdb-mysql: GaussDB(for MySQL) <p>Default value N/A</p>
cluster_id	No	String	<p>Explanation ID of the associated cluster</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
status	No	Integer	<p>Explanation Data connection status.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: normal • 1: in use <p>Default value N/A</p>

Table 6-5 Tag parameters

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • A tag key can contain letters, digits, spaces, and special characters <code>_:=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The tag key of a resource must be unique. • It can contain a maximum of 128 Unicode characters and cannot be an empty string. <p>Default value N/A</p>
value	Yes	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain letters, digits, spaces, and special characters <code>_:=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The value can contain a maximum of 255 Unicode characters and can be an empty string. <p>Default value N/A</p>

Table 6-6 NodeGroup parameters

Parameter	Mandatory	Type	Description
group_name	Yes	String	<p>Explanation Node group name.</p> <p>Constraints N/A</p> <p>Value range 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"> • master_node_default_group: master node group, which must be included in all cluster types. • core_node_analysis_group: analysis core node group, which must be included in both analysis and hybrid clusters. • core_node_streaming_group: streaming core node group, which must be included in both streaming and hybrid clusters. • task_node_analysis_group: analysis task node group, which can be selected for analysis clusters and hybrid clusters as needed. • task_node_streaming_group: streaming task node group, which can be selected for streaming clusters and hybrid clusters as needed. • node_group{x}: node group of a custom cluster. A maximum of nine such node groups can be added for a custom cluster. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
node_num	Yes	Integer	<p>Explanation Number of nodes.</p> <p>Constraints The total number of Core and Task nodes cannot exceed 500.</p> <p>Value range 0-500</p> <p>Default value N/A</p>
node_size	Yes	String	<p>Explanation Instance specifications of a node. Example: c3.4xlarge.2.linux.bigdata The host specifications supported by MRS are determined by CPU, memory, and disk space. Obtain the instance specifications of the corresponding version in the corresponding region from the cluster creation page of the MRS management console.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
root_volume	No	Volume object	<p>Explanation 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 Table 6-7.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
data_volume	No	Volume object	<p>Explanation Data disk information</p> <p>Constraints This parameter is mandatory when data_volume_count is not 0. For details about this parameter, see Table 6-7.</p>
data_volume_count	No	Integer	<p>Explanation Number of data disks of a node.</p> <p>Constraints N/A</p> <p>Value range 0-20</p> <p>Default value N/A</p>
auto_scaling_policy	No	auto_scaling_policy object	<p>Explanation Autoscaling rule corresponding to the node group. For details about the parameters, see Table 6-10.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
assigned_roles	No	Array of strings	<p>Explanation</p> <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"> • If the role is deployed on all nodes in the node group, set this parameter to <i><role name></i>, for example, DataNode. • If the role is deployed on a specified subscript node in the node group: <i><role name>:<index1>,<index2>..., <indexN></i>, for example, NameNode:1,2. The subscript starts from 1. <p>For details about available roles, see Roles and components supported by MRS.</p> <p>Constraints</p> <p>N/A</p>

Table 6-7 Volume

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Explanation Disk type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O disk • SAS: high I/O disk • SSD: ultra-high I/O disk • GPSSD: general-purpose SSD disk <p>Default value N/A</p>
size	Yes	Integer	<p>Explanation Data disk size in GB.</p> <p>Constraints N/A</p> <p>Value range 10-32768</p> <p>Default value N/A</p>

Table 6-8 BootstrapScript

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of a bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range The names of bootstrap action scripts in the same cluster must be unique. 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.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
uri	Yes	String	<p>Explanation</p> <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"> OBS bucket path: Enter a script path manually. For example, enter the path of the public sample script provided by MRS. Example: s3a://bootstrap/presto/presto-install.sh. If dualroles is installed, the parameter of the presto-install.sh script is dualroles. If worker is installed, the parameter of the presto-install.sh script is worker. Based on the Presto usage habit, you are advised to install dualroles on the active master nodes and worker on the Core nodes. Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
parameters	No	String	<p>Explanation</p> <p>Bootstrap action script parameters</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
nodes	Yes	Array of strings	<p>Explanation Name of the node group where the bootstrap action script is executed</p> <p>Constraints N/A</p>
active_master	No	Boolean	<p>Explanation Whether the bootstrap action script runs only on active Master nodes.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script runs only on active Master nodes. • false: The bootstrap action script can run on all Master nodes. <p>Default value N/A</p>
before_component_start	No	Boolean	<p>Explanation Time when the bootstrap action script is executed. Currently, the following two options are available: Before component start and After component start</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script is executed before the component starts. • false: The bootstrap action script is executed after the component starts. <p>Default value false</p>

Parameter	Mandatory	Type	Description
fail_action	Yes	String	<p>Explanation Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to be executed. You are advised to set this parameter to continue in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value N/A</p>
start_time	No	Long	<p>Explanation Execution time of one boot operation script.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
state	No	String	<p>Explanation The running status of one bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • PENDING • IN_PROGRESS • SUCCESS • FAILURE <p>Default value N/A</p>

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

Table 6-9 add_jobs parameters

Parameter	Mandatory	Type	Description
job_type	Yes	Integer	<p>Explanation Job type code</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • MapReduce • Spark • Hive Script • HiveSQL (not supported) • DistCp: imports and exports data (not supported). • Spark Script • Spark SQL: submits SQL statements (not supported). <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
job_name	Yes	String	<p>Explanation Job name</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). Identical job names are allowed but not recommended.</p> <p>Default value N/A</p>
jar_path	No	String	<p>Explanation Path of the .jar file or .sql file to be executed.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • Contains a maximum of 1023 characters, excluding special characters such as ; &><'\$. The parameter value cannot be empty or full of spaces. • Files can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> – OBS: The path must start with s3a://. Files or programs encrypted by KMS are not supported. – HDFS: The path starts with a slash (/). • Spark Script must end with .sql while MapReduce and Spark Jar must end with .jar. sql and jar are case-insensitive. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
arguments	No	String	<p>Explanation 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>Constraints N/A</p> <p>Value range The parameter can contain 0 to 150,000 characters, but special characters (; &>'<\$) are not allowed.</p> <p>Default value N/A</p>
input	No	String	<p>Explanation Address for inputting data. Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> • OBS: The path must start with s3a://. Files or programs encrypted by KMS are not supported. • HDFS: The path starts with a slash (/). <p>Constraints N/A</p> <p>Value range The parameter contains a maximum of 1023 characters, excluding special characters such as ; &>'<\$, and can be left blank.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
output	No	String	<p>Explanation</p> <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"> • OBS: The path must start with s3a://. • HDFS: The path starts with a slash (/). <p>If the specified path does not exist, the system will automatically create it.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The parameter contains a maximum of 1023 characters, excluding special characters such as ; &>'<\$, and can be left blank.</p> <p>Default value</p> <p>N/A</p>
job_log	No	String	<p>Explanation</p> <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"> • OBS: The path must start with s3a://. • HDFS: The path starts with a slash (/). <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The parameter contains a maximum of 1023 characters, excluding special characters such as ; &>'<\$, and can be left blank.</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
shutdown_cluster	No	Boolean	<p>Explanation Whether to delete the cluster after the job execution is complete.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Delete the cluster. • false: Do not delete the cluster. <p>Default value N/A</p>
file_action	No	String	<p>Explanation The action to be performed on a file.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • import: Import data. • export: Export data. <p>Default value N/A</p>
submit_job_once_cluster_run	Yes	Boolean	<p>Explanation Whether to submit a job when creating a cluster. Set this parameter to true in this example.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Submit a job during cluster creation. • false: Submit a job after the cluster is created. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
hql	No	String	<p>Explanation HiveQL statement</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
hive_script_path	No	String	<p>Explanation SQL program path. This parameter is needed by Spark Script and Hive Script jobs only,</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • Contains a maximum of 1023 characters, excluding special characters such as ; &><'\$. The address cannot be empty or full of spaces. • Files can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> – OBS: The path must start with s3a://. Files or programs encrypted by KMS are not supported. – HDFS: The path starts with a slash (/). • Ends with .sql. sql is case-insensitive. <p>Default value N/A</p>

Table 6-10 auto_scaling_policy parameters

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	<p>Explanation Whether to enable the auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable the auto scaling rule. • false: Disable the autoscaling rule. <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation Minimum number of nodes left in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation Maximum number of nodes in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
resources_plans	No	Array of resources_plan objects	<p>Explanation Resource plan list. For details, see Table 6-11. If this parameter is left blank, the resource plan is disabled.</p> <p>Constraints When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured.</p>
exec_scripts	No	Array of scale_script objects	<p>Explanation List of custom scaling automation scripts. For details, see Table 6-12. If this parameter is left blank, a hook script is disabled. This parameter is not available in the V2 API for creating and updating autoscaling policies.</p> <p>Constraints The number of records cannot exceed 10.</p>
rules	No	Array of rules objects	<p>Explanation List of auto scaling rules. For details, see Table 6-13.</p> <p>Constraints When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured. The number of records cannot exceed 10.</p>

Table 6-11 ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	<p>Explanation Cycle type of a resource plan. This parameter can be set to daily only.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
start_time	Yes	String	<p>Explanation The start time of a resource plan. The value is in the format of hour:minute, indicating that the time ranges from 00:00 to 23:59.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
end_time	Yes	String	<p>Explanation End time of a resource plan. The format is the same as that of start_time.</p> <p>Constraints The value cannot be earlier than the start_time, and the interval between start_time and start_time cannot be less than 30 minutes.</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
min_capacity	Yes	Integer	<p>Explanation Minimum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation Maximum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
effective_days	No	Array of strings	<p>Explanation 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: MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</p> <p>Constraints N/A</p>

Table 6-12 scale_script parameters

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Names of custom scaling automation scripts.</p> <p>Constraints N/A</p> <p>Value range The names in the same cluster must be unique. 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.</p> <p>Default value N/A</p>
uri	Yes	String	<p>Explanation 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"> • OBS bucket path: Enter a script path manually. for example, s3a://XXX/scale.sh. • Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
parameters	No	String	<p>Explanation</p> <p>Parameters of a custom automation script.</p> <ul style="list-style-type: none"> Multiple parameters are separated by space. The following predefined system parameters can be transferred: <ul style="list-style-type: none"> <i>\${mrs_scale_node_num}</i>: Number of the nodes to be added or removed <i>\${mrs_scale_type}</i>: Scaling type. The value can be scale_out or scale_in. <i>\$ {mrs_scale_node_hostnames }</i>: Host names of the nodes to be added or removed <i>\${mrs_scale_node_ips}</i>: IP addresses of the nodes to be added or removed <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. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
nodes	Yes	List<String>	<p>Explanation</p> <p>Type of a node where the custom automation script is executed. The node type can be Master, Core, or Task.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
active_master	No	Boolean	<p>Explanation Whether the custom automation script runs only on the active master node.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The custom automation script runs only on the active Master nodes. • false: The custom automation script can run on all Master nodes. <p>Default value false</p>
action_stage	Yes	String	<p>Explanation Time when a script is executed.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • before_scale_out: before scale-out • before_scale_in: before scale-in • after_scale_out: after scale-out • after_scale_in: after scale-in <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
fail_action	Yes	String	<p>Explanation Whether to continue to execute subsequent scripts and create a cluster after the custom automation script fails to be executed. You are advised to set this parameter to continue in the commissioning phase so the cluster can continue to be installed and started no matter whether the custom automation script is executed successfully.</p> <p>Constraints The scale-in operation cannot be undone. fail_action must be set to continue for the scripts that are executed after scale-in.</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value N/A</p>

Table 6-13 rules parameters

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). Rule names must be unique in a node group.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
description	No	String	<p>Explanation Description about an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 1024 characters.</p> <p>Default value N/A</p>
adjustment_type	Yes	String	<p>Explanation Adjustment type of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_out: cluster scale-out • scale_in: cluster scale-in <p>Default value N/A</p>
cool_down_minutes	Yes	Integer	<p>Explanation Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.</p> <p>Constraints N/A</p> <p>Value range The value ranges from 0 to 10080. 10080 indicates the number of minutes in a week.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
scaling_adjustment	Yes	Integer	<p>Explanation Number of nodes that can be adjusted once.</p> <p>Constraints N/A</p> <p>Value range 1-100</p> <p>Default value N/A</p>
trigger	Yes	Trigger object	<p>Explanation Condition for triggering a rule. For details, see Table 6-14.</p> <p>Constraints N/A</p>

Table 6-14 trigger parameters

Parameter	Mandatory	Type	Description
metric_name	Yes	String	<p>Explanation Metric name. This triggering condition makes a judgment according to the value of the metric.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 64 characters.</p> <p>Default value N/A</p>

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

Table 6-15 ComponentConfig

Parameter	Mandatory	Type	Description
component_name	Yes	String	<p>Explanation Component name</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
configs	No	Array of Config objects	<p>Explanation The component configuration item list. For details about this parameter, see Table 6-16.</p> <p>Constraints The number of records cannot exceed 100.</p>

Table 6-16 Config

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation The configuration name. Only the configuration names displayed on the MRS component configuration page are supported.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
value	Yes	String	<p>Explanation Configuration value</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
config_file_name	Yes	String	<p>Explanation Configuration file name. Only the file names displayed on the MRS component configuration page are supported.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Table 6-17 SmnNotify

Parameter	Mandatory	Type	Description
topic_urn	No	String	<p>Explanation SMN topic URN.</p> <p>Constraints Mandatory if alarm subscription needs to be enabled.</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
subscription_name	No	String	<p>Explanation Subscription rule name</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value default_alert_rule</p>

Response Parameters

Status code: 200

Table 6-18 Response parameters

Parameter	Type	Description
cluster_id	String	<p>Explanation Cluster ID, which is returned by the system after the cluster is created.</p> <p>Value range N/A</p>

Example Request

- Create an MRS 3.3.1-LTS 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 3.3.1-LTS",
  "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,Loader,FTP-Server,Solr,Flink,Oozie,GraphBase,HetuEngine,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 3.3.1-LTS 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 3.3.1-LTS",
  "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 3.3.1-LTS 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 3.3.1-LTS",
  "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,Storm,Spark2x,HBase,Hive,Hue,Loader,FTP-Server,Solr,Flink,Oozie,GraphBase,HetuEngine,Ranger,Tez,Kafka,Flume,Redis,Elasticsearch",
  "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 3.3.1-LTS. 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 3.3.1-LTS",
  "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,FTP-Server,Solr,Redis,Elasticsearch,Flink,Oozie,GraphBase,HetuEngine,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",
  "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 3.3.1-LTS. There a Master node group with five nodes and a Core node group with three nodes.

POST /v2/{project_id}/clusters

```

{
"cluster_version" : "MRS 3.3.1-LTS",
"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,FTP-
Server,Solr,Redis,Elasticsearch,Flink,Oozie,GraphBase,HetuEngine,Ranger,Tez,Ranger,Tez,ZooKeeper,Click
House",
"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", "RESTServer: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", "RESTServer", "FlinkResource" ]
} ]
}

```

- Create a cluster where data nodes are deployed independently from other nodes. The cluster version is MRS 3.3.1-LTS. 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 3.3.1-LTS",
  "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,FTP-
Server,Solr,Redis,Elasticsearch,Flink,Oozie,GraphBase,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",
  "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.

URI

PUT /v2/{project_id}/clusters/{cluster_id}/cluster-name

Table 6-19 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation The cluster ID. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-20 Request body parameter

Parameter	Mandatory	Type	Description
cluster_name	Yes	String	<p>Explanation New cluster name.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Response Parameters

Status code: 200

Table 6-21 Response body parameter

Parameter	Type	Description
result	String	<p>Explanation Result of the request for updating a mapping.</p> <p>Value range</p> <ul style="list-style-type: none"> • succeeded: The operation is successful. • failed: The operation failed.

Status code: 400

Table 6-22 Response body parameters

Parameter	Type	Description
error_code	String	<p>Explanation Error code.</p> <p>Value range N/A</p>

Parameter	Type	Description
error_msg	String	Explanation Error message. Value range N/A

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

URI

POST /v2/{project_id}/run-job-flow

Table 6-23 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Request Parameters

Table 6-24 Request body parameters

Parameter	Mandatory	Type	Description
is_dec_project	No	Boolean	<p>Explanation Whether the resource is a DeC resource.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The resource is a DeC resource. • false: The resource is not a DeC resource. <p>Default value false</p>

Parameter	Mandatory	Type	Description
cluster_version	Yes	String	<p>Explanation Cluster version, for example, MRS 3.3.1-LTS.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
cluster_name	Yes	String	<p>Explanation Cluster name.</p> <p>Constraints N/A</p> <p>Value range The cluster name must globally unique. A cluster name can contain only 1 to 64 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed.</p> <p>Default value N/A</p>
cluster_type	Yes	String	<p>Explanation The cluster type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> ● ANALYSIS: analysis cluster ● STREAMING: streaming cluster ● MIXED: hybrid cluster ● CUSTOM: custom cluster, which is supported only by MRS 3.x. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
region	Yes	String	<p>Explanation Information about the region where the cluster is located. For details, see Endpoints.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
vpc_name	Yes	String	<p>Explanation 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"> 1. Log in to the VPC management console. 2. Choose Virtual Private Cloud > My VPCs. On the Virtual Private Cloud page, obtain the VPC name from the list. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
subnet_id	No	String	<p>Explanation</p> <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"> 1. Log in to the VPC management console. 2. Choose Virtual Private Cloud > My VPCs. 3. Locate the row containing the target VPC and click the number in the Subnets column to view the subnet information. 4. Click the subnet name to obtain the network ID. <p>Constraints</p> <p>At least one of subnet_id and subnet_name must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. subnet_id is recommended.</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
subnet_name	Yes	String	<p>Explanation</p> <p>The subnet name. Obtain the subnet name by performing the following operations on the VPC management console:</p> <ol style="list-style-type: none"> 1. Log in to the management console. 2. Choose Virtual Private Cloud > My VPCs. 3. Locate the row that contains the target VPC and click the number in the Subnets column to obtain the subnet name. <p>Constraints</p> <p>At least one of subnet_id and subnet_name must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. If only subnet_name 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. subnet_id is recommended.</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
components	Yes	String	<p>Explanation List of component names, which are separated by commas (.). For details about the components that are supported, see "Components Supported by MRS" in Obtaining the MRS Cluster Information.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
external_data_sources	No	Array of ClusterDataConnectorMap objects	<p>Explanation 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 Table 6-25.</p> <p>Constraints N/A</p>
availability_zone	Yes	String	<p>Explanation The AZ name. Multi-AZ clusters are not supported. For details about AZs, see Endpoints.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
security_group_id	No	String	<p>Explanation</p> <p>The ID of the security group configured for the cluster.</p> <ul style="list-style-type: none"> • If this parameter is left blank, MRS automatically creates a security group whose name starts with mrs_{cluster_name}. • 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. • Multiple security group IDs are supported and separated by commas (,). <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
auto_create_default_security_group	No	Boolean	<p>Explanation Whether to create the default security group for the MRS cluster.</p> <p>Constraints If this parameter is set to true, the default security group will be created for the cluster regardless of whether security_groups_id is specified.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The default security group is created for the MRS cluster. • false: The default security group is not created. <p>Default value false</p>

Parameter	Mandatory	Type	Description
safe_mode	Yes	String	<p>Explanation The running mode of an MRS cluster.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SIMPLE: normal cluster. In a normal cluster, Kerberos authentication is disabled, and users can use all functions provided by the cluster. • KERBEROS: 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. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
manager_ad min_password	Yes	String	<p>Explanation Password of the MRS Manager administrator.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value must contain 8 to 26 characters. • The value must contain at least four of the following: uppercase letters, lowercase letters, numbers, and special characters (!@\$%^_-=+[{ } ; , / ?), but must not contain spaces. • The value cannot be the username or the username spelled backwards. <p>Default value N/A</p>
login_mode	Yes	String	<p>Explanation Node login mode.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • PASSWORD: password-based login. If this value is selected, node_root_password cannot be left blank. • KEYPAIR: key pair used for login. If this value is selected, node_keypair_name cannot be left blank. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
node_root_password	No	String	<p>Explanation The password of user root for logging in to a cluster node.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • Must be 8 to 26 characters long. • Must contain at least four of the following: uppercase letters, lowercase letters, numbers, and special characters (!@\$%^&*_+[]{};:./?), but must not contain spaces. • Cannot be the username or the username spelled backwards. <p>Default value N/A</p>
node_keypair_name	No	String	<p>Explanation The name of a key pair. You can use a key pair to log in to a cluster node.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
enterprise_project_id	No	String	<p>Explanation Enterprise project ID. When you create a cluster, associate the enterprise project ID with the cluster. The default value is 0, indicating the default enterprise project.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value The default value is 0, indicating the default enterprise project.</p>
eip_address	No	String	<p>Explanation 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.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
eip_id	No	String	<p>Explanation ID of the bound EIP.</p> <p>Constraints ID of the bound EIP. This parameter is mandatory when eip_address is configured.</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
mrs_ecs_default_agency	No	String	<p>Explanation</p> <p>Name of the agency bound to a cluster node by default. The value is fixed to MRS_ECS_DEFAULT_AGENCY. 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 MRS_ECS_DEFAULT_AGENCY agency has the OBS OperateAccess permission of OBS and the CES FullAccess (for users who have enabled fine-grained policies), CES Administrator, and KMS Administrator permissions in the region where the cluster is located.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
template_id	No	String	<p>Explanation</p> <p>The template used for node deployment when the cluster type is CUSTOM.</p> <ul style="list-style-type: none"> • mgmt_control_combined_v2: 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 fewer than 100 nodes, reducing costs. • mgmt_control_separated_v2: 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. • mgmt_control_data_separated_v2: 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. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value</p>

Parameter	Mandatory	Type	Description
			N/A
tags	No	Array of Tag objects	<p>Explanation Cluster tag information. For details, see Table 6-26.</p> <p>Constraints A cluster allows a maximum of 10 tags. A tag name (key) must be unique in a cluster.</p>
log_collection	No	Integer	<p>Explanation Whether to collect logs when cluster creation fails.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: Do not create an OBS bucket only for log collection when a cluster fails to be created. • 1: Create an OBS bucket only for collect logs when a cluster fails to be created. <p>Default value 1</p>
node_groups	Yes	Array of NodeGroupV2 objects	<p>Explanation Information about the node groups that form the cluster. For details about the parameters, see Table 6-27.</p> <p>Constraints N/A</p>
bootstrap_scripts	No	Array of BootstrapScript objects	<p>Explanation The bootstrap action script. For details about the parameters, see Table 6-34.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
log_uri	No	String	<p>Explanation</p> <p>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 MRS_ECS_DEFAULT_AGENCY or customize an agency with the read and write permissions on the OBS path. This parameter is available only for cluster versions that support dumping cluster logs to OBS.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>
component_configs	No	Array of ComponentConfig objects	<p>Explanation</p> <p>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 Table 6-35.</p> <p>Constraints</p> <p>The number of records cannot exceed 50.</p>

Parameter	Mandatory	Type	Description
delete_when_no_steps	No	Boolean	<p>Explanation Whether to automatically delete the cluster after the job is complete.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The cluster is deleted after the job is complete. • false: The cluster is not deleted after the job is complete. <p>Default value false</p>
steps	Yes	Array of StepConfig objects	<p>Explanation The job list. For details about this parameter, see Table 6-37.</p> <p>Constraints The number of records cannot exceed 255.</p>

Table 6-25 ClusterDataConnectorMap

Parameter	Mandatory	Type	Description
map_id	No	Integer	<p>Explanation Data connection association ID</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
connector_id	No	String	<p>Explanation Data connection ID</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
component_name	No	String	<p>Explanation Component name</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
role_type	No	String	<p>Explanation Component role type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • hive_metastore: Hive Metastore role • hive_data: Hive role • hbase_data: HBase role • ranger_data: Ranger role <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
source_type	No	String	<p>Explanation Data connection type</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • LOCAL_DB: local metadata • RDS_POSTGRES: RDS PostgreSQL database • RDS_MYSQL: RDS MySQL database • gaussdb-mysql: GaussDB(for MySQL) <p>Default value N/A</p>
cluster_id	No	String	<p>Explanation ID of the associated cluster</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
status	No	Integer	<p>Explanation Data connection status.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: normal • 1: in use <p>Default value N/A</p>

Table 6-26 Tag

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • A tag key can contain letters, digits, spaces, and special characters <code>._:=-@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The tag key of a resource must be unique. • It can contain a maximum of 128 Unicode characters and cannot be an empty string. <p>Default value N/A</p>
value	Yes	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain letters, digits, spaces, and special characters <code>._:=-@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The value can contain a maximum of 255 Unicode characters and can be an empty string. <p>Default value N/A</p>

Table 6-27 NodeGroupV2

Parameter	Mandatory	Type	Description
group_name	Yes	String	<p>Explanation Node group name.</p> <p>Constraints N/A</p> <p>Value range 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"> • master_node_default_group: master node group, which must be included in all cluster types. • core_node_analysis_group: analysis core node group, which must be included in both analysis and hybrid clusters. • core_node_streaming_group: streaming core node group, which must be included in both streaming and hybrid clusters. • task_node_analysis_group: analysis task node group, which can be selected for analysis clusters and hybrid clusters as needed. • task_node_streaming_group: streaming task node group, which can be selected for streaming clusters and hybrid clusters as needed. • node_group{x}: node group of a custom cluster. A maximum of nine such node groups can be added for a custom cluster. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
node_num	Yes	Integer	<p>Explanation Number of nodes.</p> <p>Constraints The total number of Core and Task nodes cannot exceed 500.</p> <p>Value range 0-500</p> <p>Default value N/A</p>
node_size	Yes	String	<p>Explanation The instance specifications of a node, for example, c3.4xlarge.2.linux.bigdata. 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.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
root_volume	No	Volume object	<p>Explanation 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 Table 6-28.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
data_volume	No	Volume object	<p>Explanation Data disk information. For details about the parameter, see Table 6-28.</p> <p>Constraints This parameter is mandatory when data_volume_count is not 0.</p>
data_volume_count	No	Integer	<p>Explanation Number of data disks of a node.</p> <p>Constraints N/A</p> <p>Value range 0-20</p> <p>Default value N/A</p>
auto_scaling_policy	No	AutoScalingPolicy object	<p>Explanation The auto scaling rule information. For details about this parameter, see Table 6-29.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
assigned_roles	No	Array of strings	<p>Explanation</p> <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 string array. Each string represents a role expression. Role expression definition:</p> <ul style="list-style-type: none"> • If the role is deployed on all nodes in a node group, set this parameter to <i>{role name}</i>, for example, DataNode. • If the role is deployed on a specified subscript node in the node group: <i>{role name}:{index1},{index2}...,{indexN}</i>, for example, NameNode:1,2. The subscript starts from 1. • 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, EsNode[9]. For details about the available roles, see Roles and components supported by MRS. <p>Constraints</p> <p>N/A</p>

Table 6-28 Volume

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Explanation Disk type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O disk • SAS: high I/O disk • SSD: ultra-high I/O disk • GPSSD: general-purpose SSD disk <p>Default value N/A</p>
size	Yes	Integer	<p>Explanation Data disk size in GB.</p> <p>Constraints N/A</p> <p>Value range 10-32768</p> <p>Default value N/A</p>

Table 6-29 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]

Parameter	Mandatory	Type	Description
resources_plans	No	Array of ResourcesPlan objects	The resource plan list. If this parameter is left blank, the resource plan is disabled. When auto_scaling_enable is set to true , either this parameter or rules must be configured. For details about this parameter, see Table 6-30 .
rules	No	Array of Rule objects	The list of auto scaling rules. When auto_scaling_enable is set to true , either this parameter or resources_plans must be configured. For details about this parameter, see Table 6-31 .
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 Table 6-33 .

