

Cloud Search Service

API Reference

Issue	01
Date	2025-11-28



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

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

- 1 Before You Start..... 1**
- 2 API Overview..... 4**
- 3 Calling APIs..... 11**
 - 3.1 Making an API Request..... 11
 - 3.2 Authentication..... 13
 - 3.3 Returned Values..... 14
- 4 API V2..... 15**
 - 4.1 Cluster Management..... 15
 - 4.1.1 Creating a Cluster (V2)..... 15
 - 4.1.2 Restarting a Cluster (V2)..... 50
 - 4.1.3 Rolling Restart (Deprecated)..... 53
 - 4.1.4 Rolling Restart..... 56
 - 4.2 Snapshot Management..... 59
 - 4.2.1 Enabling Automatic Snapshot Creation..... 59
 - 4.2.2 Disabling Automatic Snapshot Creation..... 64
 - 4.3 Logstash..... 66
 - 4.3.1 Delete a Configuration File (V2)..... 66
 - 4.3.2 Delete a Custom Template V2..... 68
- 5 API V1..... 71**
 - 5.1 Cluster Management..... 71
 - 5.1.1 Creating a Cluster (Discarded)..... 71
 - 5.1.2 Querying the Cluster List..... 92
 - 5.1.3 Querying Cluster Details..... 108
 - 5.1.4 Deleting a Cluster..... 124
 - 5.1.5 Modifying the Name or Description of a Cluster..... 126
 - 5.1.6 Changing the Cluster Administrator Password..... 128
 - 5.1.7 Restarting a Cluster (Deprecated)..... 131
 - 5.1.8 Query the disk usage of a cluster..... 133
 - 5.1.9 Downloading a Security Certificate..... 136
 - 5.1.10 Upgrading a Cluster Kernel..... 137
 - 5.1.11 Obtaining the ID of the Target Image to Upgrade To..... 144
 - 5.1.12 Obtaining Details About a Cluster Upgrade or AZ Switchover..... 149

5.1.13 Retrying or Terminating a Failed Upgrade or AZ Switchover Task.....	158
5.1.14 Changing the Security Group.....	161
5.1.15 Changing the Subnet of a Cluster.....	163
5.1.16 Automatically create an agency.....	166
5.2 Specifications Query.....	169
5.2.1 Obtaining the Instance Specifications List.....	169
5.2.2 Querying Available Node Specifications for a Cluster.....	173
5.2.3 Querying Details About a Flavor.....	178
5.2.4 Query available node specifications by engine.....	182
5.3 Tag Management.....	189
5.3.1 Querying All Tags.....	189
5.3.2 Querying Tags of a Specified Cluster.....	191
5.3.3 Adding Tags to a Cluster.....	194
5.3.4 Deleting a Cluster Tag.....	198
5.3.5 Adding or Deleting Cluster Tags in Batches.....	200
5.4 Modify Configuration.....	204
5.4.1 Adding Instances and Expanding Instance Storage Capacity.....	205
5.4.2 Adding Dedicated Master or Client Nodes or Cold Data Nodes.....	211
5.4.3 Changing Node Specifications (Discarded).....	216
5.4.4 Scaling In a Cluster by Removing a Specific Node.....	221
5.4.5 Change the specifications of a specified node type.....	225
5.4.6 Scaling In Nodes of a Specific Type.....	231
5.4.7 Replacing a Node.....	236
5.4.8 Changing the Security Mode.....	239
5.4.9 Scaling Out a Cluster (Discarded).....	243
5.4.10 Changing AZs for a Cluster Instance.....	246
5.5 Word Dictionary Management.....	251
5.5.1 Querying the Status of a Custom Word Dictionary.....	251
5.5.2 Loading Custom Word Dictionaries.....	255
5.5.3 Disabling a Word Dictionary.....	264
5.6 Kibana/Dashboards Public Network Access.....	266
5.6.1 Enabling Kibana/Dashboards Public Network Access.....	266
5.6.2 Disabling Kibana/Dashboards Public Network Access.....	269
5.6.3 Modify Kibana/Dashboards Public Network Bandwidth.....	272
5.6.4 Modifying Kibana/Dashboards Public Network Access.....	274
5.6.5 Disabling Kibana/Dashboards Public Network Access Control.....	276
5.7 Log Management.....	278
5.7.1 Enabling Log Backup or Log Collection.....	278
5.7.2 Disabling the Log Backup or Collection Function.....	285
5.7.3 Querying the Log Backup Task List.....	288
5.7.4 Querying Basic Log Configurations.....	293
5.7.5 Modifying Basic Settings for Log Backup or Log Collection.....	300

5.7.6 Enabling the Automatic Log Backup Policy.....	305
5.7.7 Disabling the Automatic Log Backup Policy.....	308
5.7.8 Backing Up Logs.....	310
5.7.9 Searching for Logs.....	312
5.7.10 Log Ingestion Connectivity Testing.....	320
5.8 Public Network Access.....	322
5.8.1 Enabling Public Network Access.....	323
5.8.2 Disabling Public Network Access.....	326
5.8.3 Modifying Public Network Access Bandwidth.....	329
5.8.4 Enabling the Public Network Access Control Whitelist.....	332
5.8.5 Disabling the Public Network Access Control Whitelist.....	334
5.9 Snapshot Management.....	336
5.9.1 Enabling or Modifying the Cluster Snapshot Function.....	336
5.9.2 Manually Creating a Snapshot.....	348
5.9.3 Restoring a Snapshot.....	353
5.9.4 Deleting a Snapshot.....	359
5.9.5 Querying the Automatic Snapshot Creation Policy.....	361
5.9.6 Configuring the Automatic Snapshot Creation Policy.....	369
5.9.7 Querying a Snapshot List.....	375
5.9.8 Disabling Cluster Snapshots.....	381
5.9.9 (Not Recommended) Automatically Setting Basic Configurations of a Cluster Snapshot.....	383
5.10 VPC Endpoint.....	385
5.10.1 Enabling the VPC Endpoint Service.....	385
5.10.2 Disabling the VPC Endpoint Service.....	388
5.10.3 Obtaining an Endpoint Connection.....	390
5.10.4 Updating an Endpoint Connection.....	395
5.10.5 Modifying the VPCEP Service Whitelist of a Cluster.....	398
5.11 Parameter Configuration.....	400
5.11.1 Modifying Cluster Parameters.....	400
5.11.2 Obtaining the Task List of Parameter Configurations.....	405
5.11.3 Obtaining the Parameter Configuration List.....	409
5.12 Intelligent O&M.....	412
5.12.1 Obtaining the Intelligent O&M Task List and Details.....	412
5.12.2 Creating a Cluster Detection Task.....	419
5.12.3 Deleting a Detection Task.....	423
5.12.4 Obtaining SMN Topics Available for Intelligent O&M Alarms.....	425
5.12.5 Viewing Intelligent O&M Check Items.....	427
5.12.6 Viewing Scheduled Check Settings for Intelligent O&M.....	429
5.12.7 Updating Scheduled Check Settings for Intelligent O&M.....	432
5.12.8 Disabling Scheduled Checks for Intelligent O&M.....	435
5.13 Cluster management.....	437
5.13.1 Updating Cluster Routes.....	437

5.13.2 Obtaining a Cluster Route..... 440

5.14 Load Balancing..... 444

5.14.1 Querying Dedicated Load Balancers Supported by a Cluster..... 444

5.14.2 Enabling or Disabling Load Balancers for a Cluster..... 447

5.14.3 Obtain information about the load balancers of a cluster..... 450

5.14.4 Configuring Load Balancing Listeners for a Cluster..... 459

5.14.5 Updating Load Balancing Listeners for a Cluster..... 462

5.14.6 Querying the Certificate List..... 466

5.15 Node Specifications Query..... 469

5.15.1 Listing Disk Types..... 469

6 Examples..... 473

6.1 Creating a Pay-per-Use Elasticsearch Cluster..... 473

7 Appendixes..... 476

7.1 Status Codes..... 476

7.2 Error Codes..... 479

7.3 Obtaining a Project ID and Name..... 503

7.4 Obtaining the Cluster ID..... 504

1 Before You Start

Overview

Welcome to *Cloud Search Service (CSS) API Reference*. CSS is a fully managed, distributed search service that enables you to perform quick, real-time search on both structured and unstructured data. It is fully compatible with open-source Elasticsearch and provides you search, statistical analysis, and reporting capabilities.