Table 6-30 ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	<p>Explanation Cycle type of a resource plan. This parameter can be set to daily only.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
start_time	Yes	String	<p>Explanation Start time of a resource plan. The value is in the format of hour:minute, indicating that the time ranges from 00:00 to 23:59.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
end_time	Yes	String	<p>Explanation End time of a resource plan. The format is the same as that of start_time.</p> <p>Constraints The value cannot be earlier than the start_time, and the interval between start_time and start_time cannot be less than 30 minutes.</p> <p>Value range N/A</p> <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation Minimum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
max_capacity	Yes	Integer	<p>Explanation Maximum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
effective_days	No	Array of strings	<p>Explanation 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: MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</p> <p>Constraints N/A</p>

Table 6-31 Rule

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). Rule names must be unique in a node group.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
description	No	String	<p>Explanation Description about an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 1024 characters.</p> <p>Default value N/A</p>
adjustment_type	Yes	String	<p>Explanation Adjustment type of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_out: cluster scale-out • scale_in: cluster scale-in <p>Default value N/A</p>
cool_down_minutes	Yes	Integer	<p>Explanation The cluster cooling time after an auto scaling rule is triggered, in minutes, during which period no auto scaling operation is performed.</p> <p>Constraints N/A</p> <p>Value range The value ranges from 0 to 10080. 10080 indicates the number of minutes in a week.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
scaling_adjustment	Yes	Integer	<p>Explanation Number of nodes that can be adjusted once.</p> <p>Constraints N/A</p> <p>Value range 1-100</p> <p>Default value N/A</p>
trigger	Yes	Trigger object	<p>Explanation Condition for triggering a rule. For details about this parameter, see Table 6-32.</p> <p>Constraints N/A</p>

Table 6-32 Trigger

Parameter	Mandatory	Type	Description
metric_name	Yes	String	<p>Explanation Metric name. This triggering condition makes a judgment according to the value of the metric.</p> <p>Constraints N/A</p> <p>Value range The value can contains 0 to 64 characters.</p> <p>Default value N/A</p>

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

Table 6-33 ScaleScript

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Names of custom scaling automation scripts.</p> <p>Constraints N/A</p> <p>Value range The names in the same cluster must be unique. 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.</p> <p>Default value N/A</p>
uri	Yes	String	<p>Explanation 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"> • OBS bucket path: Enter a script path, for example, obs://XXX/scale.sh. • Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
parameters	No	String	<p>Explanation 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"> • <i>`\${mrs_scale_node_num}`</i>: The number of nodes to be added or removed • <i>`\${mrs_scale_type}`</i>: The scaling type. The value can be scale_out or scale_in. • <i>`\${mrs_scale_node_hostname s}`</i>: Host names of the nodes to be added or removed • <i>`\${mrs_scale_node_ips}`</i>: IP addresses of the nodes to be added or removed • <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. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
nodes	Yes	Array of strings	<p>Explanation Name of the node group where the custom automation script is executed.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
active_master	No	Boolean	<p>Explanation Whether the custom automation script runs only on the active Master node.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The custom automation script runs only on the active Master nodes. • false: The custom automation script can run on all Master nodes. <p>Default value false</p>
fail_action	Yes	String	<p>Explanation Whether to continue executing subsequent scripts and creating a cluster after the custom automation script fails to be executed. You are advised to set this parameter to continue in the commissioning phase so the cluster can continue to be installed and started no matter whether the custom automation script is executed successfully.</p> <p>Constraints The scale-in operation cannot be undone. fail_action must be set to continue for the scripts that are executed after scale-in.</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
action_stage	Yes	String	<p>Explanation Time when a script is executed.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • before_scale_out: before scale-out • before_scale_in: before scale-in • after_scale_out: after scale-out • after_scale_in: after scale-in <p>Default value N/A</p>

Table 6-34 BootstrapScript

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of a bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range The names of bootstrap action scripts in the same cluster must be unique. 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.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
uri	Yes	String	<p>Explanation 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: obs://bootstrap/presto/presto-install.sh. If dualroles is installed, the parameter of the presto-install.sh script is dualroles. If worker is installed, the parameter of the presto-install.sh script is worker. Based on the Presto usage habit, you are advised to install dualroles on the active master nodes and worker on the core nodes. Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
parameters	No	String	<p>Explanation Bootstrap action script parameters</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
nodes	Yes	Array of strings	<p>Explanation Name of the node group where the bootstrap action script is executed</p> <p>Constraints N/A</p>
active_master	No	Boolean	<p>Explanation Whether the bootstrap action script runs only on active master nodes.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script runs only on active Master nodes. • false: The bootstrap action script can run on all Master nodes. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
fail_action	Yes	String	<p>Explanation Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to execute. The default value is errorout, indicating that the action is stopped. Note: You are advised to set this parameter to continue in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value errorout</p>
before_component_start	No	Boolean	<p>Explanation Time when the bootstrap action script is executed. Currently, the following two options are available: Before component start and After component start.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script is executed before the component starts. • false: The bootstrap action script is executed after the component starts. <p>Default value false</p>

Parameter	Mandatory	Type	Description
start_time	No	Long	<p>Explanation Execution time of one boot operation script.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
state	No	String	<p>Explanation Running state of an individual bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • PENDING: The script is suspended. • IN_PROGRESS: The script is being processed. • SUCCESS • FAILURE: The script fails to be executed. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
action_stages	No	Array of strings	<p>Explanation Select the time when the bootstrap action script is executed.</p> <p>Constraints Enumerated values:</p> <ul style="list-style-type: none"> • BEFORE_COMPONENT_FIR ST_START: before initial component starts • AFTER_COMPONENT_FIRS T_START: after initial component starts • BEFORE_SCALE_IN: before scale-in • AFTER_SCALE_IN: after scale-in • BEFORE_SCALE_OUT: before scale-out • AFTER_SCALE_OUT: after scale-out

Table 6-35 ComponentConfig

Parameter	Mandatory	Type	Description
component_name	Yes	String	<p>Explanation Component name</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
configs	No	Array of Config objects	<p>Explanation Component configuration item list. For details about this parameter, see Table 6-36.</p> <p>Constraints The number of records cannot exceed 100.</p>

Table 6-36 Config

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Configuration name. Only the configuration names displayed on the MRS component configuration page are supported.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
value	Yes	String	<p>Explanation Configuration value.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
config_file_name	Yes	String	<p>Explanation Configuration file name. Only the file names displayed on the MRS component configuration page are supported.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Table 6-37 StepConfig

Parameter	Mandatory	Type	Description
job_execution	Yes	JobExecution object	<p>Explanation Job parameter. For details about this parameter, see Table 6-38.</p> <p>Constraints N/A</p>

Table 6-38 JobExecution

Parameter	Mandatory	Type	Description
job_type	Yes	String	<p>Explanation Job type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • MapReduce • SparkSubmit • 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.) • HiveScript • HiveSql • DistCp: imports and exports data. • SparkScript • SparkSql • Flink <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
job_name	Yes	String	Explanation Job name. Constraints N/A Value range A cluster name can contain only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. Identical job names are allowed but not recommended. Default value N/A

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

Response Parameters

Status code: 200

Table 6-39 Response body parameter

Parameter	Type	Description
cluster_id	String	<p>Explanation</p> <p>Cluster ID, which is returned by the system after the cluster is created.</p> <p>Value range</p> <p>N/A</p>

Example Request

Create an MRS 3.3.1-LTS 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 3.3.1-LTS",
  "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 6-40 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-41 Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	<p>Explanation Node group name.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
count	Yes	Integer	<p>Explanation Number of nodes to add</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
skip_bootstrap_scripts	No	Boolean	<p>Explanation Whether to skip the specified bootstrap action during cluster creation on the new node during scale-out.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action specified during cluster creation is skipped. • false: The bootstrap action specified during cluster creation is not skipped. <p>Default value true</p>

Parameter	Mandatory	Type	Description
scale_without_start	No	Boolean	<p>Explanation Whether the components remain stopped on the added nodes after cluster scale-out.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Do not start components after scale-out. • false: Start components after scale-out. <p>Default value false</p>

Response Parameters

Status code: 200

Table 6-42 Response body parameter

Parameter	Type	Description
result	String	<p>Explanation Result of the request. Value succeeded indicates that the operation is successful, and value failed indicates that the operation fails.</p> <p>Value range</p> <ul style="list-style-type: none"> • succeeded: The operation is successful. • failed: The operation failed.
order_id	String	<p>Explanation Order ID.</p> <p>Value range N/A</p>

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 scale-out is successful.

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

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 6-43 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-44 Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	<p>Explanation Node group name.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
count	No	Integer	<p>Explanation Number of nodes to be removed. If a specified node is to be removed, this parameter can be left blank.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
resource_ids	No	Array of strings	<p>Explanation 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 resource_id by calling the host API.</p> <p>Constraints N/A</p>

Response Parameters

Status code: 200

Table 6-45 Response body parameter

Parameter	Type	Description
result	String	<p>Explanation Result of the request for updating a mapping.</p> <p>Value range</p> <ul style="list-style-type: none"> • succeeded: The operation is successful. • failed: The operation failed.

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

You can add components to custom clusters.

URI

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

Table 6-46 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-47 Request body parameters

Parameter	Mandatory	Type	Description
components_install_mode	Yes	Array of ComponentInstallMode objects	<p>Explanation Component installation details. For details about the parameters, see Table 6-48.</p> <p>Constraints N/A</p>

Table 6-48 ComponentInstallMode

Parameter	Mandatory	Type	Description
component	Yes	String	<p>Explanation Component name</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
node_groups	Yes	Array of AssignedNode Group objects	<p>Explanation Role deployment information of the component. For details, see Table 6-49.</p> <p>Constraints N/A</p>

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

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

Table 6-49 AssignedNodeGroup

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Node group name.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
assigned_roles	Yes	Array of strings	<p>Explanation</p> <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"> • If the role is deployed on all nodes in a node group, set this parameter to <i>{role name}</i>, for example, DataNode. • If the role is deployed on a specified subscript node in the node group: <i>{role name}:{index1},{index2}...,{indexN}</i>, for example, NameNode:1,2. The subscript starts from 1. • For details about the available roles, see Roles and Components Supported by MRS. <p>Constraints</p> <p>The number of records cannot exceed 1,000.</p>

Response Parameters

Status code: 200

Table 6-50 Response body parameter

Parameter	Type	Description
result	String	<p>Explanation</p> <p>Result of the request for updating a mapping.</p> <p>Value range</p> <ul style="list-style-type: none"> • succeeded: The operation is successful. • failed: The operation failed.

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 6-51 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Table 6-52 Query parameters.

Parameter	Mandatory	Type	Description
node_group	No	String	<p>Explanation Name of the node groups to be queried.</p> <p>Constraints N/A</p> <p>Value range The value can contain 3 to 14 characters.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
limit	No	Integer	<p>Explanation Number of records displayed on each page in the returned result.</p> <p>Constraints N/A</p> <p>Value range ≥1</p> <p>Default value 10</p>
offset	No	Integer	<p>Explanation Offset from which the job list starts to be queried.</p> <p>Constraints N/A</p> <p>Value range ≥1</p> <p>Default value 1</p>
node_name	No	String	<p>Explanation Node name. Fuzzy search is supported.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
sort_key	No	String	<p>Explanation Sorting key. You can set node names as the key.</p> <p>Constraints N/A</p> <p>Value range node_name: node name</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
sort_dir	No	String	<p>Explanation Sorting order. desc indicates descending order, and asc indicates ascending order.</p> <p>Constraints N/A</p> <p>Value range desc: descending order asc: ascending order</p> <p>Default value N/A</p>
query_node_detail	No	Boolean	<p>Explanation Whether to query node details. If this parameter is set to true, the interface performance may be affected.</p> <p>Constraints N/A</p> <p>Value range true: Query node details. false: Do not query node details.</p> <p>Default value false</p>
query_ecs_detail	No	Boolean	<p>Explanation Whether to query ECS details. ECS APIs may be called.</p> <p>Constraints N/A</p> <p>Value range true: Query ECS details. false: Do not query ECS details.</p> <p>Default value false</p>

Parameter	Mandatory	Type	Description
internal_ip	No	String	<p>Explanation Internal IP address.</p> <p>Constraints N/A</p> <p>Value range The value can contain 7 to 15 characters.</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-53 Response body parameter

Parameter	Type	Description
nodes	Array of ClusterNode objects	<p>Explanation Node list. For details, see Table 6-54.</p>
node_total	Integer	<p>Explanation Number of nodes.</p> <p>Value range N/A</p>

Table 6-54 ClusterNode

Parameter	Type	Description
node_name	String	<p>Explanation Node name, which is the same as that on Manager.</p> <p>Value range N/A</p>

Parameter	Type	Description
resource_id	String	<p>Explanation Resource ID. Unique identity of a node. For yearly/monthly nodes, the ID can be used to query bills.</p> <p>Value range N/A</p>
node_group_name	String	<p>Explanation Node group name.</p> <p>Value range N/A</p>
node_type	String	<p>Explanation Node type. Available values include Task, Core, and Master.</p> <p>Value range N/A</p>
billing_type	String	<p>Explanation on-period: yearly/monthly; on-quantity: pay-per-use.</p> <p>Value range N/A</p>
deployment_type	String	<p>Explanation Deployment type. The Server hosts are supported.</p> <p>Value range N/A</p>
server_info	ServerInfo object	<p>Explanation If the deployment type is Server, this parameter cannot be empty. For details, see Table 6-55.</p>
tags	Array of Tag objects	<p>Explanation Node tag. For details, see Table 6-57.</p>
node_detail	NodeDetail object	<p>Explanation Node monitoring information. This parameter is returned only after IAM synchronization. For details, see Table 6-58.</p>

Parameter	Type	Description
node_status	String	<p>Explanation Node status. Status of operations on the corresponding page.</p> <p>Value range N/A</p>
component_infos	Array of ComponentInfo objects	<p>Explanation Component instance information array. For details, see Table 6-59.</p>

Table 6-55 ServerInfo

Parameter	Type	Description
server_id	String	<p>Explanation Server ID.</p> <p>Value range N/A</p>
server_name	String	<p>Explanation Server name.</p> <p>Value range N/A</p>
server_type	String	<p>Explanation Server type. Available options are ECS and BMS.</p> <p>Value range N/A</p>
data_volumes	Array of VolumeInfo objects	<p>Explanation Data disk. For details, see Table 6-56.</p>
root_volume	VolumeInfo object	<p>Explanation System disk. For details, see Table 6-56.</p>
cpu_type	String	<p>Explanation CPU type. Available options are x86 and Arm.</p> <p>Value range N/A</p>

Parameter	Type	Description
cpu	String	Explanation Number of vCPUs. Value range N/A
mem	String	Explanation Memory size. Value range N/A
internal_ip	String	Explanation Internal IP address. Value range N/A

Table 6-56 VolumeInfo

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

Table 6-57 Tag

Parameter	Type	Description
key	String	Explanation Tag key Value range N/A
value	String	Explanation Tag value. Value range N/A

Table 6-58 NodeDetail

Parameter	Type	Description
running_status	String	Explanation Running status. Value range N/A
cpu_usage	String	Explanation CPU utilization rate. Value range N/A
memory_usage	String	Explanation Memory utilization rate. Value range N/A
disk_usage	String	Explanation Disk usage. Value range N/A
total_memory	String	Explanation Total memory. Unit: MB. Value range N/A

Parameter	Type	Description
available_memory	String	Explanation Available memory. Unit: MB. Value range N/A
total_hard_disk_space	String	Explanation Total hard disk space. Unit: GB Value range N/A
available_hard_disk_space	String	Explanation Available hard disk space. Unit: GB Value range N/A
network_read	String	Explanation Network read speed. Unit: Byte/s Value range N/A
network_write	String	Explanation Network write speed. Unit: Byte/s Value range N/A

Table 6-59 ComponentInfo

Parameter	Type	Description
id	String	Explanation Component ID Value range N/A
name	String	Explanation Component name Value range N/A

Parameter	Type	Description
instance_group_name	String	Explanation Name of the group to which the component belongs. Value range N/A
running_statuses	String	Explanation Running status. Value range N/A
ha_status	String	Explanation HA status. Value range N/A
config_status	String	Explanation Configuration status. Value range N/A
role_name	String	Explanation Role name. Value range N/A
role_short_name	String	Explanation Abbreviation of a role name. Value range N/A
role_type	String	Explanation Role type. Value range N/A
service_name	String	Explanation Service name. Value range N/A

Parameter	Type	Description
pair_name	String	Explanation Pair name. Value range N/A
relation_pairs	String	Explanation Associated pair Value range N/A
support_decom	String	Explanation Whether Decom is supported. Value range N/A
support_reinstall_all	String	Explanation Whether reinstallation is supported. Value range N/A
support_collect_stack_info	String	Explanation Whether stack information can be collected. Value range N/A

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 "Managing Active Clusters" > "Managing Jobs" > "Using OBS to Encrypt Data for Running Jobs" in *MapReduce Service User Guide* 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.

URI

POST /v2/{project_id}/clusters/{cluster_id}/job-executions

Table 6-60 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. For details on how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-61 Request parameters

Parameter	Mandatory	Type	Description
job_type	Yes	String	<p>Explanation The job type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • MapReduce • SparkSubmit • 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.) • HiveScript • HiveSql • DistCp, importing and exporting data • SparkScript • SparkSql • Flink <p>Default value N/A</p> <p>NOTE Spark, Hive, and Flink jobs can be added to only clusters that include Spark, Hive, and Flink components.</p>

Parameter	Mandatory	Type	Description
job_name	Yes	String	<p>Explanation Job name.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). Identical job names are allowed but not recommended.</p> <p>Default value N/A</p>
arguments	No	Array of strings	<p>Explanation 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>Constraints The value can contain a maximum of 150,000 characters. Special characters (; &>'<\$!\) are not allowed. This parameter can be left blank.</p> <p>NOTE</p> <ul style="list-style-type: none"> • 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. • A file path on OBS can start with obs://. To use this format to submit HiveScript or HiveSQL jobs, choose Components > Hive > Service Configuration on the cluster details page. Set Basic to All, and search for core.site.customized.configs. Add the endpoint configuration item fs.obs.endpoint of OBS and enter the endpoint corresponding to OBS in Value. For details, see Endpoints.

Parameter	Mandatory	Type	Description
properties	No	Map<String,String>	<p>Explanation Program system parameter.</p> <p>Constraints The parameter contains a maximum of 2,048 characters, excluding special characters such as <'&!\\, and can be left blank.</p>

Response Parameters

Status code: 200

Table 6-62 Response body parameter

Parameter	Type	Description
job_submit_result	JobSubmitResult object	<p>Explanation The job execution result. For details about the parameters, see Table 6-63.</p>

Table 6-63 JobSubmitResult parameters

Parameter	Type	Description
job_id	String	<p>Explanation Job ID</p> <p>Value range N/A</p>
state	String	<p>Explanation Job submission status.</p> <p>Value range</p> <ul style="list-style-type: none"> • COMPLETE: The job is submitted. • FAILED: Failed to submit the job.

Status code: 400

Table 6-64 Response body parameters

Parameter	Type	Description
error_code	String	Explanation Error code. Value range N/A
error_msg	String	Explanation Error message. Value range N/A

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

```
{
  "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=executuraa",
    "--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

```

{
  "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

```

{
  "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

```

{
  "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

```

{
  "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

```
{  
  "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

```
{  
  "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"  
  }  
}
```

- **Crearte 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.

URI

- Format
GET /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}
- Parameter description

Table 6-65 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
job_execution_id	Yes	String	<p>Explanation The job ID. For details about how to obtain the job ID, see Obtaining a Job ID.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 202

Table 6-66 Response parameter

Parameter	Type	Description
job_detail	JobQueryBean object	<p>Explanation Job details. For details about the parameter, see Table 6-67.</p>

Table 6-67 Job parameters

Parameter	Type	Description
job_id	String	<p>Explanation The job ID.</p> <p>Value range N/A</p>
user	String	<p>Explanation Name of the user who submits a job.</p> <p>Value range N/A</p>
job_name	String	<p>Explanation Job name.</p> <p>Value range N/A</p>
job_result	String	<p>Explanation Final result of a job.</p> <p>Value range</p> <ul style="list-style-type: none"> • FAILED: indicates that the job fails to be executed. • KILLED: indicates that the job is manually terminated during execution. • UNDEFINED: indicates that the job is being executed. • SUCCEDED: indicates that the job has been successfully executed.