This document provides CSS API description, syntax, parameters, and examples. CSS provides APIs for cluster management and snapshot management. You can call these APIs to easily create, query, delete, restart, and scale clusters.

Supported versions

In CSS, Elasticsearch 7.1.1, 7.3.2, 7.6.2, 7.9.3, and 7.10.2 are supported.

Restrictions and Limitations

- The number of CSS clusters you can create is determined by your quota. To view or increase that quota, see [Quotas](#).
- For more constraints, see the API descriptions.

Regions and Endpoints

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

Concepts

- **Account**
An account is created upon successful registration with the cloud system. The account has full access permissions for all of its cloud services and resources. It can be used to reset user passwords and grant user permissions. The account is a payment entity, which should not be used directly to perform routine management. For security purposes, create Identity and Access

Management (IAM) users and grant them permissions for routine management.

- User

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

API authentication requires information such as the account name, username, and password.

- Region

A region is a geographic area in which cloud resources are deployed.

Availability zones (AZs) in the same region can communicate with each other over an intranet, while AZs in different regions are isolated from each other.

By creating cloud resources in different regions, you can design applications to better meet customer requirements and comply with local laws and regulations.


- AZ

An AZ comprises of one or multiple physical data centers equipped with independent ventilation, fire, water, and electricity facilities. Computing, network, storage, and other resources in an AZ are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to allow you to build cross-AZ high-availability systems.

- Project

A project corresponds to a region. Default projects are defined to group and physically isolate resources (including compute, storage, and network resources) between different regions. Users can be granted permissions in a default project to access all resources under their accounts in the region associated with the project. If you need more refined access control, create sub-projects under a default project and purchase resources in sub-projects. Then you can assign users the permissions required to access only the resources in the specific sub-projects.

Figure 1-1 Isolation by project



- Checkpoint: When an application consumes data, the latest SN of the consumed data is recorded as a checkpoint. When the data is consumed again, the consumption can be continued based on this checkpoint.

- APP: Multiple applications can access data in the same stream. Checkpoints generated for each application are used to record the consumed data in the stream by each application.

2 API Overview

CSS APIs comply with RESTful API design specifications. You can use the functions described in [Table 2-2](#) by calling the corresponding APIs.

 CAUTION

Some APIs of V2 have been released. You are advised to use APIs of V2 preferentially.

APIs of V2

Table 2-1 Description of APIs of version V2

API	Description
Creating a Cluster (V2)	This API is used to create a cluster with multiple node types (for example, ess , ess-cold , ess-client , and ess-master).
Restarting a Cluster (V2)	This API is used to restart nodes of all node types or the combination of some node types in the current cluster.
Rolling Start	This API is used to restart nodes one by one, which requires a long time if the nodes have a large number of indexes.
Enabling Automatic Snapshot Creation	This API is used to enable automatic snapshot creation.
Disabling Automatic Snapshot Creation	This API is used to disable automatic snapshot creation.
Deleting a Configuration File (V2)	This API is used to delete configuration files.

API	Description
Deleting a Custom Template (V2)	This API is used to delete custom templates.

APIs of V1

Table 2-2 Description of APIs of version V1

API	Description
Cluster Management APIs	Create clusters, scale out clusters, view cluster details, and obtain instance specifications.
Kibana Public Network Access APIs	Enable, disable, and modify Kibana public network access for a security-mode cluster.
Log Management APIs	Enable, disable, modify, and query logs.
Public Network Access APIs	Enable, disable, and modify public network access for a cluster.
Snapshot Management APIs	Create, restore, and delete snapshots. Snapshots are a convenient way to back up and restore data.
VPC Endpoint APIs	Enable, disable, obtain, and update VPC endpoint connections.
Parameter Configuration APIs	Modify the parameter settings of a cluster.
Load Balancer APIs	Enable or disable load balancers, or obtain load balancer information.

Table 2-3 Cluster management APIs

API	Description
Creating a Cluster	This API is used to create a cluster.
Querying the List of Clusters	This API is used to query and display the cluster list and cluster status.

API	Description
Querying Cluster Details	This API is used to query and display the details of a cluster.
Deleting a Cluster	This API is used to delete a cluster.
Changing the Cluster Name	This API is used to change the name of a cluster.
Changing the Cluster Password	This API is used to change the password of a cluster.
Updating Cluster Routes	This API is used to update cluster routes.
Obtaining a Cluster Route	This API is used to obtain a cluster route.
Restarting a Cluster	This API is used to restart a cluster. Services will be interrupted during the restart.
Downloading a Security Certificate	This API is used to download a security certificate.
Upgrading a Cluster Kernel	This API is used to upgrade Elasticsearch from an earlier version to a later version or the same version.
Obtaining the ID of the Target Image to Upgrade To	This API is used to obtain the ID of an image that the current cluster can be upgraded to.
Obtaining Details About a Cluster Upgrade or AZ Switchover	An upgrade may take a long time. This API can be used to check phase-by-phase information during a cluster upgrade or AZ switchover.
Retrying a Failed Upgrade Task	If an upgrade takes too long, there is a chance that it has already failed due to connectivity issues. You can call this API to retry the upgrade or terminate its impact.
Changing the Subnet of a Cluster	This API is used to bind a new subnet to newly added nodes after cluster creation, or in the case of a cluster subnet change.
Changing the AZ of a Cluster Instance	This API is used to change AZs by specifying node types.

Table 2-4 Kibana public network access APIs

API	Description
Enabling Kibana Public Access	This API is used to enable public network access to Kibana.

API	Description
Disabling Kibana Public Access	This API is used to disable public network access to Kibana.
Modifying the Kibana Public Network Bandwidth	This API is used to modify the public network bandwidth of Kibana.
Modifying Kibana Public Network Access Control	This API is used to modify the Kibana public network access whitelist.
Disabling Kibana Public Network Access Control	This API is used to disable the Kibana public network access whitelist.

Table 2-5 Log management APIs

API	Description
Enabling the Log Function	This API is used to enable the log management function.
Disabling Logging	This API is used to disable the log management function.
Querying the Job List	This API is used to query the job list in the logs of a cluster.
Querying Basic Log Configurations	This API is used to query basic log configurations.
Modifying Basic Log Configurations	This API is used to modify basic log configurations.
Enabling an Automatic Log Backup Policy	This API is used to enable an automatic log backup policy.
Disabling an Automatic Log Backup Policy	This API is used to disable an automatic log backup policy.
Backing Up Logs	This API is used to back up logs.
Querying Logs	This API is used to query log information.
Testing Connectivity	This API is used to test the connectivity between two clusters.

Table 2-6 Public network access APIs

API	Description
Enabling Public Network Access	This API is used to enable public network access.
Disabling Public Network Access	This API is used to disable public network access.
Modifying Public Network Access Bandwidth	This API is used to modify the public network access bandwidth.
Enabling the Public Network Access Control Whitelist	This API is used to enable the public network access control whitelist.
Disabling the Public Network Access Control Whitelist	This API is used to disable the public network access control whitelist.

Table 2-7 Snapshot management APIs

API	Description
(Not Recommended) Automatically Configuring Basic Settings of a Cluster Snapshot	This API is used to automatically configure basic settings of a cluster snapshot, including configuring OBS buckets and IAM agency.
Modifying Basic Configurations of a Cluster Snapshot	This API is used to modify the basic configurations of a cluster snapshot, including OBS buckets and IAM agency.
Manually Creating a Snapshot	This API is used to manually create a snapshot.
Restoring a Snapshot	This API is used to manually restore a snapshot.
Deleting a Snapshot	This API is used to delete a snapshot.
Setting the Automatic Snapshot Creation Policy	This API is used to configure automatic snapshot creation. By default, a snapshot is created every day.
Querying the Automatic Snapshot Creation Policy	This API is used to query the automatic snapshot creation policy of a cluster.
Querying the Snapshot List	This API is used to query the snapshot list of a specified cluster.

API	Description
Disabling the Snapshot Function	This API is used to disable the snapshot function.

Table 2-8 VPC endpoint APIs

API	Description
Enabling the VPC Endpoint Service	This API is used to enable the VPC endpoint service.
Disabling the VPC Endpoint Service	This API is used to disable the VPC endpoint service.
Obtaining a VPC Endpoint Connection	This API is used to obtain a VPC endpoint connection.
Updating a VPC Endpoint Connection	This API is used to update a VPC endpoint connection.
Modifying the Endpoint Service Whitelist	This API is used to modify the endpoint service whitelist.