Parameter	Type	Description
job_state	String	<p>Explanation Job execution status.</p> <p>Value range</p> <ul style="list-style-type: none"> • FAILED: indicates that the job failed. • KILLED: indicates that the job is terminated. • New: indicates that the job is created. • NEW_SAVING: indicates that the job has been created and is being saved. • SUBMITTED: indicates that the job is submitted. • ACCEPTED: indicates that the job is accepted. • RUNNING: indicates that the job is running. • FINISHED: indicates that the job is completed.
job_progress	Float	<p>Explanation Job execution progress.</p> <p>Value range N/A</p>
job_type	String	<p>Explanation Job type.</p> <p>Value range</p> <ul style="list-style-type: none"> • MapReduce • SparkSubmit: Select this value when you call an API to query SparkPython jobs. • HiveScript • HiveSql • DistCp: imports and exports data. • SparkScript • SparkSql • Flink

Parameter	Type	Description
started_time	Long	Explanation Time when a job starts to execute. Unit: milliseconds Value range N/A
submitted_time	Long	Explanation Time when a job is submitted. Unit: milliseconds Value range N/A
finished_time	Long	Explanation Time when a job was completed. Unit: milliseconds Value range N/A
elapsed_time	Long	Explanation Running duration of a job. Unit: milliseconds Value range N/A
arguments	String	Explanation Running parameter. Value range N/A
properties	String	Explanation Configuration parameter, which is used to configure -d parameters. Value range N/A
launcher_id	String	Explanation Real-life job ID. Value range N/A
app_id	String	Explanation Actual job ID. Value range N/A

Parameter	Type	Description
tracking_url	String	<p>Explanation</p> <p>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.</p> <p>Value range</p> <p>N/A</p>
queue	String	<p>Explanation</p> <p>The type of the resource queue used by a job.</p> <p>Value range</p> <p>N/A</p>

Status code: 500

Table 6-68 Response body parameters

Parameter	Type	Description
error_code	String	<p>Explanation</p> <p>Error code.</p> <p>Value range</p> <p>N/A</p>
error_msg	String	<p>Explanation</p> <p>Error message.</p> <p>Value range</p> <p>N/A</p>

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.

URI

- Format
GET /v2/{project_id}/clusters/{cluster_id}/job-executions
- Parameter description

Table 6-69 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Table 6-70 Query parameters

Parameter	Mandatory	Type	Description
job_name	No	String	<p>Explanation Job name.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 128 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
job_id	No	String	<p>Explanation Job ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, and hyphens (-).</p> <p>Default value N/A</p>
user	No	String	<p>Explanation Username.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 32 characters, including only letters, digits, hyphens (-), underscores (_), and periods (.), and cannot start with a digit.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
job_type	No	String	<p>Explanation Job type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • MapReduce • SparkSubmit • SparkSubmit: Select this value when you call an API to query SparkPython jobs. • HiveScript • HiveSql • DistCp: imports and exports data. • SparkScript • SparkSql • Flink <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
job_state	No	String	<p>Explanation The job execution status.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • FAILED: indicates that the job fails to be executed. • KILLED: indicates that the job is terminated. • New: indicates that the job is created. • NEW_SAVING: indicates that the job has been created and is being saved. • SUBMITTED: indicates that the job is submitted. • ACCEPTED: indicates that the job is accepted. • RUNNING: indicates that the job is running. • FINISHED: indicates that the job is completed. <p>Default value N/A</p>
job_result	No	String	<p>Explanation Job execution result.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • FAILED: indicates that the job fails to be executed. • KILLED: indicates that the job is manually terminated during execution. • UNDEFINED: indicates that the job is being executed. • SUCCEEDED: indicates that the job has been successfully executed. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
queue	No	String	<p>Explanation Resource queue type of a job.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
limit	No	String	<p>Explanation Number of records displayed on each page in the returned result.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value 10</p>
offset	No	String	<p>Explanation Offset from which the job list starts to be queried.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value 1</p>

Parameter	Mandatory	Type	Description
sort_by	No	String	<p>Explanation Sorting method of the returned result.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • asc: indicates that the returned results are sorted in ascending order. • desc: indicates that the returned results are sorted in descending order. <p>Default value desc</p>
submitted_time_begin	No	Long	<p>Explanation UTC timestamp after which a job is submitted, in milliseconds, for example, 1562032041362.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
submitted_time_end	No	Long	<p>Explanation UTC timestamp before which a job is submitted, in milliseconds, for example, 1562032041362.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 202

Table 6-71 Response body parameters

Parameter	Type	Description
total_record	Integer	Explanation Total number of records. Value range N/A
job_list	Array of JobQueryBean objects	Explanation The job list. For details about the parameters, see Table 6-72 .

Table 6-72 JobQueryBean

Parameter	Type	Description
job_id	String	Explanation Job ID. Value range N/A
user	String	Explanation Name of the user who submits the job. Value range N/A
job_name	String	Explanation Job name. Value range N/A

Parameter	Type	Description
job_result	String	<p>Explanation Final result of a job.</p> <p>Value range</p> <ul style="list-style-type: none"> ● FAILED: indicates that the job fails to be executed. ● KILLED: indicates that the job is manually terminated during execution. ● UNDEFINED: indicates that the job is being executed. ● SUCCEDED: indicates that the job has been successfully executed.
job_state	String	<p>Explanation Job execution status.</p> <p>Value range</p> <ul style="list-style-type: none"> ● FAILED: indicates that the job fails to be executed. ● KILLED: indicates that the job is terminated. ● New: indicates that the job is created. ● NEW_SAVING: indicates that the job has been created and is being saved. ● SUBMITTED: indicates that the job is submitted. ● ACCEPTED: indicates that the job is accepted. ● RUNNING: indicates that the job is running. ● FINISHED: indicates that the job is completed.
job_progress	Float	<p>Explanation Job execution progress.</p> <p>Value range N/A</p>

Parameter	Type	Description
job_type	String	<p>Explanation Job type.</p> <p>Value range</p> <ul style="list-style-type: none"> • MapReduce • SparkSubmit: Select this value when you call an API to query SparkPython jobs. • HiveScript • HiveSql • DistCp: imports and exports data. • SparkScript • SparkSql • Flink
started_time	Long	<p>Explanation Time when a job starts to execute. Unit: milliseconds</p> <p>Value range N/A</p>
submitted_time	Long	<p>Explanation Time when a job is submitted. Unit: milliseconds</p> <p>Value range N/A</p>
finished_time	Long	<p>Explanation Time when a job was completed. Unit: milliseconds</p> <p>Value range N/A</p>
elapsed_time	Long	<p>Explanation Running duration of a job. Unit: milliseconds</p> <p>Value range N/A</p>
arguments	String	<p>Explanation Running parameter.</p> <p>Value range N/A</p>

Parameter	Type	Description
properties	String	<p>Explanation Configuration parameter, which is used to configure -d parameters.</p> <p>Value range N/A</p>
launcher_id	String	<p>Explanation Actual job ID.</p> <p>Value range N/A</p>
app_id	String	<p>Explanation Actual job ID.</p> <p>Value range N/A</p>
tracking_url	String	<p>Explanation 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.</p> <p>Value range N/A</p>
queue	String	<p>Explanation Resource queue type of a job.</p> <p>Value range N/A</p>

Status code: 500

Table 6-73 Response body parameters

Parameter	Type	Description
error_code	String	Explanation Error code. Value range N/A
error_msg	String	Explanation Error message. Value range N/A

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" : "xxxx",
    "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, yujsjhe, --driver-memory, yueujddj, --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.

URI

- Format
POST /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/kill
- Parameter description

Table 6-74 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
job_execution_id	Yes	String	<p>Explanation Job ID. For details about how to obtain the job ID, see Obtaining a Job ID.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

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.

URI

- Format
GET /v2/{project_id}/clusters/{cluster_id}/job-executions/{job_execution_id}/sql-result
- Parameter description

Table 6-75 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
job_execution_id	Yes	String	<p>Explanation Job ID. For details about how to obtain the job ID, see Obtaining a Job ID.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 202

Table 6-76 Response parameter

Parameter	Type	Description
sql_results	Object	<p>Explanation The query results of a SQL statement.</p>

Status code: 400

Table 6-77 Response body parameters

Parameter	Type	Description
error_code	String	Explanation Error code. Value range N/A
error_msg	String	Explanation Error message. Value range N/A

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: 400

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.

URI

POST /v2/{project_id}/clusters/{cluster_id}/job-executions/batch-delete

Table 6-78 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID . Constraints N/A Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed. Default value N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-79 Request parameter

Parameter	Mandatory	Type	Description
job_id_list	No	Array of strings	<p>Explanation The job ID list. For details about how to obtain the job ID, see Obtaining a Job ID.</p> <p>Constraints N/A</p>

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.

URI

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

Table 6-80 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-81 Response body parameter

Parameter	Type	Description
[Array element]	Array of AutoScalingPolicyV2 objects	<p>Explanation The auto scaling policy list. For details, see Table 6-82.</p>

Table 6-82 AutoScalingPolicyV2

Parameter	Type	Description
node_group_name	String	<p>Explanation Node group name.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
resource_pool_name	String	<p>Explanation Pool name.</p> <p>Constraints If the cluster version does not support auto scaling for a specified resource pool, set this parameter to default.</p> <p>Value range The value can contain only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
auto_scaling_policy	AutoScalingPolicyInfo object	<p>Explanation The auto scaling policy. For details, see Table 6-83.</p> <p>Constraints N/A</p>

Table 6-83 AutoScalingPolicyInfo

Parameter	Type	Description
auto_scaling_enable	Boolean	<p>Explanation Whether to enable the auto scaling policy.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable the auto scaling rule. • false: Disable the autoscaling rule. <p>Default value N/A</p>
min_capacity	Integer	<p>Explanation Minimum number of nodes left in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Integer	<p>Explanation Maximum number of nodes in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
resources_plans	Array of ResourcesPlan objects	<p>Explanation Resource plan list. If this parameter is left blank, the resource plan is disabled.</p> <p>Constraints When auto_scaling_enable is set to true, either this parameter or rules must be configured. For details about this parameter, see Table 6-84.</p>

Parameter	Type	Description
rules	Array of Rule objects	<p>Explanation Auto scaling rules.</p> <p>Constraints When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured. The number of records cannot exceed 10. For details, see Table 6-85.</p>
tags	Array of Tag objects	<p>Explanation Auto scaling tags. For details, see Table 6-87.</p> <p>Constraints The number of records cannot exceed 100.</p>

Table 6-84 ResourcesPlan

Parameter	Type	Description
period_type	String	<p>Explanation Cycle type of a resource plan. This parameter can be set to daily only.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
start_time	String	<p>Explanation Start time of a resource plan. The value is in the format of hour:minute, indicating that the time ranges from 00:00 to 23:59.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Type	Description
end_time	String	<p>Explanation End time of a resource plan. The format is the same as that of start_time.</p> <p>Constraints The value cannot be earlier than the start_time, and the interval between start_time and start_time cannot be less than 30 minutes.</p> <p>Value range N/A</p> <p>Default value N/A</p>
min_capacity	Integer	<p>Explanation Minimum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Integer	<p>Explanation Maximum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Type	Description
effective_days	Array of strings	<p>Explanation 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: MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</p> <p>Constraints N/A</p>

Table 6-85 Rule

Parameter	Type	Description
name	String	<p>Explanation Name of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). Rule names must be unique in a node group.</p> <p>Default value N/A</p>
description	String	<p>Explanation Description about an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 1024 characters.</p> <p>Default value N/A</p>

Parameter	Type	Description
adjustment_type	String	<p>Explanation Adjustment type of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_out: cluster scale-out • scale_in: cluster scale-in <p>Default value N/A</p>
cool_down_minutes	Integer	<p>Explanation Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.</p> <p>Constraints N/A</p> <p>Value range The value ranges from 0 to 10080. 10080 indicates the number of minutes in a week.</p> <p>Default value N/A</p>
scaling_adjustment	Integer	<p>Explanation Number of nodes that can be adjusted once.</p> <p>Constraints N/A</p> <p>Value range 1-100</p> <p>Default value N/A</p>
trigger	Trigger object	<p>Explanation The condition for triggering a rule. For details, see Table 6-86.</p> <p>Constraints N/A</p>

Table 6-86 Trigger

Parameter	Type	Description
metric_name	String	<p>Explanation Metric name. This triggering condition makes a judgment according to the value of the metric.</p> <p>Constraints N/A</p> <p>Default value N/A</p>
metric_value	String	<p>Explanation Metric threshold to trigger a rule. The value must be an integer or a number with two decimal places.</p> <p>Constraints N/A</p> <p>Value range Only integers or numbers with two decimal places are allowed.</p> <p>Default value N/A</p>
comparison_operator	String	<p>Explanation Metric judgment logic operator.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • LT: less than • GT: greater than • LTOE: less than or equal to • GTOE: greater than or equal to <p>Default value N/A</p>

Parameter	Type	Description
evaluation_periods	Integer	<p>Explanation Number of consecutive five-minute periods, during which a metric threshold is reached</p> <p>Constraints N/A</p> <p>Value range 1-288</p> <p>Default value N/A</p>

Table 6-87 Tag

Parameter	Type	Description
key	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • A tag key can contain letters, digits, spaces, and special characters <code>._:=-@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The tag key of a resource must be unique. • It can contain a maximum of 128 Unicode characters and cannot be an empty string. <p>Default value N/A</p>

Parameter	Type	Description
value	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> The value can contain letters, digits, spaces, and special characters <code>_:+=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>. The value can contain a maximum of 255 Unicode characters and can be an empty string. <p>Default value N/A</p>

Status code: 400

Table 6-88 Response body parameters

Parameter	Type	Description
error_code	String	<p>Explanation Error code.</p> <p>Value range N/A</p>
error_msg	String	<p>Explanation Error message.</p> <p>Value range N/A</p>

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.

URI

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

Table 6-89 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-90 Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	<p>Explanation Node group name.</p> <p>Value range This parameter is mandatory. If resource_pool_name is set to default, the auto scaling policies are created by node group. If resource_pool_name is not set to default, policies of the resource pool corresponding to the node group are created.</p>
resource_pool_name	Yes	String	<p>Explanation Pool name.</p> <p>Value range This parameter is mandatory. If the cluster version does not support auto scaling for a specified resource pool, set this parameter to default. If this parameter is not default, the auto scaling policy of the specified resource pool is deleted.</p>
auto_scaling_policy	No	AutoScalingPolicyInfo object	<p>Explanation Auto scaling rule. For details about the parameters, see Table 6-91.</p>

Table 6-91 AutoScalingPolicyInfo

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	<p>Explanation Whether to enable the auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable the auto scaling rule. • false: Disable the autoscaling rule. <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation Minimum number of nodes left in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation Maximum number of nodes in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
resources_plans	No	Array of ResourcesPlan objects	<p>Explanation Resource plan list. If this parameter is left blank, the resource plan is disabled.</p> <p>Constraints When auto_scaling_enable is set to true, either this parameter or resources_plans must be configured. For details about this parameter, see Table 6-92.</p>
rules	No	Array of Rule objects	<p>Explanation Auto scaling rules.</p> <p>Constraints When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured. The number of records cannot exceed 10. For details, see Table 6-93.</p>
tags	No	Array of Tag objects	<p>Explanation Auto scaling tags. For details, see Table 6-95.</p> <p>Constraints The number of records cannot exceed 20.</p>

Table 6-92 ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	<p>Explanation Cycle type of a resource plan. This parameter can be set to daily only.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
start_time	Yes	String	<p>Explanation Start time of a resource plan. The value is in the format of hour:minute, indicating that the time ranges from 00:00 to 23:59.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
end_time	Yes	String	<p>Explanation End time of a resource plan. The format is the same as that of start_time.</p> <p>Constraints The value cannot be earlier than the start_time, and the interval between start_time and start_time cannot be less than 30 minutes.</p> <p>Value range N/A</p> <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation Minimum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
max_capacity	Yes	Integer	<p>Explanation Maximum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
effective_days	No	Array of strings	<p>Explanation 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: MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</p> <p>Constraints N/A</p>

Table 6-93 Rule

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). Rule names must be unique in a node group.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
description	No	String	<p>Explanation Description about an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 1024 characters.</p> <p>Default value N/A</p>
adjustment_type	Yes	String	<p>Explanation Adjustment type of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_out: cluster scale-out • scale_in: cluster scale-in <p>Default value N/A</p>
cool_down_minutes	Yes	Integer	<p>Explanation Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.</p> <p>Constraints N/A</p> <p>Value range The value ranges from 0 to 10080. 10080 indicates the number of minutes in a week.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
scaling_adjustment	Yes	Integer	<p>Explanation Number of nodes that can be adjusted once.</p> <p>Constraints N/A</p> <p>Value range 1-100</p> <p>Default value N/A</p>
trigger	Yes	Trigger object	<p>Explanation The condition for triggering a rule. For details, see Table 6-94.</p> <p>Constraints N/A</p>

Table 6-94 Trigger

Parameter	Mandatory	Type	Description
metric_name	Yes	String	<p>Explanation Metric name. This triggering condition makes a judgment according to the value of the metric.</p> <p>Constraints N/A</p> <p>Default value N/A</p>

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

Table 6-95 Tag

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • A tag key can contain letters, digits, spaces, and special characters <code>._:=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The tag key of a resource must be unique. • It can contain a maximum of 128 Unicode characters and cannot be an empty string. <p>Default value N/A</p>
value	Yes	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain letters, digits, spaces, and special characters <code>._:=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The value can contain a maximum of 255 Unicode characters and can be an empty string. <p>Default value N/A</p>

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.

URI

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

Table 6-96 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-97 Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	<p>Explanation Node group name.</p> <p>Constraints If resource_pool_name is set to default, the auto scaling policies are deleted by node group. If resource_pool_name is not set to default, policies of the resource pool corresponding to the node group are deleted.</p> <p>Value range N/A</p> <p>Default value N/A</p>
resource_pool_name	Yes	String	<p>Explanation Pool name.</p> <p>Constraints If the cluster version does not support auto scaling for a specified resource pool, set this parameter to default. If this parameter is not default, the auto scaling policy of the specified resource pool is deleted.</p> <p>Value range The value can contain only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

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.

URI

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

Table 6-98 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-99 Request body parameters

Parameter	Mandatory	Type	Description
node_group_name	Yes	String	<p>Explanation Node group name.</p> <p>Value range This parameter is mandatory. If resource_pool_name is set to default, the auto scaling policies are created by node group. If resource_pool_name is not set to default, policies of the resource pool corresponding to the node group are created.</p>

Parameter	Mandatory	Type	Description
resource_pool_name	Yes	String	<p>Explanation Pool name.</p> <p>Value range This parameter is mandatory. If the cluster version does not support auto scaling for a specified resource pool, set this parameter to default. If this parameter is not default, the auto scaling policy of the specified resource pool is deleted.</p>
auto_scaling_policy	No	AutoScalingPolicyInfo object	<p>Explanation Auto scaling rule. For details about the parameters, see Table 6-100.</p>

Table 6-100 AutoScalingPolicyInfo

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	<p>Explanation Whether to enable the auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable the auto scaling rule. • false: Disable the autoscaling rule. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
min_capacity	Yes	Integer	<p>Explanation Minimum number of nodes left in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation Maximum number of nodes in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
resources_plans	No	Array of ResourcesPlan objects	<p>Explanation Resource plans. For details, see Table 6-101. If this parameter is left blank, the resource plan is disabled.</p> <p>Constraints When auto_scaling_enable is set to true, either this parameter or rules must be configured. A maximum of five records are allowed.</p>
rules	No	Array of Rule objects	<p>Explanation Auto scaling rules. For details, see Table 6-102.</p> <p>Constraints When auto_scaling_enable is set to true, either this parameter or rules must be configured. The number of records cannot exceed 10.</p>

Parameter	Mandatory	Type	Description
tags	No	Array of Tag objects	<p>Explanation Tags of an auto scaling rule. For details about the parameters, see Table 6-104.</p> <p>Constraints The number of records cannot exceed 100.</p>

Table 6-101 ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	<p>Explanation Cycle type of a resource plan. This parameter can be set to daily only.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
start_time	Yes	String	<p>Explanation Start time of a resource plan. The value is in the format of hour:minute, indicating that the time ranges from 00:00 to 23:59.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
end_time	Yes	String	<p>Explanation End time of a resource plan. The format is the same as that of start_time.</p> <p>Constraints The value cannot be earlier than the start_time, and the interval between start_time and start_time cannot be less than 30 minutes.</p> <p>Value range N/A</p> <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation Minimum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation Maximum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
effective_days	No	Array of strings	<p>Explanation 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: MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, and SUNDAY</p> <p>Constraints N/A</p>

Table 6-102 Rule

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-). Rule names must be unique in a node group.</p> <p>Default value N/A</p>
description	No	String	<p>Explanation Description about an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 1024 characters.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
adjustment_type	Yes	String	<p>Explanation Adjustment type of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_out: cluster scale-out • scale_in: cluster scale-in <p>Default value N/A</p>
cool_down_minutes	Yes	Integer	<p>Explanation Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.</p> <p>Constraints N/A</p> <p>Value range The value ranges from 0 to 10080. 10080 indicates the number of minutes in a week.</p> <p>Default value N/A</p>
scaling_adjustment	Yes	Integer	<p>Explanation Number of nodes that can be adjusted once.</p> <p>Constraints N/A</p> <p>Value range 1-100</p> <p>Default value N/A</p>
trigger	Yes	Trigger object	<p>Explanation Condition for triggering a rule. For details, see Table 6-103.</p> <p>Constraints N/A</p>

Table 6-103 Trigger

Parameter	Mandatory	Type	Description
metric_name	Yes	String	<p>Explanation Metric name. This triggering condition makes a judgment according to the value of the metric.</p> <p>Constraints N/A</p> <p>Default value N/A</p>
metric_value	Yes	String	<p>Explanation Metric threshold to trigger a rule. The value must be an integer or a number with two decimal places.</p> <p>Constraints N/A</p> <p>Value range Only integers or numbers with two decimal places are allowed.</p> <p>Default value N/A</p>
comparison_operator	No	String	<p>Explanation Metric judgment logic operator.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • LT: less than • GT: greater than • LTOE: less than or equal to • GTOE: greater than or equal to <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
evaluation_periods	Yes	Integer	<p>Explanation Number of consecutive five-minute periods, during which a metric threshold is reached.</p> <p>Constraints N/A</p> <p>Value range 1-288</p> <p>Default value N/A</p>

Table 6-104 Tag

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • A tag key can contain letters, digits, spaces, and special characters <code>._:=-@</code>, but cannot start or end with a space or start with <code>_sys_</code>. • The tag key of a resource must be unique. • It can contain a maximum of 128 Unicode characters and cannot be an empty string. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
value	Yes	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> The value can contain letters, digits, spaces, and special characters <code>_:=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>. The value can contain a maximum of 255 Unicode characters and can be an empty string. <p>Default value N/A</p>

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.

URI

GET /v2/{project_id}/clusters/{cluster_id}/files

Table 6-105 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Table 6-106 Query parameters

Parameter	Mandatory	Type	Description
path	Yes	String	<p>Explanation File category. For example, to access the /tmp/test directory, the value must be a directory. The URI is /v2/{project_id}/clusters/{cluster_id}/files?path=%2Ftmp%2Ftest.</p> <p>Constraints N/A</p> <p>Value range A single-level directory must comply with the following rules:</p> <ul style="list-style-type: none"> • The directory path cannot be left blank. • The value cannot start or end with a period (.). • The value cannot contain the following characters: <code>/:*?"<> \\;&,'!{}[]\$%+</code> • The value cannot exceed 255 characters. <p>Default value N/A</p>
offset	No	String	<p>Explanation Pagination parameter. The file list is queried from the offset.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value 1</p>

Parameter	Mandatory	Type	Description
limit	No	String	<p>Explanation Maximum number of records displayed on a page.</p> <p>Constraints N/A</p> <p>Value range ≤1000</p> <p>Default value 100</p>
sort_key	No	String	<p>Explanation Attribute used for sorting the list.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • path_suffix: file or directory name • length: file size • modification_time: modification time <p>Default value path_suffix</p>
order	No	String	<p>Explanation List sorting method.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • desc: descending order • asc: ascending order <p>Default value desc</p>

Request Parameters

None

Response Parameters

Table 6-107 Response parameters

Parameter	Type	Description
total_count	Integer	<p>Explanation The total number of files, which is irrelevant to pagination.</p> <p>Value range N/A</p>
files	Array of FileStatusV2 objects	<p>Explanation File list. For details, see Table 6-108.</p>

Table 6-108 FileStatus parameters

Parameter	Type	Description
path_suffix	String	<p>Explanation 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>.</p> <p>Value range N/A</p>
owner	String	<p>Explanation File owner.</p> <p>Value range N/A</p>
group	String	<p>Explanation File owner group.</p> <p>Value range N/A</p>
permission	String	<p>Explanation Permission information.</p> <p>Value range N/A</p>

Parameter	Type	Description
replication	Integer	Explanation Number of replicas. Value range N/A
block_size	Integer	Explanation Block size. Value range N/A
length	Integer	Explanation File length. Value range N/A
type	String	Explanation File type Value range <ul style="list-style-type: none"> • FILE: file • DIRECTORY: directory
children_num	Integer	Explanation Number of files in the directory. Value range N/A
access_time	Long	Explanation File access time. Value range N/A
modification_time	Long	Explanation File modification time Value range N/A

Example Request

Obtain the list of files from a specified directory.

```
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.

URI

- Format
POST /v2/{project_id}/clusters/{cluster_id}/sql-execution
- Parameter description

Table 6-109 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-110 Request parameters

Parameter	Mandatory	Type	Description
sql_content	Yes	String	<p>Explanation SQL statement to be executed. Currently, only a single SQL statement can be executed at a time, and the statement cannot contain a semicolon (;).</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
database	No	String	<p>Explanation Database where the SQL statement is executed on.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value default</p>
archive_path	No	String	<p>Explanation Directory for storing the dumped SQL execution results. Only the select statement dumps query results. Currently, the query results can be dumped only to OBS.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Response Parameters

Status code: 200

Table 6-111 Response parameters

Parameter	Type	Description
id	String	<p>Explanation SQL execution ID. The ID is generated only when the SELECT, SHOW, or DESC statement is executed. For other operations, the ID is empty.</p> <p>Value range N/A</p>
message	String	<p>Explanation Error message</p> <p>Value range N/A</p>
statement	String	<p>Explanation Ongoing SQL statement</p> <p>Value range N/A</p>
status	String	<p>Explanation SQL execution status</p> <p>Value range</p> <ul style="list-style-type: none"> • QUEUED: queuing • WAITING_FOR_RESOURCES: waiting for resources • PLANNING: being planning • STARTING: being started • RUNNING: running • FINISHING: to be completed • FINISHED: completed • FAILED: execution failed
result_location	String	<p>Explanation Path for archiving the final results of the SQL query statement Only the SELECT statement dumps the SQL execution results to result_location.</p> <p>Value range N/A</p>

Parameter	Type	Description
content	Array<Array<String>>	Explanation SQL execution result. Only non-SELECT statements return results in content . If the SQL statement does not return results, content is empty.

Status code: 400

Table 6-112 Response body parameters

Parameter	Type	Description
error_code	String	Explanation Error code Value range N/A
error_msg	String	Explanation Error message. Value range N/A

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.

URI

- Format
GET /v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}
- Parameter description

Table 6-113 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details on how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
sql_id	Yes	String	<p>Explanation The execution ID of a SQL statement, that is, <code>sql_id</code> in the return result in Submitting a SQL Statement.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-114 Response parameters

Parameter	Type	Description
id	String	<p>Explanation SQL execution ID. The ID is generated only when the SELECT, SHOW, or DESC statement is executed. For other operations, the ID is empty.</p> <p>Value range N/A</p>
message	String	<p>Explanation Error message</p> <p>Value range N/A</p>
statement	String	<p>Explanation Ongoing SQL statement</p> <p>Value range N/A</p>
status	String	<p>Explanation SQL execution status</p> <p>Value range</p> <ul style="list-style-type: none"> • QUEUED: queuing • WAITING_FOR_RESOURCES: waiting for resources • PLANNING: being planning • STARTING: being started • RUNNING: running • FINISHING: to be completed • FINISHED: completed • FAILED: execution failed
result_location	String	<p>Explanation Path for archiving the final results of the SQL query statement Only the SELECT statement dumps the SQL execution results to result_location.</p> <p>Value range N/A</p>

Parameter	Type	Description
content	Array<Array<String>>	Explanation SQL execution result Only non-SELECT statements return results in content . If the SQL statement does not return results, content is empty.

Status code: 400

Table 6-115 Response body parameters

Parameter	Type	Description
error_code	String	Explanation Error code Value range N/A
error_msg	String	Explanation Error message. Value range N/A

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_xxxx",
  "statement": "show tables",
  "status": "FINISHED",
  "result_location": "obs://my_bucket/uuid_date/xxx.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.

URI

POST /v2/{project_id}/clusters/{cluster_id}/sql-execution/{sql_id}/cancel

Table 6-116 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details on how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
sql_id	Yes	String	<p>Explanation The execution ID of a SQL statement, that is, <code>sql_id</code> in the return result in Submitting a SQL Statement.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-117 Response parameters

Parameter	Type	Description
message	String	<p>Explanation Error message</p> <p>Value range N/A</p>
status	String	<p>Explanation The result of canceling the execution of the SQL statement. By default, SUCCEED is returned. SUCCEED is also returned when the task is completed. Only when the running SQL statement fails to be canceled, FAILED is returned.</p> <p>Value range</p> <ul style="list-style-type: none"> • SUCCEED • FAILED

Status code: 400

Table 6-118 Response body parameter

Parameter	Type	Description
message	String	<p>Explanation Error message</p> <p>Value range N/A</p>
status	String	<p>Explanation The result of canceling the execution of the SQL statement. By default, SUCCEED is returned. SUCCEED is also returned when the task is completed. Only when the running SQL statement fails to be canceled, FAILED is returned.</p> <p>Value range</p> <ul style="list-style-type: none"> • SUCCEED • FAILED

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.

URI

- URI format
GET /v2/{project_id}/clusters/{cluster_id}/agency-mapping
- Parameter description

Table 6-119 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation The cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-120 Response parameter

Parameter	Type	Description
agency_mappings	Array of AgencyMapping objects	<p>Explanation The mapping between users or user groups and agencies. For details, see Table 6-121.</p> <p>Constraints N/A</p>

Table 6-121 agency_mappings parameters

Parameter	Type	Description
agency	String	<p>Explanation The name of the IAM agency bound to the mapping.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
identifier_type	String	<p>Explanation The agency type. Available values are User and Group.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • User: indicates that the mapping is for users. Enter the user name list in identifiers. • Group: indicates that the mapping is for user groups. Enter the user group name list in identifiers. <p>Default value N/A</p>

Parameter	Type	Description
identifiers	Array of String	<p>Explanation List of users or user groups mapped to the IAM agency. Log in to the IAM management console and choose Users or User Groups in the navigation pane to obtain the user or user group name list.</p> <p>Constraints N/A</p>
agency_id	String	<p>Explanation Unique ID of the agency bound to the mapping. Log in to the IAM management console and choose Agencies in the left navigation pane. On the Agencies page that is displayed, move your cursor over the agency name to obtain the agency ID.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Status code: 400

Table 6-122 Response parameters

Parameter	Type	Description
error_code	String	<p>Explanation Error code.</p> <p>Value range N/A</p>
error_msg	String	<p>Explanation Error message</p> <p>Value range N/A</p>

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.

URI

- URI format
PUT /v2/{project_id}/clusters/{cluster_id}/agency-mapping
- Parameter description

Table 6-123 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation The cluster ID. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-124 Request parameter

Parameter	Mandatory	Type	Description
agency_mappings	Yes	Array of AgencyMapping objects	<p>Explanation The mapping between users or user groups and agencies. For details, see Table 6-125.</p> <p>Constraints N/A</p>

Table 6-125 agency_mappings parameters

Parameter	Mandatory	Type	Description
agency	Yes	String	<p>Explanation Name of the IAM agency bound to the mapping.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
identifier_type	Yes	String	<p>Explanation The agency type. Available values are User and Group.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • User: indicates that the mapping is for users. Enter the user name list in identifiers. • Group: indicates that the mapping is for user groups. Enter the user group name list in identifiers. <p>Default value N/A</p>
identifiers	Yes	Array of String	<p>Explanation List of users or user groups mapped to the IAM agency. Log in to the IAM management console and choose Users or User Groups in the navigation pane to obtain the user or user group name list.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
agency_id	Yes	String	<p>Explanation Unique ID of the agency bound to the mapping. Log in to the IAM management console and choose Agencies in the left navigation pane. On the Agencies page that is displayed, move your cursor over the agency name to obtain the agency ID.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Response Parameters

Status code: 200

Table 6-126 Response body parameter

Parameter	Type	Description
result	String	<p>Explanation Result of the request for updating a mapping.</p> <p>Value range</p> <ul style="list-style-type: none"> • succeeded: The operation is successful. • failed: The operation failed.

Status code: 400

Table 6-127 Response body parameters

Parameter	Type	Description
error_code	String	Explanation Error code. Value range N/A
error_msg	String	Explanation Error message Value range N/A

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.

URI

POST /v2/{project_id}/data-connectors

Table 6-128 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Request Parameters

Table 6-129 Request body parameters

Parameter	Mandatory	Type	Description
data_connect or or	Yes	object	<p>Explanation Data connection. For details, see Table 6-130.</p>

Table 6-130 DataConnector

Parameter	Mandatory	Type	Description
connector_name	Yes	String	<p>Explanation Data connection name.</p> <p>Value range N/A</p>

Parameter	Mandatory	Type	Description
source_type	Yes	String	<p>Explanation Data connection type.</p> <p>Value range</p> <ul style="list-style-type: none"> • RDS_POSTGRES: RDS PostgreSQL database • RDS_MYSQL: RDS MySQL database • gaussdb-mysql: GaussDB(for MySQL)
source_info	Yes	String	<p>Explanation Data source information, which is in JSON format.</p> <p>Value range N/A</p>

Response Parameters

Status code: 201

Table 6-131 Response body parameter

Parameter	Type	Description
connector_id	String	<p>Explanation Data connection ID returned after a data connection is created.</p> <p>Value range N/A</p>

Example Request

Create a MySQL 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.

URI

GET /v2/{project_id}/data-connectors

Table 6-132 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Table 6-133 Query parameters

Parameter	Mandatory	Type	Description
connector_id	No	String	<p>Explanation Connection ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
source_type	No	String	<p>Explanation Data source type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • RDS_POSTGRES: RDS PostgreSQL database • RDS_MYSQL: RDS MySQL database • gaussdb-mysql: GaussDB(for MySQL) <p>Default value N/A</p>
connector_name	No	String	<p>Explanation Data connection name.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
limit	No	Integer	<p>Explanation Number of resources on each page.</p> <p>Constraints N/A</p> <p>Value range 1-1000</p> <p>Default value N/A</p>
offset	No	Integer	<p>Explanation Start offset of the pagination query.</p> <p>Constraints N/A</p> <p>Value range ≥0</p> <p>Default value N/A</p>
available	No	Boolean	<p>Explanation Whether the data connection is valid.</p> <p>Constraints N/A</p> <p>Value range true: The connection is valid. false: The connection is invalid.</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-134 Response body parameters

Parameter	Type	Description
total_count	Integer	<p>Explanation Total number of data connections</p> <p>Value range None</p>
data_connectors	Array of DataConnectorDetail objects	<p>Explanation Data connection details list. For details, see Table 6-135.</p> <p>Constraints N/A</p>

Table 6-135 DataConnectorDetail

Parameter	Type	Description
connector_name	String	<p>Explanation Data connection name.</p> <p>Value range N/A</p>
source_type	String	<p>Explanation Data connection type.</p> <p>Value range</p> <ul style="list-style-type: none"> • RDS_POSTGRES: RDS PostgreSQL database • RDS_MYSQL: RDS MySQL database • gaussdb-mysql: GaussDB(for MySQL)
source_info	String	<p>Explanation Data source information, which is in JSON format.</p> <p>Value range N/A</p>
connector_id	String	<p>Explanation Data connection ID.</p> <p>Value range N/A</p>
create_time	Long	<p>Explanation Time when the connection is created.</p> <p>Value range N/A</p>

Parameter	Type	Description
last_update_time	Long	Explanation Last update time of the connection. Value range N/A
create_by	String	Explanation Creator ID. Value range N/A
create_user	String	Explanation Creator username. Value range N/A
tenant_id	String	Explanation Tenant ID. Value range N/A
last_update_by	String	Explanation ID of the user who last updated the information. Value range N/A
status	Integer	Explanation Data connection status. Value range <ul style="list-style-type: none"> ● -1: deleted ● 0: normal ● -2: abnormal ● 1: in use
used_clusters	String	Explanation Clusters using the connection. Value range N/A
encrypt_type	Integer	Explanation Encryption type Value range None

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": "xxxe662a7e24cc99bfc858c4558dbf6",
    "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.

URI

PUT /v2/{project_id}/data-connectors/{connector_id}

Table 6-136 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
connector_id	Yes	String	<p>Explanation Data connection ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-137 Request body parameters

Parameter	Mandatory	Type	Description
data_connector	Yes	DataConnector object	<p>Explanation Data connection. For details, see Table 6-138.</p>

Table 6-138 DataConnector

Parameter	Mandatory	Type	Description
connector_name	Yes	String	Explanation Data connection name. Value range N/A
source_type	Yes	String	Explanation Data connection type. Value range <ul style="list-style-type: none"> • RDS_POSTGRES: RDS PostgreSQL database • RDS_MYSQL: RDS MySQL database • gaussdb-mysql: GaussDB(for MySQL)
source_info	Yes	String	Explanation Data source information, which is in JSON format. Value range N/A

Response Parameters

Status code: 200

Table 6-139 Response body parameter

Parameter	Type	Description
result	String	Explanation Result of the request for updating a mapping. Value range <ul style="list-style-type: none"> • succeeded: The operation is successful. • failed: The operation failed.

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
```

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

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.

URI

DELETE /v2/{project_id}/data-connectors/{connector_id}

Table 6-140 URI parameters

Parameter	Mandatory	Type	Description
connector_id	Yes	String	<p>Explanation Connection ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

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.

URI

GET /v2/{project_id}/metadata/versions

Table 6-141 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Request Parameter

None

Response Parameter

Status code: 200

Table 6-142 Response body parameter

Parameter	Type	Description
cluster_versions	Array of strings	<p>Explanation Cluster version list</p>

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.

URI

GET /v2/{project_id}/metadata/version/{version_name}/available-flavor

Table 6-143 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
version_name	Yes	String	<p>Explanation Cluster version, for example, MRS 3.3.1. If the request client does not support automatic escape, escape the space to %20, for example, MRS %203.3.1.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters.</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-144 Response body parameter

Parameter	Type	Description
version_name	String	<p>Explanation Version name.</p> <p>Value range N/A</p>
available_flavors	Array of AzFlavors objects	<p>Explanation Flavors supported by different AZs. For details, see Table 6-145.</p>

Table 6-145 AzFlavors

Parameter	Type	Description
az_code	String	Explanation AZ code. Value range N/A
az_name	String	Explanation AZ name Value range N/A
master	Array of Flavor objects	Explanation Specifications supported by the master node. For details, see Table 6-146 .
core	Array of Flavor objects	Explanation Specifications supported by the core node. For details, see Table 6-146 .
task	Array of Flavor objects	Explanation Specifications supported by the task node. For details, see Table 6-146 .

Table 6-146 Flavor

Parameter	Type	Description
flavor_name	String	Explanation Flavor name. Value range N/A

Example Request

Obtain details about MRS 3.3.1-LTS metadata.

```
/v2/f77c10d14a544393a24e5f0bf53202b6/metadata/version/MRS%203.3.1-LTS/available-flavor
```