Table 2-9 Parameter configuration APIs

API	Description
Modifying Parameter Settings	This API is used to modify parameters.
Obtaining the Task List of Parameter Configurations	This API is used to obtain the task list of parameter configurations.
Obtaining the Parameter Configuration List	This API is used to obtain the parameter configuration list of the current cluster.

Table 2-10 Load balancer APIs

API	Description
Querying ELB V3 Load Balancers Supported by a Cluster	This API is used to query the ELB V3 load balancers supported by a cluster.

API	Description
Enabling or Disabling the Elasticsearch Load Balancer	This API is used to enable or disable an Elasticsearch load balancer.
Configuring an Elasticsearch Listener	This API is used to configure an Elasticsearch listener.
Obtaining Elasticsearch ELB Information and Health Check Status	This API is used to obtain the ELB information and health check status of an Elasticsearch cluster.
Updating an Elasticsearch Listener	This API is used to update an Elasticsearch listener.
Querying the Certificate List	This API is used to query the certificate list.

3 Calling APIs

3.1 Making an API Request

This section describes the structure of a RESTful API request, and uses the IAM API for as an example to describe how to call an API.

Request URI

A request URI is in the following format:

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

Table 3-1 Request URL

Parameter	Description
URI-scheme	Protocol used to transmit requests. All APIs use HTTPS .
Endpoint	Domain name or IP address of the server running the REST service. The endpoint varies between services in different regions. It can be obtained from Regions and Endpoints .
resource-path	API access path for performing a specified operation. Obtain the value from the URI of an API. .
query-string	Query parameter, which is optional. Ensure that a question mark (?) is included before a query parameter that is in the format of " Parameter name=Parameter value ". For example, limit=10 indicates that a maximum of 10 pieces of data is to be viewed.

NOTE

To simplify the URI display, each API is provided with only 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

HTTP-based request methods, which are also called operations or actions, specify the type of operations that you are requesting.

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

For example, in the URI of the API for , the request method is **POST**. The request is as follows:

Request Header

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

[Table 3-2](#) lists common request header fields.

Table 3-2 Common request headers

Parameter	Mandatory	Description
Content-Type	Yes	Message body type (or format). You are advised to use the default value application/json .
X-Auth-Token	Mandatory for token authentication	User token. It is the response to the API for obtaining a user token (only this API does not require authentication). After the request is processed, the value of X-Subject-Token in the response header (Header) is the token value.
X-Project-Id	No	Subproject ID, which is used in multi-project scenarios. The X-Project-ID field is mandatory in the request header for accessing resources in a subproject through AK/SK-based authentication.

Parameter	Mandatory	Description
X-Sdk-Date	Mandatory for AK/SK authentication	Request sending time. When AK/SK authentication is enabled, this field is automatically specified when SDK is used to sign the request. For details, see Authentication . The format is YYYYMMDD'T'HHMMSS'Z'. The value is the current GMT time of the system.
Authorization	Mandatory for AK/SK authentication	Signature authentication information. When AK/SK authentication is enabled, this field is automatically specified when SDK is used to sign the request. For details, see Authentication .
X-Language	No	Request language

3.2 Authentication

You can use either of the following authentication methods when calling APIs:

- AK/SK-based authentication: Requests are encrypted using AK/SK pairs.
- Token authentication: Requests are authenticated using a token.

Authentication Using Tokens

NOTE

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

A token specifies certain permissions in a computer system. Authentication using a token adds the token to a request as its header during API calling to obtain permissions to operate APIs through IAM.

```
{
  "auth": {
    "identity": {
      "methods": [
        "password"
      ],
      "password": {
        "user": {
          "name": "username", //Username
          "password": "*****", //Login password
          "domain": {
            "name": "domainname" //Name of the account to which the user belongs
          }
        }
      }
    }
  }
}
```

```
    }  
  },  
  "scope": {  
    "project": {  
      "name": "xxxxxxx" //Project name  
    }  
  }  
}
```

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....**, **X-Auth-Token: ABCDEFJ....** can be added to a request as follows:

```
Content-Type: application/json  
X-Auth-Token: ABCDEFJ....
```

3.3 Returned Values

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

Status Code

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](#).

If status code 201 is returned for the API for , the request is successful.

Response Header

A response header corresponds to a request header, for example, **Content-Type**.

4 API V2

4.1 Cluster Management

4.1.1 Creating a Cluster (V2)

Function

This API is used to create a cluster that uses multiple different types of nodes. The cluster engine type can be Elasticsearch, OpenSearch, or Logstash.

- An Elasticsearch cluster supports the following nodes:
- Data node (ess): mandatory. Data nodes store the data of a cluster. If there are no client nodes in the cluster, data nodes will need to additionally handle cluster access and data analytics requests; if there are no master nodes, data nodes will need to provide cluster management.
- Master nodes (ess-master): optional. Master nodes manage cluster-wide operations, including metadata, indexes, and shard allocation. For large-scale deployments, using dedicated master nodes enhances cluster stability, service availability, and centralized control.
- Client nodes (ess-client): Client nodes route and coordinate search and index requests, offloading processing from data nodes for enhanced query performance and cluster scalability when there are heavy loads.
- Cold data nodes (ess-cold): Cold data nodes are used to store and query large quantities of latency-insensitive data. They offer an effective way to manage large datasets while cutting storage costs.
- OpenSearch clusters support the same nodes as Elasticsearch clusters.
- Logstash clusters support the creation of Logstash nodes (lgs). Logstash nodes ingest, parse, process, and transmit data.

Constraints

- Before creating a cluster, make sure you have planned its configuration, including the cluster size, node specifications, and storage capacity requirements.

- You have created a VPC, subnet, and security group beforehand, or you can create them when creating the cluster.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v2.0/{project_id}/clusters

Table 4-1 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Request Parameters

Table 4-2 Request body parameters

Parameter	Mandatory	Type	Description
cluster	Yes	CreateCluster Body object	Definition: Cluster object. Constraints: N/A Value range: N/A Default value: N/A

Table 4-3 CreateClusterBody

Parameter	Mandatory	Type	Description
name	Yes	String	Definition: Cluster name. Constraints: N/A Value range: The value must consist of 4 to 32 characters, and must start with a letter. It can only contain digits, letters, periods (.), hyphens (-), and underscores (_). Default value: N/A
desc	No	String	Definition: Cluster description. Constraints: Enter a maximum of 128 characters. Value range: Digits, letters, periods (.), hyphens (-), underscores (_), and other special characters are allowed. Default value: N/A
backupStrategy	No	CreateClusterBackupStrategyBody object	Definition: Automatic snapshot creation. Constraints: The automatic snapshot creation policy is enabled only when backupStrategy is not left blank. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
roles	Yes	Array of CreateClusterRolesBody objects	Definition: Cluster node information. Constraints: N/A Value range: N/A Default value: N/A
nics	Yes	CreateClusterInstanceNicsBody object	Definition: Network configuration information. Constraints: N/A Value range: N/A Default value: N/A
enterprise_project_id	No	String	Definition: Enterprise project ID. When you create a cluster, associate the enterprise project ID with the cluster. Constraints: N/A Value range: The value can contain a maximum of 36 characters. It is string 0 or in UUID format with hyphens (-). Default value: 0 indicates an enterprise project.