Example Response

Status code: 200

Available specifications of the current version

```
{
  "version_name" : "MRS 3.3.1-LTS",
  "available_flavors" : [ {
```

```

"az_code" : "az1.cn-fcs-xxx",
"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" : "az1.cn-fcs-xxx",
"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" : "az1.cn-fcs-xxx",
"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.

URI

GET /v2/{project_id}/clusters/{cluster_id}/iam-sync-user

Table 6-147 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-148 Response body parameter

Parameter	Type	Description
user_names	Array of strings	Explanation Synchronized users. Constraints N/A
group_names	Array of strings	Explanation Synchronized user groups. Constraints N/A

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.

URI

POST /v2/{project_id}/clusters/{cluster_id}/iam-sync-user

Table 6-149 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-150 Request body parameters

Parameter	Mandatory	Type	Description
is_all_sync	No	Boolean	<p>Explanation Whether to perform full synchronization.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: full synchronization • false: synchronization of specified users and user groups <p>Default value false</p>
group_names	No	Array of strings	<p>Explanation IAM user groups to be synchronized</p> <p>Constraints N/A</p>
user_names	No	Array of strings	<p>Explanation IAM user to be synchronized</p> <p>Constraints N/A</p>

Response Parameters

Status code: 202

Table 6-151 Response body parameter

Parameter	Type	Description
state	String	<p>Explanation Request result</p> <p>Value range N/A</p>

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.

URI

DELETE /v2/{project_id}/clusters/{cluster_id}/iam-sync-user

Table 6-152 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-153 Request body parameters

Parameter	Mandatory	Type	Description
group_names	No	Array of strings	<p>Explanation IAM user group whose synchronization is to be canceled</p> <p>Constraints The number of records cannot exceed 1,000.</p>
user_names	No	Array of strings	<p>Explanation IAM user whose synchronization is to be canceled.</p> <p>Constraints The number of records cannot exceed 1,000.</p>

Response Parameters

Status code: 202

Table 6-154 Response body parameter

Parameter	Type	Description
state	String	<p>Explanation Request result</p> <p>Value range N/A</p>

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.

URI

POST /v2/{project_id}/clusters/{cluster_id}/tags/switch

Table 6-155 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 6-156 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	<p>Explanation Operation, which can be create or delete.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> create: Create a cluster tag. delete: Delete a cluster tag. <p>Default value N/A</p>

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.

URI

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

Table 6-157 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-158 Response body parameter

Parameter	Type	Description
status	String	<p>Explanation Tag status.</p> <p>Value range</p> <ul style="list-style-type: none"> • processing: The operation is being processed. • succeed: The operation is successful. • failed: The operation failed.
default_tags_enable	Boolean	<p>Explanation Whether to enable default cluster tags.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable default cluster tags. • false: Disable default cluster tags.

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.

URI

GET /v2/{project_id}/clusters/{cluster_id}/tags/quota

Table 6-159 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Table 6-160 Query parameters

Parameter	Mandatory	Type	Description
auto_scaling_policy_tags	No	Boolean	<p>Explanation Whether to query auto scaling policy tags.</p> <p>Constraints N/A</p> <p>Value range true: Query the tags of the auto scaling policy. false: Do not query the tags of the auto scaling policy.</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 6-161 Response body parameter

Parameter	Type	Description
total_quota	Integer	<p>Explanation Total quotas.</p> <p>Value range N/A</p>
available_quota	Integer	<p>Explanation Available quotas.</p> <p>Value range N/A</p>

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.

A maximum of 10 clusters can be concurrently created.

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 VPC > Querying VPCs and VPC > Creating a VPC in the <i>VPC API Reference</i> .
Subnet	See operation instructions in Subnet > Querying Subnets and Subnet > Creating a Subnet in the <i>VPC API Reference</i> .
Key Pair	See operation instructions in ECS SSH Key Management > Querying SSH Key Pairs and ECS SSH Key Management > Creating and Importing an SSH Key Pair in the <i>ECS API Reference</i> .
Zone	See Endpoints for details about regions and AZs.
Version	Currently, MRS 3.3.1-LTS is supported.

Resource	How to Obtain
Component	<p>MRS 3.3.1-LTS supports the following components:</p> <ul style="list-style-type: none"> • An analysis cluster contains the following components: Hadoop, Spark, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, JobGateway, Guardian, and Doris. • A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger. • A hybrid cluster contains the following components: Hadoop, Spark, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, Flume, JobGateway, Guardian, and Doris. • A custom cluster contains the following components: CDL, Hadoop, Spark, HBase, Hive, Hue, IoTDB, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse, Guardian, JobGateway, Doris and MemArtsCC.

URI

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

Table 7-2 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Request Parameters

Table 7-3 Request parameters

Parameter	Mandatory	Type	Description
cluster_version	Yes	String	<p>Explanation Cluster version.</p> <p>Constraints N/A</p> <p>Value range MRS 3.3.1-LTS</p> <p>Default value N/A</p>
cluster_name	Yes	String	<p>Explanation Cluster name. It must be unique.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
master_node_num	No	Integer	<p>Explanation Number of Master nodes.</p> <p>Constraints If cluster HA is enabled, set this parameter to 2. If cluster HA is disabled, set this parameter to 1. This parameter cannot be set to 1 in MRS 3.x.</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
core_node_num	No	Integer	<p>Explanation Number of Core nodes. The default maximum number of core nodes is 500. If more than 500 core nodes are required, apply for a higher quota.</p> <p>Constraints N/A</p> <p>Value range 1-500</p> <p>Default value N/A</p>
data_center	Yes	String	<p>Explanation The information about the region where the cluster is located. For details, see Endpoints.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
vpc	Yes	String	<p>Explanation</p> <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"> 1. Log in to the management console. 2. Choose Virtual Private Cloud > My VPCs. <p>On the Virtual Private Cloud page, obtain the VPC name from the list.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
master_node_size	No	String	<p>Explanation</p> <p>The instance specifications of master nodes, for example, c3.4xlarge.2.linux.bigdata. The host specifications supported by MRS are determined by CPU, memory, and disk space. 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.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
core_node_size	No	String	<p>Explanation The instance specifications of core nodes, for example, c3.4xlarge.2.linux.bigdata. 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.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
component_list	Yes	Array of component_list objects	<p>Explanation The list of service components to be installed. For details about the parameters, see Table 7-4.</p> <p>Constraints N/A</p>
available_zone_id	Yes	String	<p>Explanation AZ ID.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
vpc_id	Yes	String	<p>Explanation</p> <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"> 1. Log in to the management console. 2. Choose Virtual Private Cloud > My VPCs. <p>On the Virtual Private Cloud page, obtain the VPC ID from the list.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

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

Parameter	Mandatory	Type	Description
subnet_name	Yes	String	<p>Explanation</p> <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"> 1. Log in to the management console. 2. Choose Virtual Private Cloud > My VPCs. 3. Locate the row that contains the target VPC and click the number in the Subnets column to obtain the subnet name. <p>Constraints</p> <p>At least one of subnet_id and subnet_name must be configured. If the two parameters are configured but do not match the same subnet, the cluster fails to create. If only subnet_name 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. subnet_id is recommended.</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
security_groups_id	No	String	<p>Explanation</p> <p>The ID of the security group configured for the cluster.</p> <ul style="list-style-type: none"> • If this parameter is left blank, MRS automatically creates a security group, whose name starts with mrs_{cluster_name}. • 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. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
add_jobs	No	Array of add_jobs objects	<p>Explanation</p> <p>Jobs can be submitted when a cluster is created. Currently, only one job can be created. For details about the parameters, see Table 7-5.</p> <p>Constraints</p> <p>There must be no more than 1 record.</p>

Parameter	Mandatory	Type	Description
volume_size	No	Integer	<p>Explanation</p> <p>Data disk storage space of Master and Core nodes, in GB 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"> • 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. • 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. <p>Constraints</p> <p>This parameter is not recommended. For details, see the description of the volume_type parameter.</p> <p>Value range</p> <p>100-32000</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
volume_type	No	String	<p>Explanation</p> <p>The data disk storage type of master and core nodes. Currently, SATA, SAS, SSD, and GPSSD are supported. Disk parameters can be represented by volume_type and volume_size, or multi-disk parameters. If the volume_type and volume_size parameters coexist with the multi-disk parameters, the system reads the volume_type and volume_size parameters first. You are advised to use the multi-disk parameters.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O • GPSSD: general-purpose SSD <p>Default value</p> <p>N/A</p>
master_data_volume_type	No	String	<p>Explanation</p> <p>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.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O • GPSSD: general-purpose SSD <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
master_data_volume_size	No	Integer	<p>Explanation</p> <p>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. You only need to pass in a number without the unit GB.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>100-32000</p> <p>Default value</p> <p>N/A</p>
master_data_volume_count	No	Integer	<p>Explanation</p> <p>This parameter is a multi-disk parameter, indicating the number of data disks of the master nodes.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The value can only be 1.</p> <p>Default value</p> <p>1</p>
core_data_volume_type	No	String	<p>Explanation</p> <p>This parameter is a multi-disk parameter, indicating the data disk storage type of core nodes. Currently, SATA, SAS, SSD, and GPSSD are supported.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <ul style="list-style-type: none"> ● SATA: common I/O ● SAS: high I/O ● SSD: ultra-high I/O ● GPSSD: general-purpose SSD <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
core_data_volume_size	No	Integer	<p>Explanation</p> <p>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. You only need to pass in a number without the unit GB.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>100-32000</p> <p>Default value</p> <p>N/A</p>
core_data_volume_count	No	Integer	<p>Explanation</p> <p>This parameter is a multi-disk parameter, indicating the number of data disks of the core nodes.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>1-20</p> <p>Default value</p> <p>N/A</p>
task_node_groups	No	Array of task_node_groups objects	<p>Explanation</p> <p>The list of task nodes. For details about the parameters, see Table 7-6.</p> <p>Constraints</p> <p>There must be no more than 1 record.</p>
bootstrap_scripts	No	Array of Bootstrap Script objects	<p>Explanation</p> <p>The Bootstrap action script information. For details about the parameters, see Table 7-8.</p> <p>Constraints</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
node_public_cert_name	No	String	<p>Explanation The name of a key pair. You can use a key pair to log in to a cluster node.</p> <p>Constraints If login_mode is set to 1, the request body contains the node_public_cert_name field.</p> <p>Value range N/A</p> <p>Default value N/A</p>
cluster_admin_secret	No	String	<p>Explanation Password of the MRS Manager administrator.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • Must contain 8 to 26 characters. • Cannot be the username or the username spelled backwards. • Must contain every type of the following: <ul style="list-style-type: none"> - Lowercase letters - Uppercase letters - Numbers - Special characters (!@#\$%^- _+[{ } : ; , . / ?) <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_master_secret	No	String	<p>Explanation The password of user root for logging in to a cluster node.</p> <p>Constraints If login_mode is set to 0, the request body contains the cluster_master_secret field.</p> <p>Value range A password must meet the following complexity requirements:</p> <ul style="list-style-type: none"> • Must be 8 to 26 characters long. • Must contain every type of the following: uppercase letters, lowercase letters, numbers, and special characters (!@\$%^_-=+ [{}]:./?), but must not contain spaces. • Cannot be the username or the username spelled backwards. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
safe_mode	Yes	Integer	<p>Explanation The running mode of an MRS cluster.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: normal cluster. In a normal cluster, Kerberos authentication is disabled, and users can use all functions provided by the cluster. • 1: 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. <p>Default value N/A</p>
tags	No	Array of tag objects	<p>Explanation The cluster tags. For details about the parameters, see Table 7-9.</p> <p>Constraints A maximum of 20 tags can be used in a cluster. The tag name (key) must be unique. The tag key and value can contain letters, digits, spaces, and special characters (_ := + - @), but cannot start or end with a space or start with _sys_.</p>

Parameter	Mandatory	Type	Description
cluster_type	No	Integer	<p>Explanation The cluster type. Currently, hybrid clusters cannot be created using APIs.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: analysis cluster • 1: streaming cluster <p>Default value 0</p>
log_collection	No	Integer	<p>Explanation Whether to collect logs when cluster creation fails. The default value is 1, indicating that OBS buckets are created only for collecting logs when an MRS cluster fails to create.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: Do not collect logs. • 1: Collect logs. <p>Default value 1</p>
enterprise_project_id	No	String	<p>Explanation Enterprise project ID When you create a cluster, associate the enterprise project ID with the cluster. The default value is 0, indicating the default enterprise project.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value 0</p>

Parameter	Mandatory	Type	Description
login_mode	No	Integer	<p>Explanation Cluster login mode.</p> <p>Constraints</p> <ul style="list-style-type: none"> • If login_mode is set to 0, the request body contains the cluster_master_secret field. • If login_mode is set to 1, the request body contains the node_public_cert_name field. <p>Value range</p> <ul style="list-style-type: none"> • 0: password • 1: key pair <p>Default value 1</p>
node_groups	No	Array of NodeGroupV11 objects	<p>Explanation List of nodes. For details about the parameters, see Table 7-10.</p> <p>Constraints Configure either this parameter or the following parameters: master_node_num, master_node_size, core_node_num, core_node_size, master_data_volume_type, master_data_volume_size, master_data_volume_count, core_data_volume_type, core_data_volume_size, core_data_volume_count, volume_type, volume_size, task_node_groups</p>

Table 7-4 ComponentAmbV11

Parameter	Mandatory	Type	Description
component_name	Yes	String	<p>Explanation Component name. For details, see the component information in Table 7-1.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Table 7-5 AddJobsReqV11

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

Parameter	Mandatory	Type	Description
jar_path	No	String	<p>Explanation Path of the .jar file or .sql file to be executed.</p> <p>Constraints N/A</p> <p>Value range The value must meet the following requirements:</p> <ul style="list-style-type: none"> • The value contains a maximum of 1,023 characters. It cannot contain special characters (; &>,<'\$) and cannot be left blank or all spaces. • Files can be stored in HDFS or OBS. The path varies depending on the file system. <ul style="list-style-type: none"> – OBS: The path must start with s3a://. Files or programs encrypted by KMS are not supported. – HDFS: The path starts with a slash (/). • Spark Script must end with .sql while MapReduce and Spark Jar must end with .jar.sql and jar are case-insensitive. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
arguments	No	String	<p>Explanation 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>Constraints N/A</p> <p>Value range The parameter can contain 0 to 150,000 characters, but special characters (; &>'<\$) are not allowed.</p> <p>Default value N/A</p>
input	No	String	<p>Explanation The data input path. Files can be stored in HDFS or OBS. The path varies depending on the file system.</p> <ul style="list-style-type: none"> • OBS: The path must start with s3a://. Files or programs encrypted by KMS are not supported. • HDFS: The path starts with a slash (/). <p>Constraints N/A</p> <p>Value range The value can contain 0 to 1,023 characters, but special characters (; &>'<\$) are not allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
output	No	String	<p>Explanation</p> <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"> • OBS: The path must start with s3a://. • HDFS: The path starts with a slash (/). <p>If the specified path does not exist, the system will automatically create it.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The value can contain 0 to 1,023 characters, but special characters (; &>'<\$) are not allowed.</p> <p>Default value</p> <p>N/A</p>
job_log	No	String	<p>Explanation</p> <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"> • OBS: The path must start with s3a://. • HDFS: The path starts with a slash (/). <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The value can contain 0 to 1,023 characters, but special characters (; &>'<\$) are not allowed.</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
shutdown_cluster	No	Boolean	<p>Explanation Whether to delete the cluster after the job execution is complete.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Delete the cluster after job execution is complete. • false: Do not delete the cluster after the job execution is complete. <p>Default value N/A</p>
file_action	No	String	<p>Explanation The action to be performed on a file.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • import: Import data. • export: Export data. <p>Default value N/A</p>
submit_job_once_cluster_run	Yes	Boolean	<p>Explanation Whether to submit a job when creating a cluster. Set it to true.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Submit a job during cluster creation. • false: Submit a job after the cluster is created. <p>Default value N/A</p>

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

Table 7-6 TaskNodeGroup

Parameter	Mandatory	Type	Description
node_num	Yes	Integer	<p>Explanation Number of Task nodes.</p> <p>Constraints The total number of Core and Task nodes cannot exceed 500.</p> <p>Value range 0-500</p> <p>Default value N/A</p>
node_size	Yes	String	<p>Explanation The instance specifications of task nodes, for example, c3.4xlarge.2.linux.bigdatac6.4xlarge.4.linux.mrs.</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> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
data_volume_type	Yes	String	<p>Explanation Data disk storage type of the Task node. Supported types include SATA, SAS, and SSD.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> ● SATA: common I/O ● SAS: high I/O ● SSD: ultra-high I/O ● GPSSD: general-purpose SSD <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
data_volume_count	Yes	Integer	<p>Explanation Number of data disks of a Task node.</p> <p>Constraints N/A</p> <p>Value range 0-20</p> <p>Default value N/A</p>
data_volume_size	Yes	Integer	<p>Explanation Data disk storage space of a Task node. You only need to pass in a number without the unit GB.</p> <p>Constraints N/A</p> <p>Value range 100-32000</p> <p>Default value N/A</p>
auto_scaling_policy	No	auto_scaling_policy object	<p>Explanation The auto scaling policy. For details, see Table 7-7.</p> <p>Constraints N/A</p>

Table 7-7 AutoScalingPolicy

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	<p>Explanation Whether to enable the auto scaling policy.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable the auto scaling rule. • false: Disable the autoscaling rule. <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation The minimum number of nodes reserved in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation The maximum number of nodes in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
resources_plans	No	Array of resources_plan objects	<p>Explanation</p> <p>The resource plan list. For details, see Table 7-11. If this parameter is left blank, the resource plan is disabled.</p> <p>Constraints</p> <p>When auto_scaling_enable is set to true, either this parameter or rules must be configured. There must be no more than 5 records.</p>
exec_scripts	No	Array of scale_script objects	<p>Explanation</p> <p>The list of custom scaling automation scripts. For details, see Table 7-14. If this parameter is left blank, a hook script is disabled.</p> <p>Constraints</p> <p>The number of records cannot exceed 10.</p>
rules	No	Array of rules objects	<p>Explanation</p> <p>The list of auto scaling rules. For details, see Table 7-12.</p> <p>Constraints</p> <p>When auto_scaling_enable is set to true, either this parameter or resources_plans must be configured. The number of records cannot exceed 10.</p>

Table 7-8 BootstrapScript

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of a bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range The names of bootstrap action scripts in the same cluster must be unique. The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-), and cannot start with a space.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
uri	Yes	String	<p>Explanation</p> <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"> • OBS bucket path: Enter a script path. For example, enter the path of the public sample script provided by MRS. Example: s3a://bootstrap/presto/presto-install.sh. If dualroles is installed, the parameter of the presto-install.sh script is dualroles. If worker is installed, the parameter of the presto-install.sh script is worker. Based on the Presto usage habit, you are advised to install dualroles on the active master nodes and worker on the Core nodes. • Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
parameters	No	String	<p>Explanation</p> <p>The bootstrap action script parameters.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
nodes	Yes	Array of strings	<p>Explanation The type of a node where the bootstrap action script is executed. The value can be Master, Core, or Task.</p> <p>Constraints The node type must be represented in lowercase letters.</p>
active_master	No	Boolean	<p>Explanation Whether the bootstrap action script runs only on active master nodes.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script runs only on active Master nodes. • false: The bootstrap action script can run on all Master nodes. <p>Default value false</p>
before_component_start	No	Boolean	<p>Explanation Time when the bootstrap action script is executed. Currently, the following two options are available: Before component start and After component start</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script is executed before the component starts. • false: The bootstrap action script is executed after the component starts. <p>Default value false</p>

Parameter	Mandatory	Type	Description
fail_action	Yes	String	<p>Explanation Whether to continue executing subsequent scripts and creating a cluster after the Bootstrap action script fails to be executed. You are advised to set this parameter to continue in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value errorout</p>
start_time	No	Long	<p>Explanation The execution time of one bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
state	No	String	<p>Explanation The running status of one bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> ● PENDING: The script is suspended. ● IN_PROGRESS: The script is being processed. ● SUCCESS ● FAILURE: The script fails to be executed. <p>Default value N/A</p>
action_stages	No	Array of strings	<p>Explanation Select the time when the bootstrap action script is executed.</p> <ul style="list-style-type: none"> ● BEFORE_COMPONENT_FIRST_START: before initial component starts ● AFTER_COMPONENT_FIRST_START: after initial component starts ● BEFORE_SCALE_IN: before scale-in ● AFTER_SCALE_IN: after scale-in ● BEFORE_SCALE_OUT: before scale-out ● AFTER_SCALE_OUT: after scale-out <p>Constraints N/A</p>

Table 7-9 Tag

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain a maximum of 128 characters and cannot be an empty string. • The tag key of a resource must be unique. • A tag key can contain letters, digits, spaces, and special characters <code>._:=-@</code>, but cannot start or end with a space or start with <code>_sys_</code>. <p>Default value N/A</p>
value	Yes	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain a maximum of 255 characters and can be an empty string. • The value can contain letters, digits, spaces, and special characters <code>._:=-@</code>, but cannot start or end with a space or start with <code>_sys_</code>. <p>Default value N/A</p>

Table 7-10 NodeGroupV11

Parameter	Mandatory	Type	Description
group_name	Yes	String	<p>Explanation The node group name.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • master_node_default_group • core_node_analysis_group • core_node_streaming_group • task_node_analysis_group • task_node_streaming_group <p>Default value N/A</p>
node_num	Yes	Integer	<p>Explanation Number of nodes.</p> <p>Constraints The total number of Core and Task nodes cannot exceed 500.</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
node_size	Yes	String	<p>Explanation The instance specifications of cluster nodes, for example, c3.4xlarge.2.linux.bigdata. The host specifications supported by MRS are determined by CPU, memory, and disk space. 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.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
root_volume_size	No	String	<p>Explanation The system disk storage space of a node.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
root_volume_type	No	String	<p>Explanation System disk storage type of a node. Supported types include SATA, SAS, and SSD.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O • GPSSD: general-purpose SSD <p>Default value N/A</p>
data_volume_type	No	String	<p>Explanation Data disk storage type of nodes. Supported types include SATA, SAS, and SSD.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O • GPSSD: general-purpose SSD <p>Default value N/A</p>
data_volume_count	No	Integer	<p>Explanation Number of data disks of a node.</p> <p>Constraints N/A</p> <p>Value range 0-20</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
data_volume_size	No	Integer	<p>Explanation Data disk storage space of a node. Unit: GB.</p> <p>Constraints N/A</p> <p>Value range 100-32000</p> <p>Default value N/A</p>
auto_scaling_policy	No	auto_scaling_policy object	<p>Explanation The auto scaling policy.</p> <p>Constraints The auto scaling rule information. This parameter is available only when group_name is set to task_node_analysis_group or task_node_streaming_group.</p> <p>For details about the parameters, see Table 7-7.</p>

Table 7-11 ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	<p>Explanation Cycle type of a resource plan. This parameter can be set to daily only.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
start_time	Yes	String	<p>Explanation The start time of a resource plan. The value is in the format of hour:minute, indicating that the time ranges from 00:00 to 23:59.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
end_time	Yes	String	<p>Explanation End time of a resource plan. The format is the same as that of start_time.</p> <p>Constraints The value cannot be earlier than the start_time, and the interval between start_time and start_time cannot be less than 30 minutes.</p> <p>Value range N/A</p> <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation Minimum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
max_capacity	Yes	Integer	<p>Explanation Maximum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Table 7-12 Rule

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range It contains only 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. Rule names must be unique in a node group.</p> <p>Default value N/A</p>
description	No	String	<p>Explanation Description about an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range It contains a maximum of 1024 characters.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
adjustment_type	Yes	String	<p>Explanation Auto scaling rule adjustment type.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_out: cluster scale-out • scale_in: cluster scale-in <p>Default value N/A</p>
cool_down_minutes	Yes	Integer	<p>Explanation Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.</p> <p>Constraints N/A</p> <p>Value range The value ranges from 0 to 10080. 10080 indicates the number of minutes in a week.</p> <p>Default value N/A</p>
scaling_adjustment	Yes	Integer	<p>Explanation Number of nodes that can be adjusted once.</p> <p>Constraints N/A</p> <p>Value range 1-100</p> <p>Default value N/A</p>
trigger	Yes	trigger object	<p>Explanation Condition for triggering a rule. For details, see Table 7-13.</p> <p>Constraints N/A</p>

Table 7-13 Trigger

Parameter	Mandatory	Type	Description
metric_name	Yes	String	<p>Explanation Metric name. This triggering condition makes a judgment according to the value of the metric.</p> <p>Constraints N/A</p> <p>Value range A metric name contains a maximum of 64 characters.</p> <p>Default value N/A</p>
metric_value	Yes	String	<p>Explanation Metric threshold to trigger a rule The value must be an integer or a number with two decimal places.</p> <p>Constraints N/A</p> <p>Value range Only integers or numbers with two decimal places are allowed.</p> <p>Default value N/A</p>
comparison_operator	No	String	<p>Explanation Metric judgment logic operator.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • LT: less than • GT: greater than • LTOE: less than or equal to • GTOE: greater than or equal to <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
evaluation_periods	Yes	Integer	<p>Explanation The number of consecutive five-minute periods, during which a metric threshold is reached.</p> <p>Constraints N/A</p> <p>Value range 1-200</p> <p>Default value N/A</p>

Table 7-14 ScaleScript

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Names of custom scaling automation scripts.</p> <p>Constraints N/A</p> <p>Value range The names in the same cluster must be unique. The value can contain 1 to 64 characters, including only digits, letters, spaces, hyphens (-), and underscores (_), and cannot start with a space.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
uri	Yes	String	<p>Explanation</p> <p>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"> • OBS bucket path: Enter a script path manually, for example, s3a://XXX/scale.sh. • Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
parameters	No	String	<p>Explanation</p> <p>Parameters of a custom automation script. Separate multiple parameters by spaces. The following predefined parameters can be transferred:</p> <ul style="list-style-type: none"> • <i>`\${mrs_scale_node_num}`</i>: Number of the nodes to be added or removed • <i>`\${mrs_scale_type}`</i>: Scaling type. The value can be scale_out or scale_in. • <i>`\${mrs_scale_node_hostnames}`</i>: Host names of the nodes to be added or removed • <i>`\${mrs_scale_node_ips}`</i>: IP addresses of the nodes to be added or removed • <i>`\${mrs_scale_rule_name}`</i>: Name of the rule that triggers auto scaling <p>Other user-defined parameters are used in the same way as those of common shell scripts. Parameters are separated by space.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
nodes	Yes	Array of string	<p>Explanation</p> <p>Type of a node where the custom automation script is executed. The node type can be Master, Core, or Task.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
active_master	No	Boolean	<p>Explanation Whether the custom automation script runs only on the active master node.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The custom automation script runs only on the active Master nodes. • false: The custom automation script can run on all Master nodes. <p>Default value false</p>
action_stage	Yes	String	<p>Explanation Time when a script is executed.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • before_scale_out: before scale-out • before_scale_in: before scale-in • after_scale_out: after scale-out • after_scale_in: after scale-in <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
fail_action	Yes	String	<p>Explanation Whether to continue to execute subsequent scripts and create a cluster after the custom automation script fails to be executed. You are advised to set this parameter to continue in the commissioning phase so the cluster can continue to be installed and started no matter whether the custom automation script is executed successfully. The scale-in operation cannot be undone. fail_action must be set to continue for the scripts that are executed after scale-in.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value N/A</p>

Response Parameters

Status code: 200

Table 7-15 Response body parameters

Parameter	Type	Description
cluster_id	String	<p>Explanation Cluster ID, which is returned by the system after the cluster is created.</p> <p>Value range N/A</p>

Parameter	Type	Description
result	Boolean	<p>Explanation Operation result.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The operation is successful. • false: The operation failed.
msg	String	<p>Explanation System message, which can be empty.</p> <p>Value range N/A</p>

Example Request

- Use the **node_groups** parameter group to create a cluster with the HA function enabled. The cluster version is MRS 3.3.1-LTS.

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

```
{
  "billing_type" : 12,
  "data_center" : "",
  "available_zone_id" : "az1.cn-fcs-1a",
  "cluster_name" : "mrs_HEbK",
  "cluster_version" : "MRS 3.3.1-LTS",
  "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",
  "tags" : [ {
    "key" : "key1",
    "value" : "value1"
  }, {
    "key" : "key2",
    "value" : "value2"
  } ],
  "node_groups" : [ {
    "group_name" : "master_node_default_group",
```

```
"node_num" : 2,
"node_size" : "c3.2xlarge.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" : "c3.2xlarge.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" : "c3.2xlarge.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/ouput/",
  "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 3.3.1-LTS.

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

```

{
  "billing_type": 12,
  "data_center": "",
  "master_node_num": 2,
  "master_node_size": "c3.2xlarge.2.linux.bigdata",
  "core_node_num": 3,
  "core_node_size": "c3.2xlarge.2.linux.bigdata",
  "available_zone_id": "az1.cn-fcs-1a",

```



```
"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 3.3.1-LTS",
"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" : "c3.2xlarge.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
    }
  ]
}, {
  "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/ouput/",
  "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 3.3.1-LTS.

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

```
{
  "billing_type": 12,
  "data_center": "",
  "available_zone_id": "az1.cn-fcs-1a",
  "cluster_name": "mrs_HEbK",
  "cluster_version": "MRS 3.3.1-LTS",
  "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",
  "tags": [ {
    "key": "key1",
    "value": "value1"
  }, {
    "key": "key2",
    "value": "value2"
  } ],
  "node_groups": [ {
    "group_name": "master_node_default_group",
    "node_num": 1,
    "node_size": "c3.2xlarge.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": "c3.2xlarge.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/ouput/",
"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 3.3.1-LTS.
POST https://{endpoint}/v1.1/{project_id}/run-job-flow

```
{
  "billing_type" : 12,
  "data_center" : "",
  "master_node_num" : 1,
  "master_node_size" : "c3.2xlarge.2.linux.bigdata",
  "core_node_num" : 1,
  "core_node_size" : "c3.2xlarge.2.linux.bigdata",
  "available_zone_id" : "az1.cn-fcs-1a",
  "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" : "",
  "tags" : [ {
    "key" : "key1",
    "value" : "value1"
  }, {
    "key" : "key2",
    "value" : "value2"
  } ],
  "cluster_version" : "MRS 3.3.1-LTS",
  "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.

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	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 7-18 Request parameters

Parameter	Mandatory	Type	Description
service_id	No	String	<p>Explanation Service ID. This parameter is reserved for extension. You do not need to set this parameter.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
plan_id	No	String	<p>Explanation Plan ID. This parameter is reserved for extension. You do not need to set this parameter.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
parameters	Yes	ClusterScalingParams object	<p>Explanation Request parameters. For details, see Table 7-19.</p> <p>Constraints N/A</p>
previous_values	No	Map<String,String>	<p>Explanation This parameter is an extended API and needs to be reserved. You do not need to set this parameter.</p> <p>Constraints N/A</p>

Table 7-19 ClusterScalingParams

Parameter	Mandatory	Type	Description
order_id	No	String	<p>Explanation Order ID obtained by the system during scale-out or scale-in. You do not need to set the parameter.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
scale_type	Yes	String	<p>Explanation Operation type on the cluster node.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_in: cluster scale-in • scale_out: cluster scale-out <p>Default value N/A</p>
node_id	Yes	String	<p>Explanation ID of the newly added or removed node. The parameter value is fixed to node_orderadd.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
node_group	No	String	<p>Explanation Node group to be scaled out or in</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • core_node_default_group: indicates the Core node group. • task_node_default_group: indicates the Task node group. <p>Default value core_node_default_group</p>

Parameter	Mandatory	Type	Description
task_node_info	No	Object	<p>Explanation Task node specifications. For more parameter description, see Table 7-20.</p> <p>Constraints</p> <ul style="list-style-type: none"> • When the number of Task nodes is 0, this parameter is used to specify Task node specifications. • When the number of Task nodes is greater than 0, this parameter is unavailable. If this parameter is left blank, its mandatory field becomes unavailable and will pass an empty button for the null object.

Parameter	Mandatory	Type	Description
instances	Yes	Integer	<p>Explanation Number of nodes to be added or removed</p> <p>Constraints</p> <ul style="list-style-type: none"> • 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. • 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. • 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. <p>Value range ≥1</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
skip_bootstrap_scripts	No	String	<p>Explanation Whether to skip the bootstrap action. 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.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Skip the bootstrap action. • false: Execute the bootstrap action. <p>Default value false</p>
scale_without_start	No	Boolean	<p>Explanation Whether to start components on the added nodes after cluster scale-out</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Do not start components after scale-out. • false: Start components after scale-out. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
server_ids	No	Array of strings	<p>Explanation ID list of Task nodes to be deleted during task node scale-in.</p> <p>Constraints</p> <ul style="list-style-type: none"> • This parameter does not take effect when scale_type is set to scale-out. • If scale_type is set to scale-in and cannot be left blank, the system deletes the specified Task nodes. • When scale_type is set to scale-in and server_ids is left blank, the system automatically deletes the Task nodes based on the system rules.

Table 7-20 task_node_info parameters

Parameter	Mandatory	Type	Description
node_size	Yes	String	<p>Explanation Instance specifications of a Task node, for example, c3.4xlarge.2.linux.bigdata. 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.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
data_volume_type	Yes	String	<p>Explanation Data disk storage type of the Task node. Supported types include SATA, SAS, and SSD.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O • GPSSD: general-purpose SSD <p>Default value N/A</p>
data_volume_count	Yes	Integer	<p>Explanation Number of data disks of a Task node.</p> <p>Constraints N/A</p> <p>Value range 0-20</p> <p>Default value N/A</p>
data_volume_size	Yes	Integer	<p>Explanation Data disk storage space of a Task node. You only need to pass in a number without the unit GB.</p> <p>Constraints N/A</p> <p>Value range 100-32000</p> <p>Default value N/A</p>

Response Parameters

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

Table 7-21 Response parameter

Parameter	Type	Description
result	String	<p>Explanation Operation result</p> <p>Value range</p> <ul style="list-style-type: none"> ● succeeded: The operation is successful. ● Error Codes describes the error codes returned upon operation failures.

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 c3.2xlarge.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": "c3.2xlarge.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.

URI

- Format
GET /v1.1/{project_id}/cluster_infos
- Parameter description

Table 7-23 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Table 7-24 Query parameters

Parameter	Mandatory	Type	Description
pageSize	No	String	<p>Explanation Maximum number of clusters displayed on a page</p> <p>Constraints N/A</p> <p>Value range 1-2147483646</p> <p>Default value 10</p>
currentPage	No	String	<p>Explanation Current page number</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value 1</p>

Parameter	Mandatory	Type	Description
clusterName	No	String	Explanation The cluster name. Constraints N/A Value range N/A Default value N/A

Parameter	Mandatory	Type	Description
clusterState	No	String	<p>Explanation Query clusters by status.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • existing: Query existing clusters, including all clusters except those in the Deleted state and the yearly/monthly clusters in the Order processing or Preparing state. • history: 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. • starting: Query a list of clusters that are being started. • running: Query a list of running clusters. • terminated: Query a list of terminated clusters. • failed: Query a list of failed clusters. • abnormal: Query a list of abnormal clusters. • terminating: Query a list of clusters that are being terminated. • frozen: Query a list of frozen clusters. • scaling-out: Query a list of clusters that are being scaled out. • scaling-in: Query a list of clusters that are being scaled in. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
tags	No	String	<p>Explanation Tag list. You can search for a cluster by its tag. If you specify multiple tags, the relationship between them is AND.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> The format of the tags parameter is tags=k1*v1,k2*v2,k3*v3. When the values of some tags are null, the format is tags=k1,k2,k3*v3. <p>Default value N/A</p>
enterpriseProjectId	No	String	<p>Explanation The enterprise project ID used to query clusters in a specified enterprise project.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value 0</p>

Request Parameters

None

Response Parameters

Table 7-25 Response parameters

Parameter	Type	Description
clusterTotal	Integer	<p>Explanation Total number of clusters in a list</p> <p>Value range N/A</p>

Parameter	Type	Description
clusters	Array of Cluster objects	Explanation Cluster parameters. For details, see Table 7-26 .

Table 7-26 clusters parameters

Parameter	Type	Description
clusterId	String	Explanation Cluster ID. Value range N/A
clusterName	String	Explanation Cluster name. Value range N/A
masterNodeNum	String	Explanation Number of Master nodes deployed in a cluster. Value range N/A
coreNodeNum	String	Explanation Number of Core nodes deployed in a cluster. Value range N/A
totalNodeNum	String	Explanation Total number of nodes deployed in a cluster. Value range N/A

Parameter	Type	Description
clusterState	String	<p>Explanation Cluster status.</p> <p>Value range</p> <ul style="list-style-type: none"> • starting: The cluster is being started. • running: The cluster is running. • terminated: The cluster has been terminated. • failed: The cluster fails. • abnormal: The cluster is abnormal. • terminating: The cluster is being terminated. • frozen: The cluster has been frozen. • scaling-out: The cluster is being scaled out. • scaling-in: The cluster is being scaled in.
createAt	String	<p>Explanation Cluster creation time, which is a 10-bit timestamp.</p> <p>Value range N/A</p>
updateAt	String	<p>Explanation Cluster update time, which is a 10-bit timestamp.</p> <p>Value range N/A</p>
dataCenter	String	<p>Explanation Cluster work region.</p> <p>Value range N/A</p>
vpc	String	<p>Explanation VPC name.</p> <p>Value range N/A</p>
vpcId	String	<p>Explanation VPC ID.</p> <p>Value range N/A</p>

Parameter	Type	Description
hadoopVersion	String	Explanation Hadoop version Value range N/A
masterNodeSize	String	Explanation Instance specifications of a Master node. Value range N/A
coreNodeSize	String	Explanation Instance specifications of a Core node. Value range N/A
componentList	Array	Explanation Component list. For details, see Table 7-27 .
externallp	String	Explanation External IP address. This IP address is opened to the outside of the cluster as access to the active Master process. Value range N/A
externalAlternatelp	String	Explanation Backup external IP address. This IP address is used outside the cluster as access to the standby Master process. Value range N/A
internallp	String	Explanation Internal IP address. This IP address is used within the cluster as access to Master processes. Value range N/A

Parameter	Type	Description
deploymentId	String	Explanation Cluster deployment ID. The ID is used by cloud services to query details about cluster deployment tasks. Value range N/A
remark	String	Explanation Cluster remarks. By default, this parameter is left blank. Value range N/A
orderId	String	Explanation Cluster creation order ID. Value range N/A
azId	String	Explanation AZ ID. Value range N/A
masterNodeProductId	String	Explanation Product ID of a Master node. Value range N/A
masterNodeSpecId	String	Explanation Specification ID of a Master node. Value range N/A
coreNodeProductId	String	Explanation Product ID of a Core node. Value range N/A
coreNodeSpecId	String	Explanation Specification ID of a Core node Value range N/A

Parameter	Type	Description
azName	String	Explanation AZ name. Value range N/A
azCode	String	Explanation AZ name (en). Value range N/A
availabilityZoneId	String	Explanation Availability zone. Value range N/A
instanceId	String	Explanation Instance ID. Value range N/A
vnc	String	Explanation URI for remotely logging in to an ECS. Value range N/A
tenantId	String	Explanation Project ID. Value range N/A
volumeSize	Integer	Explanation Disk storage space. Value range N/A
volumeType	String	Explanation Disk type. Value range N/A
subnetId	String	Explanation Subnet ID. Value range N/A

Parameter	Type	Description
clusterType	Integer	<p>Explanation Cluster type.</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: analysis cluster • 1: streaming cluster • 2: hybrid cluster • 3: custom cluster • 4: Offline cluster
subnetName	String	<p>Explanation Subnet name.</p> <p>Value range N/A</p>
securityGroupsId	String	<p>Explanation Security group ID.</p> <p>Value range N/A</p>
slaveSecurityGroupsId	String	<p>Explanation 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 securityGroupsId does for compatibility consideration.</p> <p>Value range N/A</p>
bootstrapScripts	Array of BootstrapScript objects	<p>Explanation Bootstrap action script. For details, see Table 7-28.</p>

Parameter	Type	Description
stageDesc	String	<p>Explanation Cluster process.</p> <p>Value range The cluster installation progress includes:</p> <ul style="list-style-type: none"> • Verifying cluster parameters: Cluster parameters are being verified. • Applying for cluster resources: Cluster resources are being applied for. • Creating VMs: The VMs are being created. • Initializing VMs: The VMs are being initialized. • Installing MRS Manager: MRS Manager is being installed. • Deploying the cluster: The cluster is being deployed. • Cluster installation failed: Failed to install the cluster. <p>The cluster scale-out progress includes:</p> <ul style="list-style-type: none"> • Preparing for scale-out: Cluster scale-out is being prepared. • Creating VMs: The VMs are being created. • Initializing VMs: The VMs are being initialized. • Adding nodes to the cluster: The nodes are being added to the cluster. • Scale-out failed: Failed to scale out the cluster. <p>The cluster scale-in progress includes:</p> <ul style="list-style-type: none"> • Preparing for scale-in: Cluster scale-in is being prepared. • Decommissioning instance: The instance is being decommissioned. • Deleting VMs: The VMs are being deleted. • Deleting nodes from the cluster: The nodes are being deleted from the cluster.

Parameter	Type	Description
		<ul style="list-style-type: none"> Scale-in failed: Failed to scale in the cluster. <p>If the cluster installation, scale-out, or scale-in fails, stageDesc will display the failure cause.</p>
isMrsManagerFinish	Boolean	<p>Explanation</p> <p>Whether MRS Manager installation is finished during cluster creation.</p> <p>Value range</p> <ul style="list-style-type: none"> true: MRS Manager installation is finished. false: MRS Manager installation is not finished.
safeMode	Integer	<p>Explanation</p> <p>Run mode of an MRS cluster.</p> <p>Value range</p> <ul style="list-style-type: none"> 0: Normal cluster 1: Security cluster
clusterVersion	String	<p>Explanation</p> <p>Cluster version.</p> <p>Value range</p> <p>N/A</p>
nodePublicCertificateName	String	<p>Explanation</p> <p>Name of the key file.</p> <p>Value range</p> <p>N/A</p>
masterNodeIp	String	<p>Explanation</p> <p>IP address of a Master node.</p> <p>Value range</p> <p>N/A</p>
privateIpFirst	String	<p>Explanation</p> <p>Preferred private IP address.</p> <p>Value range</p> <p>N/A</p>
errorInfo	String	<p>Explanation</p> <p>Error message.</p> <p>Value range</p> <p>N/A</p>

Parameter	Type	Description
tags	String	<p>Explanation Tag information.</p> <p>Value range N/A</p>
logCollection	Integer	<p>Explanation Whether to collect logs when cluster installation fails.</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: Do not collect logs. • 1: Collect logs.
taskNodeGroups	List<NodeGroup>	<p>Explanation List of Task nodes. For more parameter description, see Table 7-29.</p>
nodeGroups	List<NodeGroup>	<p>Explanation List of Master, Core and Task nodes. For more parameter description, see Table 7-29.</p>
masterDataVolumeType	String	<p>Explanation Data disk storage type of the Master node. Currently, SATA, SAS, and SSD are supported.</p> <p>Value range N/A</p>
masterDataVolumeSize	Integer	<p>Explanation Data disk storage space of Master nodes, in GB To increase data storage capacity, you can add disks at the same time when creating a cluster.</p> <p>Value range 100-32000</p>
masterDataVolumeCount	Integer	<p>Explanation Number of data disks of the Master node</p> <p>Value range The value can only be 1.</p>

Parameter	Type	Description
coreDataVolumeType	String	<p>Explanation Data disk storage type of the Core node. Currently, SATA, SAS, and SSD are supported.</p> <p>Value range N/A</p>
coreDataVolumeSize	Integer	<p>Explanation Data disk storage space of Core nodes, in GB To increase data storage capacity, you can add disks at the same time when creating a cluster.</p> <p>Value range 100-32000</p>
coreDataVolumeCount	Integer	<p>Explanation Number of data disks of the Core node.</p> <p>Value range 1-10</p>

Parameter	Type	Description
scale	String	<p>Explanation Status of node changes If this parameter is left blank, no change operation is performed on a cluster node.</p> <p>Value range</p> <ul style="list-style-type: none"> • Scaling-out: The cluster is being scaled out. • Scaling-in: The cluster is being scaled in. • scaling-error: 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. • scaling-up: The master node specifications are being scaled up. • scaling_up_first: The standby master node specifications are being scaled up. • scaled_up_first: The standby master node specifications have been scaled up. • scaled-up-success: The master node specifications have been scaled up.
eipId	String	<p>Explanation Unique ID of the cluster EIP</p> <p>Value range N/A</p>
eipAddress	String	<p>Explanation IPv4 address of the cluster EIP</p> <p>Value range N/A</p>
eipv6Address	String	<p>Explanation IPv6 address of the cluster EIP. This parameter is not returned when an IPv4 address is used.</p> <p>Value range N/A</p>

Table 7-27 componentAmb

Parameter	Type	Description
componentId	String	Explanation Component ID Value range For example, the component_id of Hadoop is MRS 3.3.1-LTS_001 .
componentName	String	Explanation Component name. Value range N/A
componentVersion	String	Explanation Component version. Value range N/A
componentDesc	String	Explanation Component description. Value range N/A

Table 7-28 bootstrapScripts

Parameter	Type	Description
name	String	Explanation Name of a bootstrap action script. Constraints N/A Value range The names of bootstrap action scripts in the same cluster must be unique. The value can contain 1 to 64 characters, including only digits, letters, spaces, hyphens (-), and underscores (_), and cannot start with a space. Default value N/A

Parameter	Type	Description
uri	String	<p>Explanation</p> <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"> • OBS bucket path: Enter a script path. For example, enter the path of the public sample script provided by MRS. Example: s3a://bootstrap/presto/presto-install.sh. If dualroles is installed, the parameter of the presto-install.sh script is dualroles. If worker is installed, the parameter of the presto-install.sh script is worker. Based on the Presto usage habit, you are advised to install dualroles on the active master nodes and worker on the core nodes. • Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
parameters	String	<p>Explanation</p> <p>The bootstrap action script parameters.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
nodes	Array of strings	<p>Explanation</p> <p>Type of a node where the bootstrap action script is executed. The value can be Master, Core, or Task.</p> <p>Constraints</p> <p>The node type must be represented in lowercase letters.</p>

Parameter	Type	Description
active_master	Boolean	<p>Explanation Whether the bootstrap action script runs only on active master nodes.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script runs only on active Master nodes. • false: The bootstrap action script can run on all Master nodes. <p>Default value false</p>
fail_action	String	<p>Explanation Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to execute. You are advised to set this parameter to continue in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value errorout</p>

Parameter	Type	Description
before_component_start	Boolean	<p>Explanation Time when the bootstrap action script is executed. Currently, the following two options are available: Before component start and After component start.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script is executed before the component starts. • false: The bootstrap action script is executed after the component starts. <p>Default value false</p>
start_time	Long	<p>Explanation The execution time of one bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
state	String	<p>Explanation The running status of one bootstrap action script.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • PENDING: The script is suspended. • IN_PROGRESS: The script is being processed. • SUCCESS • FAILURE: The script fails to be executed.

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

Table 7-29 NodeGroupV10

Parameter	Type	Description
GroupName	String	<p>Explanation Node group name.</p> <p>Value range N/A</p>
NodeNum	Integer	<p>Explanation 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.</p> <p>Value range 0-500</p>
NodeSize	String	<p>Explanation Instance specifications of a node.</p> <p>Value range N/A</p>
NodeSpecId	String	<p>Explanation Instance specification ID of a node.</p> <p>Value range N/A</p>

Parameter	Type	Description
NodeProductId	String	Explanation Instance product ID of a node. Value range N/A
VmProductId	String	Explanation VM specifications of a node. Value range N/A
VmSpecCode	String	Explanation VM specifications of a node. Value range N/A
RootVolumeSize	Integer	Explanation System disk size of a node. This parameter is not configurable and its default value is 40 GB . Value range N/A
RootVolumeProductId	String	Explanation System disk product ID of a node. Value range N/A
RootVolumeType	String	Explanation System disk type of a node. Value range N/A
RootVolumeResourceSpecCode	String	Explanation System disk product specifications of a node. Value range N/A
RootVolumeResourceType	String	Explanation System disk product type of a node. Value range N/A

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

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": ,
"vpc": null,
"duration": "0",
"fee": null,
"hadoopVersion": null,
"masterNodeSize": null,
"coreNodeSize": null,
"componentList": [ {
  "componentId": "MRS 3.3.1-LTS_001",
  "componentName": "Hadoop",
  "componentVersion": "3.1.1",
  "componentDesc": "A framework that allows for the distributed processing of large data sets across
clusters."
}, {
  "componentId": "MRS 3.3.1-LTS_003",
  "componentName": "HBase",
  "componentVersion": "2.1.1",
  "componentDesc": "A scalable, distributed database that supports structured data storage for large
tables."
}, {
  "componentId": "MRS 3.3.1-LTS_002",
  "componentName": "Spark",
  "componentVersion": "2.3.2",
  "componentDesc": "A fast and general engine for large-scale data processing."
}, {
  "componentId": "MRS 3.3.1-LTS_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.cn-fcs-1a",
"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://bootstrap/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" : "c3.2xlarge.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" : "c3.2xlarge.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" : "c3.2xlarge.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": "c3.2xlarge.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",
  } ]
}

```

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.

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	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Table 7-32 Response body parameter

Parameter	Type	Description
cluster	Cluster object	Explanation Cluster parameters. For details, see Table 7-33 .

Table 7-33 Response parameters

Parameter	Type	Description
clusterId	String	Explanation Cluster ID. Value range N/A
clusterName	String	Explanation Cluster name. Value range N/A
totalNodeNum	String	Explanation Total number of nodes deployed in a cluster. Value range N/A

Parameter	Type	Description
clusterState	String	<p>Explanation Cluster status.</p> <p>Value range</p> <ul style="list-style-type: none"> ● starting: The cluster is being started. ● running: The cluster is running. ● terminated: The cluster has been deleted. ● failed: The cluster is failed. ● abnormal: The cluster is abnormal. ● terminating: The cluster is being deleted. ● frozen: The cluster is frozen. ● scaling-out: The cluster is being scaled out. ● scaling-in: The cluster is being scaled in.

Parameter	Type	Description
stageDesc	String	<p>Explanation Cluster process.</p> <p>Value range The cluster installation progress includes:</p> <ul style="list-style-type: none"> • Verifying cluster parameters: Cluster parameters are being verified. • Applying for cluster resources: Cluster resources are being applied for. • Creating VMs: The VMs are being created. • Initializing VMs: The VMs are being initialized. • Installing MRS Manager: MRS Manager is being installed. • Deploying the cluster: The cluster is being deployed. • Cluster installation failed: Failed to install the cluster. <p>The cluster scale-out progress includes:</p> <ul style="list-style-type: none"> • Preparing for scale-out: Cluster scale-out is being prepared. • Creating VMs: The VMs are being created. • Initializing VMs: The VMs are being initialized. • Adding nodes to the cluster: The nodes are being added to the cluster. • Scale-out failed: Failed to scale out the cluster. <p>The cluster scale-in progress includes:</p> <ul style="list-style-type: none"> • Preparing for scale-in: Cluster scale-in is being prepared. • Decommissioning instance: The instance is being decommissioned. • Deleting VMs: The VMs are being deleted. • Deleting nodes from the cluster: The nodes are being deleted from the cluster.

Parameter	Type	Description
		<ul style="list-style-type: none"> Scale-in failed: Failed to scale in the cluster. <p>If the cluster installation, scale-out, or scale-in fails, stageDesc will display the failure cause.</p>
createAt	String	<p>Explanation Cluster creation time, which is a 10-bit timestamp.</p> <p>Value range N/A</p>
updateAt	String	<p>Explanation Cluster update time, which is a 10-bit timestamp.</p> <p>Value range N/A</p>
dataCenter	String	<p>Explanation Cluster work region.</p> <p>Value range N/A</p>
vpc	String	<p>Explanation VPC name.</p> <p>Value range N/A</p>
vpcId	String	<p>Explanation VPC ID.</p> <p>Value range N/A</p>
hadoopVersion	String	<p>Explanation Hadoop version.</p> <p>Value range N/A</p>
componentList	Array	<p>Explanation Component list. For details, see Table 7-34.</p>

Parameter	Type	Description
externallp	String	<p>Explanation External IP address This IP address is opened to the outside of the cluster as access to the active Master process.</p> <p>Value range N/A</p>
externalAlternatelp	String	<p>Explanation Backup external IP address. This IP address is used outside the cluster as access to the standby Master process.</p> <p>Value range N/A</p>
internallp	String	<p>Explanation Internal IP address This IP address is used within the cluster as access to Master processes.</p> <p>Value range N/A</p>
deploymentId	String	<p>Explanation Cluster deployment ID. The ID is used by cloud services to query details about cluster deployment tasks.</p> <p>Value range N/A</p>
remark	String	<p>Explanation Cluster remarks. By default, this parameter is left blank.</p> <p>Value range N/A</p>
orderId	String	<p>Explanation Cluster creation order ID.</p> <p>Value range N/A</p>
azId	String	<p>Explanation AZ ID</p> <p>Value range N/A</p>

Parameter	Type	Description
azName	String	Explanation AZ name. Value range N/A
azCode	String	Explanation AZ name (en). Value range N/A
availabilityZoneId	String	Explanation Availability zone. Value range N/A
instanceId	String	Explanation Instance ID. Value range N/A
vnc	String	Explanation URI for remotely logging in to an ECS. Value range N/A
tenantId	String	Explanation Project ID. Value range N/A
volumeSize	Integer	Explanation Disk storage space. Value range N/A
volumeType	String	Explanation Disk type. Value range N/A
subnetId	String	Explanation Subnet ID. Value range N/A

Parameter	Type	Description
subnetName	String	Explanation Subnet name. Value range N/A
securityGroupsId	String	Explanation Security group ID. Value range N/A
slaveSecurityGroupsId	String	Explanation 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 securityGroupsId does for compatibility consideration. Value range N/A
bootstrapScripts	Array of BootstrapScript objects	Explanation Bootstrap action script information. For more parameter description, see Table 7-36 .
safeMode	Integer	Explanation Run mode of an MRS cluster. Value range <ul style="list-style-type: none"> • 0: Normal cluster • 1: Security cluster
clusterVersion	String	Explanation Cluster version. Value range N/A
nodePublicKeyName	String	Explanation Name of the key file. Value range N/A

Parameter	Type	Description
masterNodeIp	String	Explanation IP address of a Master node. Value range N/A
privateIpFirst	String	Explanation Preferred private IP address. Value range N/A
errorInfo	String	Explanation Error message. Value range N/A
tags	String	Explanation Tag information. Value range N/A
masterNodeNum	String	Explanation Number of Master nodes deployed in a cluster. Value range N/A
coreNodeNum	String	Explanation Number of Core nodes deployed in a cluster. Value range N/A
masterNodeSize	String	Explanation Instance specifications of a Master node. Value range N/A
coreNodeSize	String	Explanation Instance specifications of a Core node. Value range N/A

Parameter	Type	Description
masterNodeProductId	String	Explanation Product ID of a Master node. Value range N/A
masterNodeSpecId	String	Explanation Specification ID of a Master node. Value range N/A
coreNodeProductId	String	Explanation Product ID of a Core node. Value range N/A
coreNodeSpecId	String	Explanation Specification ID of a Core node. Value range N/A
masterDataVolumeType	String	Explanation Data disk storage type of the Master node. Currently, SATA, SAS, and SSD are supported. Value range N/A
masterDataVolumeSize	Integer	Explanation 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. Unit: GB. Value range 100-32000
masterDataVolumeCount	Integer	Explanation Number of data disks of the Master node. Value range The value can only be 1.