Parameter	Mandatory	Type	Description
tags	No	Array of CreateClusterTagsBody objects	Definition: Cluster tags. Constraints: N/A Value range: N/A Default value: N/A
availability_zone	Yes	String	Definition: AZ of the node. The AZ name needs to be specified. Constraints: If multiple AZs are selected, the name of each AZ must be unique, and the number of nodes must be greater than or equal to the number of AZs. If the number of nodes is evenly divisible by the number of AZs, the nodes will be evenly distributed among all AZs. If the number of nodes is not evenly divisible by the number of AZs, the absolute difference between node quantity in any two AZs is 1 at most. To learn more about AZs, see Basic Concepts in CSS Before You Start. Value range: When specifying the AZs of nodes, use commas (,) to separate different AZ names. For example, if multiple AZs are selected for general AZs, the value of availability_zone is region-1a,region-1b,region-1c. Default value: N/A

Parameter	Mandatory	Type	Description
datastore	Yes	CreateClusterDatastoreBody object	Definition: Engine type of the cluster to be created. Constraints: N/A Value range: The value can be elasticsearch, opensearch, or logstash. Default value: N/A
authorityEnable	No	Boolean	Definition: Whether to enable security authentication for a cluster. Constraints: This parameter is available only for Elasticsearch and OpenSearch clusters. • Elasticsearch: Only Elasticsearch 6.5.4 and later versions support security-mode clusters. • OpenSearch: No constraints. • true: Authentication is enabled for the cluster. • false: Authentication is disabled for the cluster. Value range: Default value: false

Parameter	Mandatory	Type	Description
httpsEnable	No	Boolean	Definition: Whether communication is encrypted on the cluster. Constraints: When httpsEnable is set to true, authorityEnable must be set to true. Only Elasticsearch 6.5.4 security-mode clusters and later support this parameter. HTTPS access can be disabled only for OpenSearch 1.3.6 and 2.19.0 security-mode clusters. For other versions, HTTPS access is forcibly enabled and cannot be disabled. Value range: • true: Communication is encrypted on the cluster. • false: Communication is not encrypted for the cluster. Default value: false
adminPwd	No	String	Definition: Password of the administrator admin for a security-mode cluster. Constraints: This parameter is mandatory only when authorityEnable is set to true during cluster creation. Value range: A string of 8 to 32 characters. Must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters ~!@#\$\$%&*()-_=[{}];,;<.>/? [{}];,;<.>/? Default value: N/A

Parameter	Mandatory	Type	Description
publicIPReq	No	CreateClusterPublicIPReq object	Definition: Public network access information. Constraints: This parameter is valid only when httpsEnable is set to true. Value range: N/A Default value: N/A
loadBalance	No	CreateClusterLoadBalance object	Definition: VPC endpoint service information. Constraints: N/A Value range: N/A Default value: N/A
publicKibanaReq	No	CreateClusterPublicKibanaReq object	Definition: Kibana public network access information. Constraints: This parameter is valid only when authorityEnable is set to true. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
ipv6_enable	No	Boolean	Definition: Whether to enable automatic IPv6 address assignment for the cluster. Constraints: N/A Value range: • true: Automatic IPv6 address assignment is enabled for the cluster. • false: Automatic IPv6 address assignment is disabled for the cluster. Default value: false
diskEncryption	No	diskEncryptionInfo object	Definition: Disk encryption information. Constraints: N/A Value range: N/A Default value: N/A

Table 4-4 CreateClusterBackupStrategyBody

Parameter	Mandatory	Type	Description
period	Yes	String	Definition: Time when a snapshot is created every day. Constraints: Snapshots can only be created on the hour. The time format is the time followed by the time zone, specifically, HH:mm z. In the format, HH:mm refers to the hour time and z refers to the time zone. For example, 00:00 GMT+08:00 and 01:00 GMT+08:00. Value range: N/A Default value: 00:00 GMT+08:00. NOTE When frequency is set to HOUR, you do not need to specify period. By default, snapshot creation starts from the next hour after the settings are successful.

Parameter	Mandatory	Type	Description
frequency	No	String	Definition: Frequency of automatically creating snapshots. Constraints: The number of retained snapshots depends on the execution frequency and indexes set in the automatic snapshot creation policy. If the execution interval is short or the index data volume is large, the number of retained automatic snapshots may not reach the preset value. Set this value based on actual conditions. Value range: • HOUR: Execute once every hour on the hour. • DAY: Execute once every day. • SUN, MON, TUE, WED, THU, FRI, and SAT: Execute once on the specified day of a week. For example, SUN indicates that the task is executed once every Sunday. Default value: DAY
prefix	Yes	String	Definition: Prefix of a snapshot that is automatically created, which is manually entered. Constraints: Enter up to 32 characters and start with a lowercase letter. Lowercase letters, digits, hyphens (-), and underscores (_) are allowed. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
keepday	Yes	Integer	Definition: Customize the number of snapshots to be retained. Constraints: Expired snapshots will be automatically deleted on the half hour. The deletion policy applies only to automated snapshots that are executed at the same frequency as the current automatic snapshot creation policy. Value range: 1–90 Default value: N/A
bucket	Yes	String	Definition: Name of the OBS bucket used for backup. Constraints: Select an OBS bucket that meets the following requirements: • Storage class: Standard. • Region: the same as that of the cluster. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
basePath	Yes	String	Definition: Storage path of the snapshot in the OBS bucket. Constraints: The backup path cannot: • Contain the following characters: :?* "<> '{'} • Start with a slash (/). • Start or end with a period (.). • Contain more than two consecutive slashes (/) or periods (.). • Exceed 512 characters. Value range: N/A Default value: N/A
agency	No	String	Definition: Select an IAM agency to grant the current account the permission to access and use OBS. Constraints: N/A Value range: The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Default value: N/A NOTE If none of the bucket, basePath, and agency parameters are specified, the system will automatically create an OBS bucket and an IAM agency. If the creation fails, you need to manually configure correct parameters.

Parameter	Mandatory	Type	Description
maxSnapshotBytesPerSeconds	No	String	Parameter description: This parameter sets the maximum backup speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Constraints N/A Value range: The following values and formats are allowed: • Number + Unit The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). 0mb means there is no speed limit. An overly high backup speed may lead to excessive resource usage, which may impact system stability. Therefore, set this parameter carefully. • 0 There is no speed limit. An overly high backup speed may lead to excessive resource usage, which may impact system stability. Therefore, set this parameter carefully. • -1 There is no speed limit. An overly high backup speed may lead to excessive resource usage, which may impact system stability. Therefore, set this parameter carefully. Default value:

Parameter	Mandatory	Type	Description
			40mb

Parameter	Mandatory	Type	Description
maxRestoreBytesPerSeconds	No	String	Parameter description: This parameter sets the maximum recovery speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Constraints For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter. If Maximum recovery speed (per second) is lower than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the former is reached. If Maximum recovery speed (per second) is higher than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the latter is reached. Value range: The following values and formats are allowed: • Number + Unit The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). 0mb means there is no speed limit. (However, for OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter.) An overly high recovery speed may lead

Parameter	Mandatory	Type	Description
			to excessive resource usage, which may impact system stability. Therefore, set this parameter carefully. • 0 There is no speed limit. (However, for OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter.) An overly high recovery speed may lead to excessive resource usage, which may impact system stability. Therefore, set this parameter carefully. • -1 There is no speed limit. (However, for OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter.) An overly high recovery speed may lead to excessive resource usage, which may impact system stability. Therefore, set this parameter carefully. Default value: For Elasticsearch clusters of 7.6.2 or earlier, the default setting is 40mb. For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the default setting is no limit, but the recovery speed is still limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter.

Table 4-5 CreateClusterRolesBody

Parameter	Mandatory	Type	Description
flavorRef	Yes	String	Definition: Name flavor name. For details about the current instance flavor, see the name attribute in Obtaining the Instance Specifications List . Constraints: For details about the current instance flavor, see the name attribute in Obtaining the Instance Specifications List . Value range: N/A Default value: N/A
volume	Yes	CreateClusterInstanceVolumeBody object	Definition: Volume information. Constraints: If flavorRef is set to a local disk flavor, leave this parameter blank. Obtain the local disk flavor by referring to the diskrange attribute in Obtaining the Instance Specifications List . Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
type	Yes	String	Definition: Node type. Constraints: N/A Value range: • ess-master indicates a master node. • ess-client indicates a client node. • ess-cold indicates a cold data node. • ess indicates a data node. • lgs indicates a Logstash node. • english indicates an English language model. • arabic indicates an Arabic model. • thai indicates a Thai model. • turkish indicates a Turkish language model. • tools indicates a tools model. • portuguese indicates a Portuguese model. • spanish indicates a Spanish model. • chinese-english: general model for Chinese and English. • spanish-portuguese: general model for Spanish and Portuguese.] (tag:non_helpcenter) [- chinese indicates a Chinese language model. Default value: N/A

Parameter	Mandatory	Type	Description
instanceNum	Yes	Integer	Definition: Number of instances. Constraints: See the value range. Value range: • If ess is selected, you can set the value to 1 to 32. – If both ess and ess-master are selected, you can set the value to 1 to 200. – If both ess and ess-client are selected, you can set the value to 1 to 32. – If both ess and ess-cold are selected, you can set the value to 32. • If ess-master is selected, you can set the value to 3 to 9. • If ess-client is selected, you can set the value to 1 to 64. • If ess-cold is selected, you can set the value to 1 to 32. • For Logstash nodes, the value ranges from 1 to 100. Default value: N/A

Table 4-6 CreateClusterInstanceVolumeBody

Parameter	Mandatory	Type	Description
volume_type	Yes	String	Definition: Disk type. Constraints: N/A Value range: • COMMON: common I/O • HIGH: high I/O • ULTRAHIGH: ultra-high I/O • ESSD: ultra-fast SSD Default value: N/A
size	Yes	Integer	Definition: Disk size. Constraints: The value must be greater than 0 and a common multiple of 4 and 10, in GB. Value range: You can obtain the disk size from the diskrange attribute in Obtaining the Instance Specifications List . Default value: • ess and ess-cold nodes: 100 GB or the minimum disk capacity supported by the selected node flavor, whichever is larger. • ess-master and ess-client nodes: The default volume size is 40 GB and cannot be changed. NOTE For ess and ess-cold nodes, 100 GB or a larger value is recommended.

Table 4-7 CreateClusterInstanceNicsBody

Parameter	Mandatory	Type	Description
vpcId	Yes	String	Definition: VPC ID, which is used for configuring cluster network. Constraints: N/A Value range: N/A Default value: N/A
netId	Yes	String	Definition: Subnet ID (network ID). Constraints: N/A Value range: N/A Default value: N/A
securityGroupId	Yes	String	Definition: Security group ID. Constraints: N/A Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
ips	No	Array of strings	Definition: Node IP address, which is configured when a cluster is created. Constraints: N/A Value range: Only Elasticsearch and OpenSearch support this parameter. Assign automatically: Private IPv4 addresses will be automatically assigned to cluster nodes. Assign manually: Manually assign private IPv4 addresses to cluster nodes. Before assigning an IP address, click View In-Use IP Address to check IP addresses that are already in use. Default value: N/A

Table 4-8 CreateClusterTagsBody

Parameter	Mandatory	Type	Description
key	Yes	String	Definition: Tag key. Constraints: N/A Value range: The value can contain 1 to 36 characters. Only digits, letters, hyphens (-), and underscores (_) are allowed. Default value: N/A

Parameter	Mandatory	Type	Description
value	Yes	String	Definition: Tag value. Constraints: N/A Value range: The value can contain 0 to 43 characters. It can only contain digits, letters, hyphens (-), and underscores (_). Default value: N/A

Table 4-9 CreateClusterDatastoreBody

Parameter	Mandatory	Type	Description
version	Yes	String	Definition: Version of the CSS cluster engine. For details, see the supported versions in Before You Start . Constraints: N/A Value range: N/A Default value: N/A
type	Yes	String	Definition: Engine type of the cluster to be created. Constraints: N/A Value range: • elasticsearch: Elasticsearch cluster. • opensearch: OpenSearch cluster. • logstash: Logstash cluster. Default value: N/A

Table 4-10 CreateClusterPublicIpReq

Parameter	Mandatory	Type	Description
eip	Yes	CreateClusterPublicEip object	Definition: Public network bandwidth. Constraints: N/A Value range: N/A Default value: N/A
elbWhiteListReq	No	CreateClusterElbWhiteList object	Definition: EIP whitelist. Constraints: N/A Value range: N/A Default value: N/A
publicBindType	Yes	String	Definition: Whether to automatically bind an EIP. Constraints: N/A Value range: Currently, only auto_assign is supported. Default value: N/A
eipId	No	String	Definition: EIP ID. Constraints: N/A Value range: N/A Default value: N/A

Table 4-11 CreateClusterPublicEip

Parameter	Mandatory	Type	Description
bandWidth	Yes	CreateClusterPublicEipSize object	Definition: Set the public network bandwidth. Constraints: N/A Value range: N/A Default value: N/A

Table 4-12 CreateClusterPublicEipSize

Parameter	Mandatory	Type	Description
size	Yes	Integer	Definition: Bandwidth size, in Mbit/s. Constraints: N/A Value range: 1-200 Default value: 1

Table 4-13 CreateClusterElbWhiteList

Parameter	Mandatory	Type	Description
enableWhiteList	Yes	Boolean	Definition: Whether to enable public network access control. Constraints: N/A Value range: • true: A whitelist is set. Only IP addresses or CIDR blocks on this whitelist can access the cluster from the public network. • false: No whitelist is set. All public IP addresses can access the cluster. Click +Add. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: 192.168.1.1,10.0.0.0/24. Examples of invalid values: 0.0.0.0, x.x.x.x0, 172.16.0.0-172.16.255.255, non-standard formats (e.g., 192.168.1), and duplicate values. Default value: false
whiteList	No	String	Definition: Public network access control whitelist. Constraints: Separate the whitelisted CIDR blocks or IP addresses with commas (,), and make sure each of them is unique. Value range: N/A Default value: N/A

Table 4-14 CreateClusterLoadBalance

Parameter	Mandatory	Type	Description
endpointWithDnsName	Yes	Boolean	Definition: Specifies the private domain name for accessing the VPC endpoint. Constraints: N/A Value range: • true: Internal domain names are supported. The system automatically assigns a private domain name to the VPC endpoint. After cluster creation, you can check this private domain name on the VPC Endpoint tab of the cluster details page. • false: Internal domain names are not supported. No private domain name is assigned for the VPC endpoint. The cluster can only be accessed through an IP address assigned to the VPC endpoint. Default value: false
vpcPermissions	No	Array of strings	Definition: Access control. Constraints: N/A Value range: N/A. Default value: N/A

Parameter	Mandatory	Type	Description
professionVpc ep	No	Boolean	Definition: Create a professional VPC endpoint. Constraints: N/A Value range: • true: enable. Create professional VPC endpoints that support the IPv4/IPv6 dual-stack network. • false: disable. Create basic VPC endpoints that support IPv4 only. Default value: false
dualstackEna ble	No	Boolean	Definition: Whether to enable IPv4/IPv6 dual-stack networking. Constraints: The IPv4/IPv6 dual-stack network can be enabled only when a professional VPC endpoint is created and the VPC of the cluster supports IPv6. Value range: • true: enable. The system automatically assigns an IPv6 address to the VPC endpoint. • false: disable. The VPC endpoint does not support IPv6 addresses. Default value: false

Table 4-15 CreateClusterPublicKibanaReq

Parameter	Mandatory	Type	Description
eipSize	Yes	Integer	Definition: Bandwidth size, in Mbit/s. Constraints: N/A Value range: 1-200 Default value: 1
elbWhiteList	No	CreateClusterPublicKibanaElbWhiteList object	Definition: Kibana whitelist information. Constraints: N/A Value range: N/A Default value: N/A

Table 4-16 CreateClusterPublicKibanaElbWhiteList

Parameter	Mandatory	Type	Description
whiteList	Yes	String	Definition: Whitelists Constraints: Separate the whitelisted CIDR blocks or IP addresses with commas (,), and make sure each of them is unique. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
enableWhiteList	Yes	Boolean	Definition: Whether to enable Kibana access control. Constraints: N/A Value range: • true: A whitelist is set. Only IP addresses or CIDR blocks on this whitelist can access the cluster's Kibana console from the public network. • false: No whitelist is set. All public IP addresses can access the cluster through Kibana. Click +Add. In the displayed text box, enter IP addresses or CIDR blocks that are allowed to access the cluster's Kibana console from the public network. Separate them using commas (,). Each value must be unique. An example of valid values: 192.168.1.1,10.0.0.0/24. Examples of invalid values: 0.0.0.0, x.x.x.x0, 172.16.0.0-172.16.255.255, non-standard formats (e.g., 192.168.1), and duplicate values. Default value: false

Table 4-17 diskEncryptionInfo

Parameter	Mandatory	Type	Description
systemEncrypted	No	String	Definition: Whether disk encryption is enabled. KMS is used to encrypt the data disks of cluster nodes to ensure the security of stored data. Constraints: Disk encryption and decryption do not alter cluster management or O&M processes. However, they do increase the system's processing load, potentially affecting the system's operational performance. Value range: • 0: Disable disk encryption. • 1: Enable disk encryption. Default value: 0

Parameter	Mandatory	Type	Description
systemCmkid	No	String	Definition: Disk key ID. Constraints: This parameter is valid only when systemEncrypted is set to 1. Select a KMS key from the drop-down list. KMS keys that are unavailable in this drop-down list are not supported by the cluster. Set a key that meets the following requirements: • Cryptographic algorithm: AES or SM4. • Key usage: ENCRYPT_DECRYPT. After cluster creation, the KMS key cannot be changed. If the KMS key used by the cluster is disabled, the cluster cannot be scaled or upgraded, its node specifications or AZs cannot be changed, and its nodes cannot be replaced (by specifying the nodes that need replacement). To solve this problem, you will have to create a new cluster and migrate your data to that new cluster. Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Table 4-18 Response body parameters

Parameter	Type	Description
cluster	CreateClusterResp object	Definition: Cluster object. For a pay-per-use cluster, only the cluster parameter is returned. Value range: N/A.