Parameter	Type	Description
coreDataVolumeType	String	<p>Explanation Data disk storage type of the Core node. Currently, SATA, SAS, and SSD are supported.</p> <p>Value range N/A</p>
coreDataVolumeSize	Integer	<p>Explanation 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. Unit: GB.</p> <p>Value range 100-32000</p>
coreDataVolumeCount	Integer	<p>Explanation Number of data disks of the Core node.</p> <p>Value range 1-10</p>
isMrsManagerFinish	Boolean	<p>Explanation Whether MRS Manager installation is finished during cluster creation.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: MRS Manager installation is finished. • false: MRS Manager installation is not finished.
clusterType	Integer	<p>Explanation Cluster type.</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: analysis cluster • 1: streaming cluster • 2: hybrid cluster • 3: custom cluster • 4: Offline cluster

Parameter	Type	Description
logCollection	Integer	<p>Explanation Whether to collect logs when cluster installation fails</p> <p>Value range</p> <ul style="list-style-type: none"> • 0: Do not collect. • 1: Collect.
scale	String	<p>Explanation Node change status. If this parameter is left blank, no change operation is performed on a cluster node.</p> <p>Value range</p> <ul style="list-style-type: none"> • scaling-out: The cluster is being scaled out. • scaling-in: The cluster is being scaled in. • scaling-error: 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. • scaling-up: The Master node specifications are being scaled up. • scaling_up_first: The standby Master node specifications are being scaled up. • scaled_up_first: The standby Master node specifications have been scaled up successfully. • scaled-up-success: The master node specifications have been scaled up.
nodeGroups	List<NodeGroup>	<p>Explanation List of Master, Core and Task nodes. For details about this parameter, see Table 7-35.</p>
taskNodeGroups	List<NodeGroup>	<p>Explanation List of Task nodes. For more parameter description, see Table 7-35.</p>
eipId	String	<p>Explanation Unique ID of the cluster EIP</p> <p>Value range N/A</p>

Parameter	Type	Description
eipAddress	String	Explanation IPv4 address of the cluster EIP. Value range N/A
eipv6Address	String	Explanation IPv6 address of the cluster EIP. This parameter is not returned when an IPv4 address is used. Value range N/A

Table 7-34 componentAmb

Parameter	Type	Description
componentId	String	Explanation Component ID Value range For example, the component_id of Hadoop is MRS 3.3.1-LTS_001 .
componentName	String	Explanation Component name. Value range N/A
componentVersion	String	Explanation Component version. Value range N/A
componentDesc	String	Explanation Component description. Value range N/A

Table 7-35 NodeGroupV10

Parameter	Type	Description
GroupName	String	Explanation Node group name. Value range N/A
NodeNum	Integer	Explanation 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. Value range 0-500
NodeSize	String	Explanation Instance specifications of a node. Value range N/A
NodeSpecId	String	Explanation Instance specification ID of a node. Value range N/A
NodeProductId	String	Explanation Instance product ID of a node. Value range N/A
VmProductId	String	Explanation VM specifications of a node. Value range N/A
VmSpecCode	String	Explanation VM specifications of a node. Value range N/A

Parameter	Type	Description
RootVolumeSize	Integer	<p>Explanation System disk size of a node. This parameter is not configurable and its default value is 40 GB.</p> <p>Value range N/A</p>
RootVolumeProductId	String	<p>Explanation System disk product ID of a node.</p> <p>Value range N/A</p>
RootVolumeType	String	<p>Explanation System disk type of a node.</p> <p>Value range N/A</p>
RootVolumeResourceSpecCode	String	<p>Explanation System disk product specifications of a node.</p> <p>Value range N/A</p>
RootVolumeResourceType	String	<p>Explanation System disk product type of a node.</p> <p>Value range N/A</p>
DataVolumeType	String	<p>Explanation Data disk storage type of a node. Currently, SATA, SAS, and SSD are supported.</p> <p>Value range</p> <ul style="list-style-type: none"> • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O
DataVolumeCount	Integer	<p>Explanation Number of data disks of a node.</p> <p>Value range N/A</p>

Parameter	Type	Description
DataVolumeSize	Integer	Explanation Data disk storage space of a node. Value range N/A
DataVolumeProductId	String	Explanation Data disk product ID of a node. Value range N/A
DataVolumeResourceSpecCode	String	Explanation Data disk specifications of a node. Value range N/A
DataVolumeResourceType	String	Explanation Data disk type of a node. Value range N/A

Table 7-36 BootstrapScripts

Parameter	Type	Description
name	String	Explanation Name of a bootstrap action script. Constraints N/A Value range The names of bootstrap action scripts in the same cluster must be unique. The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-), and cannot start with a space. Default value N/A

Parameter	Type	Description
uri	String	<p>Explanation</p> <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"> OBS bucket path: Enter a script path manually. For example, enter the path of the public sample script provided by MRS. Example: s3a://bootstrap/presto/presto-install.sh. If dualroles is installed, the parameter of the presto-install.sh script is dualroles. If worker is installed, the parameter of the presto-install.sh script is worker. Based on the Presto usage habit, you are advised to install dualroles on the active Master nodes and worker on the Core nodes. Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>
parameters	String	<p>Explanation</p> <p>Bootstrap action script parameters.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>N/A</p> <p>Default value</p> <p>N/A</p>
nodes	Array String	<p>Explanation</p> <p>Type of a node where the bootstrap action script is executed. The value can be Master, Core, or Task.</p> <p>Constraints</p> <p>The node type must be represented in lowercase letters.</p>

Parameter	Type	Description
active_master	Boolean	<p>Explanation Whether the bootstrap action script runs only on active Master nodes.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script runs only on active Master nodes. • false: The bootstrap action script can run on all Master nodes. <p>Default value false</p>
before_component_start	Boolean	<p>Explanation Time when the bootstrap action script is executed. Currently, the following two options are available: Before component start and After component start</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The bootstrap action script is executed before the component starts. • false: The bootstrap action script is executed after the component starts. <p>Default value false</p>

Parameter	Type	Description
fail_action	String	<p>Explanation</p> <p>Whether to continue executing subsequent scripts and creating a cluster after the bootstrap action script fails to be executed. You are advised to set this parameter to continue in the commissioning phase so that the cluster can continue to be installed and started no matter whether the bootstrap action is successful.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value</p> <p>errorout</p>
action_stages	Array of strings	<p>Explanation</p> <p>Select the time when the bootstrap action script is executed.</p> <ul style="list-style-type: none"> • BEFORE_COMPONENT_FIRST_START: before initial component starts • AFTER_COMPONENT_FIRST_START: after initial component starts • BEFORE_SCALE_IN: before scale-in • AFTER_SCALE_IN: after scale-in • BEFORE_SCALE_OUT: before scale-out • AFTER_SCALE_OUT: after scale-out <p>Constraints</p> <p>N/A</p>

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" : "",
"vpc" : "vpc-autotest",
"vpId" : "e2978efd-ca12-4058-9332-1ca0bfbab592",
"duration" : "0",
"fee" : "0",
"hadoopVersion" : "",
"masterNodeSize" : "c3.2xlarge.2.linux.bigdata",
"coreNodeSize" : "c3.2xlarge.2.linux.bigdata",
"componentList" : [ {
  "componentId" : "MRS 3.3.1-LTS_001",
  "componentName" : "Hadoop",
  "componentVersion" : "3.1.1",
  "componentDesc" : "A framework that allows for the distributed processing of large data sets across
clusters."
}, {
  "componentId" : "MRS 3.3.1-LTS_002",
  "componentName" : "Spark",
  "componentVersion" : "2.3.2",
  "componentDesc" : "A fast and general engine for large-scale data processing."
}, {
  "componentId" : "MRS 3.3.1-LTS_004",
  "componentName" : "Hive",
  "componentVersion" : "3.1.0",
  "componentDesc" : "A data warehouse infrastructure that provides data summarization and ad hoc
querying."
}, {
  "componentId" : "MRS 3.3.1-LTS_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.cn-fcs-1a",
"instanceId" : "4ac46ca7-a488-4b91-82c2-e4d7aa9c40c2",
"vnc" : null,
"tenantId" : "3f99e3319a8943ceb15c584f3325d064",
"volumeSize" : "600",
"volumeType" : "SATA",
"subnetId" : "6b96eec3-4f8d-4c83-93e2-6ec625001d7c",
"subnetName" : "subnet-ftest",
"securityGroupsId" : "930e34e2-195d-401f-af07-0b64ea6603f8",
"slaveSecurityGroupsId" : "2ef3343e-3477-4a0d-80fe-4d874e4f81b8",
"bootstrapScripts" : [ {
  "name" : "test1-success",
  "uri" : "s3a://bootstrap/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 3.3.1-LTS",
"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" : "c3.2xlarge.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" : "c3.2xlarge.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" : "c3.2xlarge.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": "c3.2xlarge.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",
}
}

```

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.

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	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Table 7-39 Query parameters

Parameter	Mandatory	Type	Description
pageSize	No	Integer	<p>Explanation Maximum number of clusters displayed on a page.</p> <p>Constraints N/A</p> <p>Value range 1-2147483646</p> <p>Default value 10</p>
currentPage	No	Integer	<p>Explanation Current page number.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value 1</p>

Request Parameters

None

Response Parameters

Table 7-40 Response body parameters

Parameter	Type	Description
total	Integer	<p>Explanation Host parameters.</p>
hosts	Array of HostModel objects	<p>Explanation Host parameters For details, see Table 7-41.</p> <p>Value range N/A</p>

Table 7-41 HostModel

Parameter	Type	Description
id	String	Explanation VM ID. Value range N/A
ip	String	Explanation VM IP address. Value range N/A
availability_zone_id	String	Explanation Availability zone. Value range N/A
tags	Array of TagPlain objects	Explanation The tag list. For details, see Table 7-42 .
flavor	String	Explanation VM specification ID. Value range N/A
type	String	Explanation VM type. The value can be MasterNode , CoreNode , or TaskNode . Value range N/A
name	String	Explanation VM name. Value range N/A
status	String	Explanation Current VM state. Value range N/A

Parameter	Type	Description
resource_id	String	Explanation Node resource ID. Value range N/A
mem	String	Explanation Memory. Value range N/A
cpu	String	Explanation Number of CPU cores. Value range N/A
root_volume_size	String	Explanation OS disk capacity. Value range N/A
data_volume_type	String	Explanation Data disk type. Value range N/A
data_volume_size	Integer	Explanation Data disk capacity. Value range N/A
data_volume_count	Integer	Explanation Number of data disks. Value range N/A
node_group_name	String	Explanation Node group name. Value range N/A

Table 7-42 TagPlain parameters

Parameter	Type	Description
key	String	<p>Explanation Tag key.</p> <p>Value range A tag key can contain letters, digits, spaces, and special characters <code>._:=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>.</p>
value	String	<p>Explanation Tag value.</p> <p>Value range The value can contain letters, digits, spaces, and special characters <code>._:=+@</code>, but cannot start or end with a space or start with <code>_sys_</code>.</p>

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

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	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID For details about how to obtain the cluster ID, see Obtaining the MRS Cluster Information.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

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.

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	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID For details about how to obtain the cluster ID, see Obtaining the MRS Cluster Information.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 7-47 Request body parameters

Parameter	Mandatory	Type	Description
node_group	Yes	String	<p>Explanation Type of the node to which an auto scaling rule applies. Currently, only task nodes support auto scaling rules.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> task_node_default_group: Task node <p>Default value N/A</p>
auto_scaling_policy	Yes	AutoScalingPolicy object	<p>Explanation The auto scaling policy. Table 7-48 describes these parameters.</p> <p>Constraints N/A</p>

Table 7-48 AutoScalingPolicy

Parameter	Mandatory	Type	Description
auto_scaling_enable	Yes	Boolean	<p>Explanation Whether to enable the auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> true: Enable the auto scaling rule. false: Disable the autoscaling rule. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
min_capacity	Yes	Integer	<p>Explanation Minimum number of nodes left in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation Maximum number of nodes in the node group.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
resources_plans	No	List	<p>Explanation Resource plan list. For details, see Table 7-49. If this parameter is left blank, the resource plan is disabled.</p> <p>Constraints When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured.</p>
exec_scripts	No	List	<p>Explanation List of custom scaling automation scripts. For details, see Table 7-50. If this parameter is left blank, a hook script is disabled.</p> <p>Constraints The number of records cannot exceed 10.</p>

Parameter	Mandatory	Type	Description
rules	No	List	<p>Explanation List of auto scaling rules. For details, see Table 7-51.</p> <p>Constraints When auto scaling is enabled, either a resource plan or an auto scaling rule must be configured. The number of records cannot exceed 10.</p>

Table 7-49 ResourcesPlan

Parameter	Mandatory	Type	Description
period_type	Yes	String	<p>Explanation Cycle type of a resource plan. This parameter can be set to daily only.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
start_time	Yes	String	<p>Explanation Start time of a resource plan. The value is in the format of hour:minute.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
end_time	Yes	String	<p>Explanation End time of a resource plan. The format is the same as that of start_time.</p> <p>Constraints The value cannot be earlier than the start_time, and the interval between start_time and start_time cannot be less than 30 minutes.</p> <p>Value range N/A</p> <p>Default value N/A</p>
min_capacity	Yes	Integer	<p>Explanation Minimum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>
max_capacity	Yes	Integer	<p>Explanation Maximum number of the preserved nodes in a node group in a resource plan.</p> <p>Constraints N/A</p> <p>Value range 0-500</p> <p>Default value N/A</p>

Table 7-50 ScaleScript

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Names of custom scaling automation scripts.</p> <p>Constraints N/A</p> <p>Value range The names in the same cluster must be unique. 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.</p> <p>Default value N/A</p>
uri	Yes	String	<p>Explanation 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"> • OBS bucket path: Enter a script path manually. for example, s3a://XXX/scale.sh. • Local VM path: Enter a script path. The script path must start with a slash (/) and end with .sh. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
parameters	No	String	<p>Explanation</p> <p>Parameters of a custom automation script.</p> <ul style="list-style-type: none"> Multiple parameters are separated by space. The following predefined system parameters can be transferred: <ul style="list-style-type: none"> <i>#{mrs_scale_node_num}</i>: Number of the nodes to be added or removed <i>#{mrs_scale_type}</i>: Scaling type. The value can be scale_out or scale_in. <i>#{mrs_scale_node_hostnames}</i>: Host names of the nodes to be added or removed <i>#{mrs_scale_node_ips}</i>: IP addresses of the nodes to be added or removed <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. <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>
nodes	Yes	Array of strings	<p>Explanation</p> <p>Type of a node where the custom automation script is executed. The node type can be Master, Core, or Task.</p> <p>Constraints N/A</p>

Parameter	Mandatory	Type	Description
active_master	No	Boolean	<p>Explanation Whether the custom automation script runs only on the active Master node.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • true: The custom automation script runs only on the active Master nodes. • false: The custom automation script can run on all Master nodes. <p>Default value false</p>
action_stage	Yes	String	<p>Explanation Time when a script is executed.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • before_scale_out: before scale-out • before_scale_in: before scale-in • after_scale_out: after scale-out • after_scale_in: after scale-in <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
fail_action	Yes	String	<p>Explanation Whether to continue to execute subsequent scripts and create a cluster after the custom automation script fails to be executed.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • continue: Continue to execute subsequent scripts. • errorout: Stop the action. <p>Default value N/A</p> <p>NOTE</p> <ul style="list-style-type: none"> • You are advised to set this parameter to continue 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. • The scale-in operation cannot be undone. Therefore, fail_action must be set to continue for the scripts that are executed after scale-in.

Table 7-51 Rule

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation Name of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range 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.</p> <p>Default value N/A</p>
description	No	String	<p>Explanation Description about an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range A string of 1 to 1024 characters</p> <p>Default value N/A</p>
adjustment_type	Yes	String	<p>Explanation Adjustment type of an auto scaling rule.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • scale_out: cluster scale-out • scale_in: cluster scale-in <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cool_down_minutes	Yes	Integer	<p>Explanation Cluster cooling time after an auto scaling rule is triggered, when no auto scaling operation is performed. The unit is minute.</p> <p>Constraints N/A</p> <p>Value range 0 to 10080. 10080 indicates the number of minutes in a week.</p> <p>Default value N/A</p>
scaling_adjustment	Yes	Integer	<p>Explanation Number of nodes that can be adjusted once.</p> <p>Constraints N/A</p> <p>Value range 1-100</p> <p>Default value N/A</p>
trigger	Yes	Trigger object	<p>Explanation Condition for triggering a rule. For details, see Table 7-52.</p> <p>Constraints N/A</p>

Table 7-52 Trigger

Parameter	Mandatory	Type	Description
metric_name	Yes	String	<p>Explanation Metric name. This triggering condition makes a judgment according to the value of the metric.</p> <p>Constraints N/A</p> <p>Value range The value contains a maximum of 64 characters.</p> <p>Default value N/A</p>
metric_value	Yes	String	<p>Explanation Metric threshold to trigger a rule The value must be an integer or a number with two decimal places.</p> <p>Constraints N/A</p> <p>Value range Only integers or numbers with two decimal places are allowed.</p> <p>Default value N/A</p>
comparison_operator	No	String	<p>Explanation Metric judgment logic operator.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • LT: less than • GT: greater than • LTOE: less than or equal to • GTOE: greater than or equal to <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
evaluation_periods	Yes	Integer	<p>Explanation Number of consecutive five-minute periods, during which a metric threshold is reached</p> <p>Constraints N/A</p> <p>Value range 1-200</p> <p>Default value N/A</p>

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.

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	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 7-55 Request body parameter

Parameter	Mandatory	Type	Description
tag	Yes	Tag object	<p>Explanation Cluster tag.</p> <p>Constraints N/A</p>

Table 7-56 tags parameters

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> The value can contain a maximum of 128 characters and cannot be an empty string. The tag key of a resource must be unique. A tag key can contain letters, digits, spaces, and special characters (_.:+=-@), but cannot start or end with a space or start with _sys_. <p>Default value N/A</p>
value	Yes	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> The value can contain a maximum of 255 characters and can be an empty string. A tag value can contain letters, digits, spaces, and special characters (_.:+=-@), but cannot start or end with a space or start with _sys_. <p>Default value N/A</p>

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.

URI

GET /v1.1/{project_id}/clusters/{cluster_id}/tags

Table 7-58 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 7-59 Response parameter

Parameter	Type	Description
tags	Array of TagPlain objects	Explanation Tag list. For details, see Table 7-60 .

Table 7-60 TagPlain parameters

Parameter	Type	Description
key	String	Explanation Tag key. Value range A tag key can contain letters, digits, spaces, and special characters (_.:=-+@), but cannot start or end with a space or start with _sys_.
value	String	Explanation Tag value. Value range A tag value can contain letters, digits, spaces, and special characters (_.:=-+@), but cannot start or end with a space or start with _sys_.

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.

URI

- Format
DELETE /v1.1/{project_id}/clusters/{cluster_id}/tags/{key}
- Parameter description

Table 7-62 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation Cluster ID For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range N/A</p> <p>Default value N/A</p>

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.

URI

POST /v1.1/{project_id}/clusters/{cluster_id}/tags/action

Table 7-63 URI parameters

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	<p>Explanation</p> <p>Cluster ID If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value</p> <p>N/A</p>

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Request Parameters

Table 7-64 Request parameters

Parameter	Mandatory	Type	Description
action	Yes	String	<p>Explanation Operation to be performed. The value can be create only.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> create: Create a tag. <p>Default value N/A</p>
tags	Yes	Array of tag objects	<p>Explanation Tag list. For details, see Table 7-65.</p> <p>Constraints N/A</p>

Table 7-65 tags parameters

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain a maximum of 128 characters and cannot be an empty string. • The tag key of a resource must be unique. • A tag key can contain letters, digits, spaces, and special characters (_.:+=-@), but cannot start or end with a space or start with _sys_. <p>Default value N/A</p>
value	Yes	String	<p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain a maximum of 255 characters and can be an empty string. • A tag value can contain letters, digits, spaces, and special characters (_.:+=-@), but cannot start or end with a space or start with _sys_. <p>Default value N/A</p>

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.

URI

```
POST /v1.1/{project_id}/clusters/{cluster_id}/tags/action
```

Table 7-66 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>
cluster_id	Yes	String	<p>Explanation Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value N/A</p>

Request Parameters

Table 7-67 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	<p>Explanation Operation to be performed. The value can be delete only.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> delete: Delete a tag. <p>Default value N/A</p>
tags	Yes	Array of tag objects	<p>Explanation The tag list. For details about the parameter, see Table 7-68.</p> <p>Constraints N/A</p>

Table 7-68 Tag

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Tag key.</p> <ul style="list-style-type: none"> The value can contain a maximum of 128 characters and cannot be an empty string. The tag key of a resource must be unique. A tag key can contain letters, digits, spaces, and special characters (<code>_</code>, <code>:=+@</code>), but cannot start or end with a space or start with <code>_sys_</code>.

Parameter	Mandatory	Type	Description
value	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain a maximum of 128 characters and cannot be an empty string. • The tag key of a resource must be unique. • A tag key can contain letters, digits, spaces, and special characters (_ := + - @), but cannot start or end with a space or start with <code>_sys_</code>. <p>Default value N/A</p> <p>Explanation Tag value.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • The value can contain a maximum of 255 characters and can be an empty string. • A tag value can contain letters, digits, spaces, and special characters (_ := + - @), but cannot start or end with a space or start with <code>_sys_</code>. <p>Default value N/A</p>

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.

URI

GET /v1.1/{project_id}/clusters/tags

Table 7-69 URI parameter

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 7-70 Response parameter

Parameter	Type	Description
tags	Array of TagWithMultiValue objects	Explanation Tag list. For details, see Table 7-71 .

Table 7-71 tags parameters

Parameter	Type	Description
key	String	Explanation Tag key. Constraints N/A Value range A tag key can contain letters, digits, spaces, and special characters ([:=+-@), but cannot start or end with a space or start with _sys_. Default value N/A
values	Array of strings	Explanation Tag value. Constraints A tag value can contain letters, digits, spaces, and special characters ([:=+-@), but cannot start or end with a space or start with _sys_.

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.

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	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Request Parameters

Table 7-73 Request body parameters

Parameter	Mandatory	Type	Description
tags	No	Array of TagWithMultiValue objects	<p>Explanation The returned result contains the resources corresponding to all tags in this parameter.</p> <p>Constraints The structure body is mandatory. A maximum of 10 tag keys are allowed in each query operation. The tag key cannot be left blank or set to the empty string. One tag key can have up to 10 tag values. For details about the parameters, see Table 7-74.</p>

Parameter	Mandatory	Type	Description
tags_any	No	Array of TagWithMultiValue objects	<p>Explanation</p> <p>The returned result contains the resources corresponding to any tag in this parameter.</p> <p>Constraints</p> <p>The structure body is mandatory. A maximum of 10 tag keys are allowed in each query operation. The tag key cannot be left blank or set to the empty string. One tag key can have up to 10 tag values. Keys must be unique and values of a key must be unique. For details about the parameters, see Table 7-74.</p>
not_tags	No	Array of TagWithMultiValue objects	<p>Explanation</p> <p>The returned result does not contain the resources corresponding to all tags in this parameter.</p> <p>Constraints</p> <p>The structure body is mandatory. A maximum of 10 tag keys are allowed in each query operation. The tag key cannot be left blank or set to the empty string. One tag key can have up to 10 tag values. Keys must be unique and values of a key must be unique. For details about the parameters, see Table 7-74.</p>

Parameter	Mandatory	Type	Description
not_tags_any	No	Array of TagWithMultiValue objects	<p>Explanation The returned result does not contain the resources corresponding to any tag in this parameter.</p> <p>Constraints The structure body is mandatory. A maximum of 10 tag keys are allowed in each query operation. The tag key cannot be left blank or set to the empty string. One tag key can have up to 10 tag values. Keys must be unique and values of a key must be unique. For details about the parameters, see Table 7-74.</p>
limit	No	Integer	<p>Explanation Number of query records.</p> <p>Constraints If you set action to count, this parameter is not required. If action is set to filter, the default value of this parameter is 1000.</p> <p>Value range 1-1000</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Explanation Index position. The query starts from the next data record indexed by this parameter. You do not need to specify this parameter when you query resources on the first page. When you query resources on subsequent pages, set this parameter to the location returned in the response body for the previous query.</p> <p>Constraints If you set action to count, this parameter is not required. If action is set to filter, the default value is 0.</p> <p>Value range ≥0</p> <p>Default value N/A</p>
action	Yes	String	<p>Explanation Operation ID (only filter and count is available). filter: The query is paginated. count: The total number of records is returned by search condition.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • filter: Perform pagination query. • count: Query the total number of resources. <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
matches	No	Array of Match objects	<p>Explanation Search field. key indicates the field to be matched, for example, resource_name. value indicates the matched value. This field is a fixed dictionary value. Determine whether fuzzy match is required based on different fields. For example, if key is resource_name, fuzzy search is used by default. If value is an empty string, exact match is used. For details about the parameters, see Table 7-75.</p> <p>Constraints N/A</p>

Table 7-74 TagWithMultiValue

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Explanation Tag key.</p> <p>Constraints N/A</p> <p>Value range A tag key can contain letters, digits, spaces, and special characters (_.:=-@), but cannot start or end with a space or start with <code>_sys_</code>.</p> <p>Default value N/A</p>
values	No	Array of strings	<p>Explanation Tag value.</p> <p>Constraints A tag value can contain letters, digits, spaces, and special characters (_.:=-@), but cannot start or end with a space or start with <code>_sys_</code>.</p>

Table 7-75 Match

Parameter	Mandatory	Type	Description
key	No	String	<p>Explanation Key. Currently, only resource_name is available, indicating the cluster name. Other keys will be supported later.</p> <p>Constraints N/A</p> <p>Value range The value can contain 1 to 64 characters.</p> <p>Default value N/A</p>
value	No	String	<p>Explanation Value. A value can contain a maximum of 64 Unicode characters.</p> <p>Constraints N/A</p> <p>Value range The value can contain 0 to 64 characters.</p> <p>Default value N/A</p>

Response Parameters

Table 7-76 Response parameters

Parameter	Type	Description
resources	Array of MRSResource objects	<p>Explanation Resource details. For details, see Table 7-77.</p>
total_count	Integer	<p>Explanation Total number of resources.</p> <p>Value range N/A</p>

Table 7-77 MRSResource

Parameter	Type	Description
resource_detail	String	Explanation Resource details. Value range N/A
resource_id	String	Explanation Resource ID. Value range N/A
resource_name	String	Explanation Resource name. Value range N/A
tags	Array of TagPlain objects	Explanation Tag list. For details, see Table 7-78 .

Table 7-78 TagPlain

Parameter	Type	Description
key	String	Explanation Tag key. Value range A tag key can contain letters, digits, spaces, and special characters (<code>_</code> , <code>:=+@</code>), but cannot start or end with a space or start with <code>_sys_</code> .
value	String	Explanation Tag value. Value range A tag value can contain letters, digits, spaces, and special characters (<code>_</code> , <code>:=+@</code>), but cannot start or end with a space or start with <code>_sys_</code> .

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

```

{
  "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	Explanation Region ID Constraints N/A Value range The value can contain 1 to 255 characters. Default value N/A

Table 7-80 Query parameters

Parameter	Mandatory	Type	Description
scope	No	String	<p>Explanation AZ scope.</p> <p>Constraints N/A</p> <p>Value range</p> <ul style="list-style-type: none"> • Center: central AZ • Edge: edge AZ <p>Default value N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 7-81 Response body parameter

Parameter	Type	Description
available_zones	Array of AvailableZoneV2 objects	<p>Explanation AZ list. For details about AvailableZoneV2, see Table 7-82.</p>
default_az_code	String	<p>Explanation Default AZ code</p> <p>Value range N/A</p>
support_physical_az_group	Boolean	<p>Explanation Whether physical AZ group is supported.</p> <p>Value range</p> <ul style="list-style-type: none"> • True: Physical availability groups are supported. • False: Physical availability groups are not supported.