Table 4-19 CreateClusterResp

Parameter	Type	Description
id	String	Definition: Cluster ID. Value range: N/A.
name	String	Definition: Cluster name. Value range: N/A.

Example Requests

- Example request for creating a pay-per-use Elasticsearch cluster

POST https://{Endpoint}/v2.0/{project_id}/clusters

```
{
  "cluster" : {
    "name" : "css-cluster",
    "desc" : "Cluster testing",
    "backupStrategy" : {
      "period" : "16:00 GMT+08:00",
      "prefix" : "snapshot",
      "keepday" : 7,
      "frequency" : "DAY",
      "bucket" : "OBS-obsName",
      "basePath" : "css_repository/obs-path",
      "agency" : "css_obs_agency",
      "maxSnapshotBytesPerSeconds" : 40,
      "maxRestoreBytesPerSeconds" : 40
    },
    "roles" : [ {
      "flavorRef" : "ess.spec-4u8g",
      "volume" : {
        "volume_type" : "COMMON",
        "size" : 100
      },
      "type" : "ess",
      "instanceNum" : 1
    } ],
  }
}
```

```
"nics" : {
  "vpclId" : "{VPC ID}",
  "netId" : "{NET ID}",
  "securityGroupId" : "{Security group ID}"
},
"enterprise_project_id" : 0,
"tags" : [ {
  "key" : "k1",
  "value" : "v1"
}, {
  "key" : "k2",
  "value" : "v2"
} ],
"availability_zone" : "{Az code}",
"datastore" : {
  "version" : "{cluster-version}",
  "type" : "elasticsearch"
},
"authorityEnable" : true,
"httpsEnable" : true,
"adminPwd" : "{password}",
"publicIPReq" : {
  "eip" : {
    "bandWidth" : {
      "size" : 5
    }
  },
  "elbWhiteListReq" : {
    "enableWhiteList" : true,
    "whiteList" : "127.0.0.1"
  },
  "publicBindType" : "auto_assign",
  "eipId" : null
},
"loadBalance" : {
  "endpointWithDnsName" : false,
  "vpcPermissions" : [ "{Account ID}" ]
},
"publicKibanaReq" : {
  "eipSize" : 5,
  "elbWhiteList" : {
    "whiteList" : "127.0.0.1",
    "enableWhiteList" : true
  }
}
}
```

Example Responses

Status code: 200

Request succeeded.

- Example response for creating a pay-per-use cluster.

```
{
  "cluster" : {
    "id" : "ef683016-871e-48bc-bf93-74a29d60d214",
    "name" : "ES-Test"
  }
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

4.1.2 Restarting a Cluster (V2)

Function

This API is used to restart nodes of all or some types in the current cluster.

NOTE

If a cluster is available, ensure that it has stopped handling service requests, such as importing and searching for data. Otherwise, the cluster restart may cause data loss.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v2.0/{project_id}/clusters/{cluster_id}/restart

Table 4-20 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster you want to restart. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 4-21 Request body parameters

Parameter	Mandatory	Type	Description
type	Yes	String	Definition: Restart mode. Value range: - node: restarts nodes. - role: restarts roles.

Parameter	Mandatory	Type	Description
value	Yes	String	Operation parameters. Parameter description: - When the operation role is node, the value is the node ID. You can obtain the ID attribute in instances by referring to Querying Cluster Details. - When the operation role is role, the value is one or multiple node types (such as ess, ess-master, ess-client, and ess-cold).

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Restart nodes of all or some types in the current cluster.

```
POST https://{Endpoint}/v2.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/restart
{
  "type" : "role",
  "value" : "ess"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.

Status Code	Description
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

4.1.3 Rolling Restart (Deprecated)

Function

This API is used to restart nodes one by one, which requires a long time when the nodes have a large number of indexes.

NOTE

Rolling restart is supported only when the cluster has more than three nodes (including master nodes, client nodes, and cold data nodes). When the cluster is available, ensure that the cluster has stopped handling service requests (such as importing data and searching for data). Otherwise, the cluster restart may cause data loss.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v2.0/{project_id}/clusters/{cluster_id}/rolling_restart

Table 4-22 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster you want to restart. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 4-23 Request body parameters

Parameter	Mandatory	Type	Description
type	Yes	String	Operation role. Its type can only be role .

Parameter	Mandatory	Type	Description
value	Yes	String	Instance type. (At least one data node is required when you configure instance types.) Use commas (,) to separate multiple types. For example: ● ess-master indicates a master node. ● ess-client indicates a client node. ● ess-cold indicates a cold data node. ● ess indicates a data node. ● all indicates all nodes.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Restart a node.

```
POST https://{Endpoint}/v2.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/rolling_restart
{
  "type" : "role",
  "value" : "ess"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.

Status Code	Description
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

4.1.4 Rolling Restart

Function

This API is used to restart nodes one by one, which requires a long time when the nodes have a large number of indexes.

NOTE

Rolling restart is supported only when the cluster has more than three nodes (including master nodes, client nodes, and cold data nodes). When the cluster is available, ensure that the cluster has stopped handling service requests (such as importing data and searching for data). Otherwise, the cluster restart may cause data loss.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v2.0/{project_id}/clusters/{cluster_id}/rolling-restart

Table 4-24 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster you want to restart. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 4-25 Request body parameters

Parameter	Mandatory	Type	Description
type	Yes	String	Operation role. Its type can only be role .

Parameter	Mandatory	Type	Description
value	Yes	String	Instance type. (At least one data node is required when you configure instance types.) Use commas (,) to separate multiple types. For example: ● ess-master indicates a master node. ● ess-client indicates a client node. ● ess-cold indicates a cold data node. ● ess indicates a data node. ● all indicates all nodes.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Restart the node.

```
POST https://{Endpoint}/v2.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/rolling-restart
{
  "type" : "role",
  "value" : "ess"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. Do not retry the request without modifying it.

Status Code	Description
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request failed to be processed because of the update of the conflict request.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

4.2 Snapshot Management

4.2.1 Enabling Automatic Snapshot Creation

Function

This API is used to enable automatic snapshot creation for clusters. A snapshot saves a point-in-time copy of a cluster in an OBS bucket. Using this snapshot you can restore the cluster to a previous state.

When automatic snapshot creation is enabled, snapshots are automatically created daily or weekly based on a preset time schedule, ensuring ongoing data protection.

Only Elasticsearch and OpenSearch clusters support automatic snapshot creation.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v2.0/{project_id}/clusters/{cluster_id}/snapshots/policy/open

Table 4-26 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: Cluster ID. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 4-27 Request body parameters

Parameter	Mandatory	Type	Description
indices	No	String	Definition: Name of the index to be backed up. - You can specify indexes you want to back up. To specify multiple indexes, use commas (,) to separate them, for example, index1,index2,index3. - You can use an asterisk (*) <i>to match multiple indexes</i>. For example, <i>index</i> indicates that all indexes with the prefix index will be restored. Constraints: N/A Value range: The value can contain 0 to 1024 characters, and cannot contain spaces, uppercase letters, or the following special characters: "\< >/?". Default value: The default value is empty, indicating that all indexes will be backed up.

Parameter	Mandatory	Type	Description
keepday	Yes	Integer	Definition: Set the number of snapshots to be retained. Expired snapshots will be automatically deleted on the half hour. The deletion policy applies only to auto snapshots that are executed at the same frequency as the current automatic snapshot creation policy. Constraints: If the snapshot creation interval is short or if the data size of indexes is large, the number of automatic snapshots retained may not reach the value set using this parameter. Value range: 1~90 Default value: N/A
frequency	No	String	Definition: Frequency of automatically creating snapshots. Constraints: N/A Value range: ● HOUR: Execute once every hour on the hour. ● DAY: Execute once every day. ● SUN, MON, TUE, WED, THU, FRI, and SAT: Execute the task at the specified day of every week. For example, SUN indicates that the task is executed once every Sunday. Default value: DAY

Parameter	Mandatory	Type	Description
period	No	String	Definition: Time when a snapshot is created every day. Constraints: When frequency is set to HOUR, you do not need to specify period. By default, snapshot creation starts from the next hour after the settings are successful. Value range: Snapshots can only be created on the hour. The time format is the time followed by the time zone, specifically, HH:mm z . In the format, HH:mm refers to the hour time and z refers to the time zone. For example, 00:00 GMT+08:00 and 01:00 GMT+08:00. Default value: 00:00 GMT+08:00
prefix	Yes	String	Definition: Prefix of the name of an automatically created snapshot. A snapshot name consists of the snapshot name prefix and timestamp, for example, snapshot-1566921603720. Constraints: N/A Value range: Enter up to 32 characters and start with a lowercase letter. Lowercase letters, digits, hyphens (-), and underscores (_) are allowed. Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Enable automatic backup. Seven snapshots are automatically created every day at 16:00. The snapshot name prefix is snapshot.XXXX.

```
POST https://{Endpoint}/v2.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/snapshots/policy/open
```

```
{
  "indices" : "*",
  "keepday" : 7,
  "frequency" : "DAY",
  "period" : "16:00 GMT+08:00",
  "prefix" : "snapshot"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
403	Request rejected.The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

4.2.2 Disabling Automatic Snapshot Creation

Function

This API is used to disable automatic snapshot creation for a cluster. If a cluster does not need automatic backup, you are advised to disable automatic snapshot creation to reduce storage costs.

You can only use this API on Elasticsearch and OpenSearch clusters with automatic snapshot creation enabled.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v2.0/{project_id}/clusters/{cluster_id}/snapshots/policy/close

Table 4-28 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: Cluster ID. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Disable the automatic backup function.

```
PUT https://{Endpoint}/v2.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/snapshots/  
policy/close
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
403	Request rejected.The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

4.3 Logstash

4.3.1 Delete a Configuration File (V2)

Function

Delete a configuration file.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v2.0/{project_id}/clusters/{cluster_id}/lgsconf/delete

Table 4-29 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose configuration file is to be deleted. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 4-30 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Configuration file name.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Deleting a Configuration File

```
POST https://{Endpoint}/v2.0/{project_id}/clusters/5c2bab0d-ec5f-491f-9418-b8da38a26760/lgsconf/delete
{
  "name" : "confName"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. Do not retry the request without modifying it.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request failed to be processed because of the update of the conflict request.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

4.3.2 Delete a Custom Template V2

Function

This API is used to delete a custom template.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v2.0/{project_id}/lgsconf/deletetemplate

Table 4-31 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Request Parameters

Table 4-32 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Template name.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Delete a custom template.

```
POST https://{Endpoint}/v2.0/{project_id}/lgsconf/deletetemplate
{
  "name" : "template1"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. Do not retry the request without modifying it.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request failed to be processed because of the update of the conflict request.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5 API V1

5.1 Cluster Management

5.1.1 Creating a Cluster (Discarded)

Function

This API can only be used to create clusters that contain a single type of nodes. To create clusters that contain multiple types of nodes (ess, ess-cold, ess-client, and ess-master), use the API for [Creating a Cluster v2](#).

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters

Table 5-1 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Request Parameters

Table 5-2 Request body parameters

Parameter	Mandatory	Type	Description
cluster	Yes	CreateCluster Body object	Parameter description: Cluster object. Constraints: N/A Options: N/A Default value: N/A

Table 5-3 CreateClusterBody

Parameter	Mandatory	Type	Description
instance	Yes	CreateClusterInstanceBody object	Parameter description: Array of Buffer instances, whose memory is to be shared. Constraints: N/A Options: N/A Default value: N/A
datastore	Yes	CreateClusterDatastoreBody object	Definition: Search engine. Constraints: N/A Value range: N/A Default value: N/A
name	Yes	String	Definition: Cluster name. Constraints: N/A Value range: It can contain 4 to 32 characters. Only digits, letters, hyphens (-), and underscores (_) are allowed. The value must start with a letter. Default value: N/A

Parameter	Mandatory	Type	Description
desc	No	String	Definition: Cluster description. Constraints: Enter a maximum of 128 characters. Value range: The value can contain digits, letters, periods (.), hyphens (-), underscores (_), and other special characters. Default value: N/A
instanceNum	Yes	Integer	Definition: Number of cluster instances. Constraints: N/A Value range: 1–32. Default value: N/A
backupStrategy	No	CreateClusterBackupStrategyBody object	Parameter description: Automatic snapshot creation policy, which is disabled by default. The automatic snapshot creation policy is enabled only when backupStrategy is not left blank. Constraints: N/A Options: N/A Default value: N/A

Parameter	Mandatory	Type	Description
httpsEnable	No	Boolean	Definition: Whether communication is encrypted on the cluster. The value can be true or false. By default, communication encryption is disabled. When httpsEnable is set to true, authorityEnable must be set to true. Constraints: Only Elasticsearch 6.5.4 security-mode clusters and later support this parameter. HTTPS access can be disabled only for OpenSearch 1.3.6 and 2.19.0 security-mode clusters. For other versions, HTTPS access is forcibly enabled and cannot be disabled. Value range: • true: Communication is encrypted on the cluster. • false: Communication is not encrypted for the cluster. Default value: false
authorityEnable	No	Boolean	Definition: Whether to enable authentication. Available values include true and false. Authentication is disabled by default. Constraints: Only Elasticsearch 6.5.4 security-mode clusters and later support this parameter. Value range: • true: Authentication is enabled for the cluster. • false: Authentication is disabled for the cluster. Default value: false

Parameter	Mandatory	Type	Description
adminPwd	No	String	Parameter description: Password of the cluster user admin in security mode. This parameter is mandatory only when authorityEnable is set to true. Weak password verification is required for a security cluster. You are advised to set a strong password. Constraints: The administrator password must meet the following requirements: • The password can contain 8 to 32 characters. • Passwords must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters ~!@#\$\$%^&*()-_+=\ [{ }];:;<.>/? Options: See Constraints. Default value: N/A
enterprise_project_id	No	String	Parameter description: Enterprise project ID. When creating a cluster, associate the enterprise project ID with the cluster. The value can contain a maximum of 36 characters. It is string 0 or in UUID format with hyphens (-). Value 0 indicates the default enterprise project. Constraints: N/A Options: N/A Default value: N/A

Parameter	Mandatory	Type	Description
tags	No	Array of CreateClusterTagsBody objects	Parameter description: Cluster tags. Constraints: N/A Options: N/A Default value: N/A
ipv6_enable	No	Boolean	Definition: Whether to enable automatic IPv6 address assignment for the cluster. The default value is false. Constraints: N/A Value range: N/A Default value: false
diskEncryption	No	diskEncryptionInfo object	Parameter description: Disk encryption information. Constraints: N/A Options: N/A Default value: N/A

Table 5-4 CreateClusterInstanceBody

Parameter	Mandatory	Type	Description
flavorRef	Yes	String	Parameter description: Instance flavor name. For details about the current instance flavor, see the name attribute in Obtaining the Instance Specifications List . Constraints: N/A Options: N/A Default value: N/A
volume	Yes	CreateClusterInstanceVolumeBody object	Parameter description: Volume information. You do not need to set this parameter when flavorRef is set to a local disk flavor. Currently, the following local disk flavors are supported: • ess.spec-i3small • ess.spec-i3medium • ess.spec-i3.8xlarge.8 • ess.spec-ds.xlarge.8 • ess.spec-ds.2xlarge.8 • ess.spec-ds.4xlarge.8 Constraints: N/A Options: N/A Default value: N/A
nics	Yes	CreateClusterInstanceNicsBody object	Parameter description: Subnet information. Constraints: N/A Options: N/A Default value: N/A

Parameter	Mandatory	Type	Description
availability_zone	Yes	String	Definition: AZ of the node. The AZ name needs to be specified. Constraints: When specifying the AZs of nodes, use commas (,) to separate different AZ names. For example, if multiple AZs are selected for general AZs, the value of availability_zone is region-1a,region-1b,region-1c. Value range: N/A Default value: availability_zone is empty by default, indicating that a single AZ is used. NOTE If multiple AZs are selected, the name of each AZ must be unique, and the number of nodes must be greater than or equal to the number of AZs. If the number of nodes is evenly divisible by the number of AZs, the nodes will be evenly distributed among all AZs. If the number of nodes is not evenly divisible by the number of AZs, the absolute difference between node quantity in any two AZs is 1 at most. AZ name. Obtain it from Regions and Endpoints in section "Before You Start".