Table 7-82 AvailableZoneV2

Parameter	Type	Description
id	String	Explanation ID Value range N/A
az_code	String	Explanation AZ code. Value range N/A
az_name	String	Explanation AZ name. Value range N/A
az_id	String	Explanation AZ ID. Value range N/A
status	String	Explanation AZ status Value range N/A
region_id	String	Explanation Region ID. Value range N/A
az_group_id	String	Explanation AZ group ID. Value range N/A

Parameter	Type	Description
az_type	String	<p>Explanation AZ type.</p> <p>Value range</p> <ul style="list-style-type: none"> • Core • Satellite • Dedicated • Virtual • Edge • EdgeCental: central and edge
az_tags	AvailableTag object	<p>Explanation AZ tag. For details about AvailableTag, see Table 7-83.</p>

Table 7-83 AvailableTag

Parameter	Type	Description
mode	String	<p>Explanation Resource mode, which can be dedicated or shared.</p> <p>Value range</p> <ul style="list-style-type: none"> • dedicated: dedicated AZ. • shared: shared AZ.
alias	String	<p>Explanation Alias of an AZ.</p> <p>Value range N/A</p>
public_border_group	String	<p>Explanation Group to which the AZ belongs. The default value is center.</p> <p>Value range N/A</p>

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.

URI

GET /v1.1/{project_id}/metadata/versions/{version_name}

Table 7-85 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation Project ID. For details about how to obtain the project ID, see Obtaining a Project ID.</p> <p>Constraints N/A</p> <p>Value range The value must consist of 1 to 64 characters. Only letters and digits are allowed.</p> <p>Default value N/A</p>

Parameter	Mandatory	Type	Description
version_name	Yes	String	<p>Explanation</p> <p>The cluster version, for example: MRS 3.3.1-LTS. If the request client does not support automatic escape, escape the space to %20, for example, MRS %203.3.1-LTS.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The value can contain 1 to 64 characters.</p> <p>Default value</p> <p>N/A</p>

Table 7-86 Query parameters

Parameter	Mandatory	Type	Description
cluster_id	No	String	<p>Explanation</p> <p>Cluster ID. If this parameter is specified, the latest metadata of the cluster that has been patched will be obtained. For details about how to obtain the cluster ID, see Obtaining a Cluster ID.</p> <p>Constraints</p> <p>N/A</p> <p>Value range</p> <p>The value can contain 1 to 64 characters, including only letters, digits, underscores (_), and hyphens (-).</p> <p>Default value</p> <p>N/A</p>

Request Parameters

None

Response Parameters

Status code: 200

Table 7-87 Response body parameter

Parameter	Type	Description
other	Map<String, Object>	Explanation Other
name	String	Explanation Image version name Value range N/A
template_type	String	Explanation Template type Value range N/A
image_id	String	Explanation Image ID. Value range N/A
status	String	Explanation Version status. Value range N/A
features	Array of strings	Explanation Feature list.
cluster_types	Array of strings	Explanation Cluster types.
version_type	String	Explanation Version type Value range N/A
components	Array of VersionComponent objects	Explanation Components. For details, see Table 7-88 .
resource_requirement	Array of strings	Explanation Description of required resources, such as IP addresses
constraints	VersionConstraint object	Explanation Version restrictions. For details about the parameters, see Table 7-90 .

Parameter	Type	Description
flavors	FlavorLists object	Explanation Specifications. For details, see Table 7-93 .
role_deploy_meta	Array of RoleDeployMeta objects	Explanation Role deployment policy of component instances. For details, see Table 7-94 .

Table 7-88 VersionComponent

Parameter	Type	Description
other	Map<String, Object>	Explanation Other
name	String	Explanation Component name. Value range N/A
version	String	Explanation Supported versions. Value range N/A
depend_on	Array of strings	Explanation Component dependencies.
description	String	Explanation Component description. Value range N/A
available_cluster_types	Array of strings	Explanation Cluster types that support the component.
external_data_sources	Array of ComponentExternalData source objects	Explanation External data source. For details, see Table 7-89 .
resource_requirement	Array of strings	Explanation Description of required resources, such as IP addresses.

Parameter	Type	Description
valid_roles	Array of strings	Explanation Valid roles.
visible	Boolean	Explanation Visible or not. Value range <ul style="list-style-type: none"> • True: visible • False: invisible
children_components	Array of strings	Explanation Sub-components
multi_az_support_status	String	Explanation Whether multi-AZ is supported. Value range N/A

Table 7-89 ComponentExternalDatasource

Parameter	Type	Description
name	String	Explanation Name of the external data source. Value range N/A
types	Array of strings	Explanation Type of the external data source.

Table 7-90 VersionConstraint

Parameter	Type	Description
other	Map<String, Object>	Explanation Other restrictions
node_constraint	NodeConstraints object	Explanation Node restrictions. For details, see Table 7-91 .
safe_mode_kerberos_exclude_components	Array of strings	Explanation Components excluded from clusters in security mode (Kerberos authentication enabled)

Table 7-91 NodeConstraints

Parameter	Type	Description
other	Map<String, Object>	Explanation Restrictions on other nodes, including the number of nodes and disks
master	NodeConstraint object	Explanation Restrictions on Master nodes, including the number of nodes and disks. For details, see Table 7-92 .
core	NodeConstraint object	Explanation Restrictions on Core nodes, including the number of nodes and disks. For details, see Table 7-92 .
task	NodeConstraint object	Explanation Restrictions on Task nodes, including the number of nodes and disks. For details, see Table 7-92 .
core_separate	NodeConstraint object	Explanation Core node restrictions when management and control nodes are different nodes, including the number of nodes and disks. For details, see Table 7-92 .
core_combine	NodeConstraint object	Explanation Core node restrictions when management and control nodes are the same nodes, including the number of nodes and disks. For details, see Table 7-92 .
task_separate	NodeConstraint object	Explanation Task node restrictions when management and control nodes are different nodes, including the number of nodes and disks. For details, see Table 7-92 .
task_combine	NodeConstraint object	Explanation Task node restrictions when management and control nodes are the same nodes, including the number of nodes and disks. For details, see Table 7-92 .

Parameter	Type	Description
node_group_task	NodeConstraint object	Explanation Restrictions on task node groups, including the number of nodes and disks. For details, see Table 7-92 .

Table 7-92 NodeConstraint

Parameter	Type	Description
other	Map<String, Object>	Explanation Other restrictions.
min_node_num	Integer	Explanation Minimum number of nodes. Value range N/A
max_node_num	Integer	Explanation Maximum number of nodes. Value range N/A
min_core_num	Map<String, Integer>	Explanation Minimum number of cores.
min_mem_size	Map<String, Integer>	Explanation Minimum memory.
min_disk_size	Integer	Explanation Minimum disk capacity. Value range N/A
max_node_group_num	Integer	Explanation Maximum number of nodes. Value range N/A
min_data_volume_total_size	Map<String, Integer>	Explanation Minimum data volume capacity.
disk_type_constraint	Map<String, String>	Explanation Disk type restrictions, including the types supported by the current node group.

Parameter	Type	Description
min_root_disk_size	Integer	Explanation Minimum system disk size. Value range N/A

Table 7-93 FlavorLists

Parameter	Type	Description
master	Array of strings	Explanation Specifications supported by Master nodes.
core	Array of strings	Explanation Specifications supported by Core nodes.
task	Array of strings	Explanation Specifications supported by Task nodes.

Table 7-94 RoleDeployMeta

Parameter	Type	Description
other	Map<String, Object>	Explanation Other properties.
name	String	Explanation Role name. Value range N/A
code_name	String	Explanation Short name of the role. Value range N/A
component	String	Explanation Component to which the role belongs. Value range N/A

Parameter	Type	Description
node_preference	String	Explanation Node deployment preference. Value range N/A
count	String	Explanation Maximum number of roles. Value range N/A
affinity	String	Explanation Affinity. Value range N/A
affinity_target	String	Explanation Target in the affinity relationship. Value range N/A
multi_instance	Integer	Explanation Multi-instance. Value range N/A
role_kind	String	Explanation Role type. Value range N/A
constraints	Array of strings	Explanation Role restrictions, including function restrictions of the current component role, for example, "no_scale_in".
multi_az_placement	String	Explanation Multi-AZ deployment. Value range N/A

Parameter	Type	Description
arbitration_deploy ment	Boolean	<p>Explanation Whether to use arbitration deployment.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Use arbitration deployment. • false: Do not use arbitration deployment.
support_elb	Boolean	<p>Explanation Whether ELB is supported.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: ELB is supported. • false: ELB is not supported.
multi_affinity_group_enable	Boolean	<p>Explanation Whether to enable the multi-affinity groups.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable multi-affinity groups. • false: Disable multi-affinity groups.
local_disks_anti_affinity	Boolean	<p>Explanation Whether to enable local disk anti-affinity.</p> <p>Value range</p> <ul style="list-style-type: none"> • true: Enable local disk anti-affinity. • false: Disable local disk anti-affinity.
multi_instance_name_pattern	String	<p>Explanation Multi-instance name pattern.</p> <p>Value range N/A</p>
private_ip	String	<p>Explanation Private IP addresses.</p> <p>Value range N/A</p>
weight	String	<p>Explanation Weight.</p> <p>Value range N/A</p>

Example Request

Obtain the metadata of MRS 3.3.1-LTS.

```
GET /v1.1/{project_id}/metadata/versions/MRS%203.3.1-LTS
```

Example Response

Status code: 200

Version metadata details. Some contents are omitted in this example. For details, call the API.

```
{
  "other" : { },
  "name" : "MRS 3.3.1-LTS",
  "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", "omm_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", "RESTServer",
"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 Obtaining a Project ID .

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"> • 1: MapReduce • 2: Spark • 3: Hive Script • 4: HiveQL (not supported currently) • 5: DistCp, importing and exporting data. For details, see Table 8-3. • 6: Spark Script • 7: Spark SQL, submitting Spark SQL statements. For details, see Table 8-4. (Not supported in this API currently.) <p>NOTE 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>NOTE 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"> • Contains a maximum of 1,023 characters, excluding special characters such as ; &><'\$. The address cannot be empty or full of spaces. • Starts with / or s3a://. The OBS path does not support files or programs encrypted by KMS. • Spark Script must end with .sql while MapReduce and Spark Jar must end with .jar.sql and jar are case-insensitive.
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 ; &><'\$, and can be left blank.</p> <p>NOTE 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 @password=yyy</p>

Parameter	Mandatory	Type	Description
input	No	String	<p>Path for inputting data, which must start with / or s3a://. 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 ; &>'<\$, and can be left blank.</p>
output	No	String	<p>Path for outputting data, which must start with / or s3a://. 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 ; &>'<\$, 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 s3a://. A correct OBS path is required.</p> <p>The parameter contains a maximum of 1,023 characters, excluding special characters such as ; &>'<\$, 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"> • Contains a maximum of 1,023 characters, excluding special characters such as ; &>'<\$. The address cannot be empty or full of spaces. • The path must start with / or s3a://. The OBS path does not support files or programs encrypted by KMS. • The path must end with .sql.sql is case-insensitive.

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>NOTE Identical job names are allowed but not recommended.</p>
input	No	String	<p>Data source path</p> <ul style="list-style-type: none"> • When you import data, the parameter is set to an OBS path. Files or programs encrypted by KMS are not supported. • When you export data, the parameter is set to an HDFS path.
output	No	String	<p>Data receiving path</p> <ul style="list-style-type: none"> • When you import data, the parameter is set to an HDFS path. • When you export data, the parameter is set to an OBS path.
file_action	Yes	String	<p>Types of file operations, including:</p> <ul style="list-style-type: none"> • export: Export data from HDFS to OBS. • import: Import data from OBS to HDFS.

Table 8-4 Spark SQL parameters

Parameter	Mandatory	Type	Description
hql	Yes	String	<p>Spark SQL statement, which needs Base64 encoding and decoding. ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789+/- is a standard encoding table. MRS uses ABCDEFGHIJKLMN OPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789+/- for Base64 encoding. The value of the hql 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"> Obtain the Base64 encoding tool. Enter the show tables; Spark SQL statement in the encoding tool to perform Base64 encoding. Obtain the encoded character string c2hvdYB0YWlsZXM7. At the beginning of c2hvdYB0YWlsZXM7, add any letter, for example, g. Then, the character string becomes gc2hvdYB0YWlsZXM7, that is, the value of the hql parameter.
job_name	Yes	String	<p>Job name. It contains 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> <p>NOTE 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 Table 8-6 .

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 Obtaining a Project ID .
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"> • true • false 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"> • true • false The current version does not support this function.
group_id	String	Group ID of a job
jar_path	String	Path of the .jar 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"> • 1: MapReduce • 2: Spark • 3: Hive Script • 4: HiveQL (not supported currently) • 5: DistCp • 6: Spark Script • 7: Spark SQL (not supported in this API currently)
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"> • -1: Terminated • 1: Starting • 2: Running • 3: Completed • 4: Abnormal • 5: Error

Parameter	Type	Description
job_final_status	Integer	Final job status <ul style="list-style-type: none"> • 0: unfinished • 1: terminated due to an execution error • 2: executed successfully • 3: canceled
hive_script_path	String	Address of the Hive script
create_by	String	User ID for creating jobs This parameter is not used in the current version, but is retained for compatibility with earlier versions.
finished_step	Integer	Number of completed steps 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 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 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. 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 This parameter is not used in the current version, but is retained for compatibility with earlier versions.
step_num	Integer	Number of steps This parameter is not used in the current version, but is retained for compatibility with earlier versions.
task_num	Integer	Number of tasks 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 The current version does not support this function.
user_id	String	User ID for creating jobs 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 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 Obtaining a Project ID .

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 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"> • -1: Terminated • 2: Running • 3: Completed • 4: Abnormal

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 For details, see Table 8-11 .

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 Obtaining a Project ID .
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 .jar file or .sql 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"> • 1: MapReduce • 2: Spark • 3: Hive Script • 4: HiveQL (not supported currently) • 5: DistCp • 6: Spark Script • 7: Spark SQL (not supported in this API currently)
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"> • -1: Terminated • 2: Running • 3: Completed • 4: Abnormal
job_final_status	Integer	Final job status. <ul style="list-style-type: none"> • 0: unfinished • 1: terminated due to an execution error • 2: executed successfully • 3: canceled
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¤t_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 Obtaining a Project ID .
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 Table 8-15 .

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 Obtaining a Project ID .

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 .jar file or .sql 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"> • 1: MapReduce • 2: Spark • 3: Hive Script • 4: HiveQL (not supported currently) • 5: DistCp • 6: Spark Script • 7: Spark SQL (not supported in this API currently)
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"> • -1: Terminated • 1: Starting • 2: Running • 3: Completed • 4: Abnormal • 5: Error

Parameter	Type	Description
job_final_status	Integer	Final job status <ul style="list-style-type: none"> 0: unfinished 1: terminated due to an execution error 2: executed successfully 3: canceled
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 Obtaining a Project ID .
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 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	√	√

10 Appendix

10.1 Status Codes

[Table 10-1](#) describes status codes.

Table 10-1 Status codes

Status Code	Message	Description
100	Continue	The client should continue with its request. 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. For example, the current HTTPS protocol is switched to a later version.
200	OK	The request has been fulfilled.
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.

Status Code	Message	Description
204	NoContent	The request has been fulfilled, but the HTTPS response does not contain a response body. 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. 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.
306	Unused	The HTTPS status code is no longer used.
400	BadRequest	The request is invalid. The client should not repeat the request without modifications.

Status Code	Message	Description
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. The client should not repeat the request without modifications.
404	NotFound	The requested resource cannot be found. The client should not repeat the request without modifications.
405	MethodNotAllowed	A request method is not supported for the requested resource. 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. The client may repeat the request without modifications at a later time.
409	Conflict	The request could not be processed due to a conflict with the current state of the resource. 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.

Status Code	Message	Description
411	Length Required	The server is refusing to process the request without a defined Content-Length .
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 Retry-After 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 Expect request-header field.
422	UnprocessableEntity	The request is well-formed but is unable to be processed due to semantic errors.
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 Retry-After 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.

Status Code	Message	Description
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 Timeout parameter is specified in the request.
505	HTTPS Version not supported	The server does not support the HTTPS protocol version used in the request.

10.2 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 Running .

Status Code	Error Code	Error Message	Description	Measure
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. 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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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 Dashboard 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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
400	12000009	The parameter is invalid.	Invalid parameter.	Check whether the parameters are correct according to the API document.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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 Running .
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 Running .
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.

Status Code	Error Code	Error Message	Description	Measure
400	120000213	The region [xxx] does not exist.	The specified region [xxx] does not exist.	Modify the request parameters based on the error message.
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	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 Running .
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	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, mains cannot be empty.	Modify the request parameters based on the error message.
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, libs cannot be empty.	Modify the request parameters based on the error message.

Status Code	Error Code	Error Message	Description	Measure
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: 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(`~!@#\$\$% %^&*()-_+= [{}];',<.>?). 3.cannot be the username or the username spelled backwards.	The password xxx must: - Contain xxx to xxx characters. - Contain at least xxx of the following character types: uppercase letters, lowercase letters, digits, and special character `~!@#\$\$% %^&*()-_+= [{}];',<.>? - 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 jar_path parameter cannot be empty.	Added the jar_path 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 Running .
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 s3a:// or s3n:// . Use a file path supported by the cluster, such as obs:// .	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. 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. 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 Running or Abnormal .
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 Stopped .

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 hive_script_path 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.
400	12000122	EPS service error.	EPS error.	Try to call the API again.

Status Code	Error Code	Error Message	Description	Measure
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 Running .
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 Running state.	Wait until the cluster task is complete and the status changes to Running .

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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 Deleting or Deleted state cannot be changed.	Wait until the cluster task is complete and the status changes to Running .
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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 Running .
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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 local_disks_anti_affinity is true : 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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
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.

Status Code	Error Code	Error Message	Description	Measure
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.
400	MRS.00005096	The role metadata is abnormal.	Role metadata is abnormal.	Initiate the call again or contact technical support.
400	MRS.00005097	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.	<p>Invalid connector name. The name can contain only uppercase and lowercase letters, digits, hyphens (-), and underscores(_). The connector ID cannot be empty.</p> <p>Unsupported data type: xxx.</p> <p>Unsupported component type: xxx. The data connection (ID: xxx) is not found.</p> <p>The mapping between the cluster (ID: xxx) and the data connection is not found.</p> <p>The mapping between the cluster (connection ID: xxx) and the data connection mapping is not found.</p> <p>The RDS instance ID cannot be left blank. The instance ID, database name, username,</p>	Locate and rectify the fault based on the error message.

Status Code	Error Code	Error Message	Description	Measure
			<p>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</p>	

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.

10.3 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 Settings** from the drop-down list.

3. On the **My Settings** 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.

Obtaining a Project ID by Calling an API

You can obtain the project ID by calling the IAM 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"
  }
}
```

10.4 Obtaining Tenant ID

An account ID (**domain-id**) is required for some URLs when an API is called. To obtain a tenant ID on the management console, perform the following steps:

1. Log in to the management console.
2. Click the username and choose **My Settings** from the drop-down list.
3. On the **My Settings** page, view the tenant ID in the project list.

10.5 Obtaining the MRS Cluster Information

Components Supported by MRS

MRS 3.3.1-LTS supports the following components:

- An analysis cluster contains the following components: Hadoop, Spark, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, JobGateway, Guardian, and Doris.

- A streaming cluster contains the following components: Kafka, Flume, ZooKeeper, and Ranger.
- A hybrid cluster contains the following components: Hadoop, Spark, HBase, Hive, Hue, Loader, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, Kafka, Flume, JobGateway, Guardian, and Doris.
- A custom cluster contains the following components: CDL, Hadoop, Spark, HBase, Hive, Hue, IoTDB, Loader, Kafka, Flume, Flink, Oozie, ZooKeeper, HetuEngine, Ranger, Tez, and ClickHouse, Guardian, JobGateway, Doris and MemArtsCC.

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.

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.

10.6 Roles and components supported by MRS

Table 10-2 Roles and components supported by MRS

Role Name	Component
OMSServer	OMSServer
NameNode	HDFS
Zkfc	HDFS
JournalNode	HDFS
DataNode	HDFS
ResourceManager	Yarn
NodeManager	Yarn

Role Name	Component
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
Broker	Kafka
Supervisor	Storm
Logviewer	Storm
Nimbus	Storm

Role Name	Component
UI	Storm
FlinkResource	Flink
ClickHouseServer	ClickHouse
ClickHouseBalancer	ClickHouse
HSBroker	HetuEngine
HSConsole	HetuEngine
QAS	HetuEngine
CDLConnector	CDL
CDLService	CDL
IoTDBServer	IoTDB
ConfigNode	IoTDB
FE	Doris
BE	Doris
DBroker	Doris
TokenServer	Guardian
JobServer	JobGateway
JobBalancer	JobGateway

10.7 Change History

Release Date	What's New
2024-11-30	This issue is the first official release.