Table 5-5 CreateClusterInstanceVolumeBody

Parameter	Mandatory	Type	Description
volume_type	Yes	String	Parameter description: Volume type. Constraints: N/A Options: • COMMON: common I/O • HIGH: high I/O • ULTRAHIGH: ultra-high I/O • ESSD: ultra-fast SSD Default value: N/A
size	Yes	Integer	Definition: Disk size. Constraints: The value must be greater than 0 and a common multiple of 4 and 10, in GB. Value range: You can obtain the disk size from the diskrange attribute in Obtaining the Instance Specifications List . Default value: • ess and ess-cold nodes: 100 GB or the minimum disk capacity supported by the selected node flavor, whichever is larger. • ess-master and ess-client nodes: The default volume size is 40 GB and cannot be changed. NOTE For ess and ess-cold nodes, 100 GB or a larger value is recommended.

Table 5-6 CreateClusterInstanceNicsBody

Parameter	Mandatory	Type	Description
vpcId	Yes	String	Parameter description: VPC ID, which is used for configuring cluster network. Constraints: N/A Options: N/A Default value: N/A
netId	Yes	String	Parameter description: Subnet ID (network ID). Constraints: N/A Options: N/A Default value: N/A
securityGroup Id	Yes	String	Parameter description: Security group ID Constraints: N/A Options: N/A Default value: N/A

Table 5-7 CreateClusterDatastoreBody

Parameter	Mandatory	Type	Description
version	Yes	String	Parameter description: Elasticsearch/Logstash engine version. For details, see the supported versions in Before You Start . Constraints: N/A Options: N/A Default value: N/A
type	Yes	String	Definition: Engine type of the cluster to be created. Constraints: N/A Value range: • elasticsearch: Elasticsearch cluster. • opensearch: OpenSearch cluster. • logstash: Logstash cluster. Default value: N/A

Table 5-8 CreateClusterBackupStrategyBody

Parameter	Mandatory	Type	Description
period	No	String	Definition: Time when a snapshot is generated every day. Constraints: N/A Value range: Snapshots can only be created on the hour. The time format is the time followed by the time zone, specifically, HH:mm z . In the format, HH:mm refers to the hour time and z refers to the time zone. For example, 00:00 GMT+08:00 and 01:00 GMT+08:00. Default value: 00:00 GMT+08:00 NOTE When frequency is set to HOUR, you do not need to specify period. By default, snapshot creation starts from the next hour after the settings are successful.
prefix	Yes	String	Definition: Prefix of a snapshot that is automatically created, which is manually entered. Constraints: N/A Value range: Enter up to 32 characters and start with a lowercase letter. Lowercase letters, digits, hyphens (-), and underscores (_) are allowed. Default value: N/A

Parameter	Mandatory	Type	Description
keepday	Yes	Integer	Definition: Set the number of snapshots to be retained. Expired snapshots will be automatically deleted on the half hour. The deletion policy applies only to automated snapshots that are executed at the same frequency as the current automatic snapshot creation policy. Constraints: N/A Value range: 1–90 Default value: N/A
bucket	No	String	Parameter description: Name of the OBS bucket used for backup. Constraints: N/A Options: N/A Default value: N/A
basePath	No	String	Parameter description: Storage path of the snapshot in the OBS bucket. Constraints: N/A Options: N/A Default value: N/A

Parameter	Mandatory	Type	Description
agency	No	String	Definition: Agency name. You can create an agency to allow CSS to call other cloud services. An agency name cannot contain special characters or Chinese characters. Constraints: N/A Value range: Only a-z, A-Z, 0-9, hyphens (-) and underscores (_) are allowed. Default value: N/A NOTE If none of the bucket, basePath, and agency parameters are specified, the system will automatically create an OBS bucket and an IAM agency. If the creation fails, you need to manually configure correct parameters.

Parameter	Mandatory	Type	Description
frequency	No	String	Definition: Frequency of automatically creating snapshots. Constraints: N/A Value range: • HOUR: Execute once every hour on the hour. • DAY: Execute once every day. • SUN, MON, TUE, WED, THU, FRI, and SAT: Execute the task at the specified day of every week. For example, SUN indicates that the task is executed once every Sunday. • The number of retained snapshots depends on the execution frequency and indexes set in the automatic snapshot creation policy. If the execution interval is short or the index data volume is large, the number of retained automatic snapshots may not reach the preset value. Set this value based on actual conditions. Default value: DAY

Parameter	Mandatory	Type	Description
maxSnapshotBytesPerSeconds	No	String	Definition: This parameter sets the maximum backup speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Constraints: N/A Value range: The following values and formats are allowed: • Number + Unit • 0 • -1 The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). Setting this parameter to 0mb, 0, -1 means there is no speed limit. An overly high backup speed may lead to excessive resource usage, which may impact system stability. Default value: The default value is 40 MB. If this parameter is left blank, the default value is used.

Parameter	Mandatory	Type	Description
maxRestoreBytesPerSeconds	No	String	Definition: This parameter sets the maximum recovery speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Constraints: For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter. If Maximum recovery speed (per second) is lower than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the former is reached. If Maximum recovery speed (per second) is higher than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the latter is reached. Value range: The following values and formats are allowed: • Number + Unit • 0 • -1 The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). Setting this parameter to 0mb, 0, -1 means there is no speed limit. (However, for OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed

Parameter	Mandatory	Type	Description
			is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter.) An overly high recovery speed may lead to excessive resource usage, which may impact system stability. Default value: For Elasticsearch clusters of 7.6.2 or earlier, the default setting is 40mb. For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the default setting is no limit, but the recovery speed is still limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter. If this parameter is left blank, the default setting is used.

Table 5-9 CreateClusterTagsBody

Parameter	Mandatory	Type	Description
key	Yes	String	Definition: Tag key. Constraints: N/A Value range: The value can contain 1 to 36 characters. Only digits, letters, hyphens (-), and underscores (_) are allowed. Default value: N/A

Parameter	Mandatory	Type	Description
value	Yes	String	Definition: Tag value. Constraints: N/A Value range: - Length: 0 to 255 characters. - The value can contain UTF-8 letters, digits, spaces, and special characters <code>_./:=+-@</code> Default value: N/A

Table 5-10 diskEncryptionInfo

Parameter	Mandatory	Type	Description
systemEncrypted	No	String	Whether to enable disk encryption. The value can be 0 or 1. The default value is 0. 0: Disable disk encryption. 1: Enable disk encryption.
systemCmkid	No	String	KMS key ID. This parameter is valid only when systemEncrypted is set to 1.

Response Parameters

Status code: 200

Table 5-11 Response body parameters

Parameter	Type	Description
cluster	CreateClusterClusterResponse object	Parameter description: Cluster object. For a pay-per-use cluster, only the cluster parameter is returned. Options: N/A

Table 5-12 CreateClusterClusterResponse

Parameter	Type	Description
id	String	Parameter description: Cluster ID. Options: N/A
name	String	Parameter description: Cluster name. Options: N/A

Example Requests

- Example request for creating a pay-per-use cluster.

POST https://{Endpoint}/v1.0/{project_id}/clusters

```
{
  "cluster" : {
    "instance" : {
      "flavorRef" : "ess.spec-4u16g",
      "volume" : {
        "volume_type" : "COMMON",
        "size" : 120
      },
    },
    "nics" : {
      "vpclId" : "{VPC ID}",
      "netId" : "{NET ID}",
      "securityGroupId" : "{Security group ID}"
    },
    "availability_zone" : "{AZ CODE}"
  },
  "datastore" : {
    "version" : "{cluster-version}",
    "type" : "elasticsearch"
  },
  "name" : "cluster-name",
  "desc" : "Cluster testing",
  "instanceNum" : 3,
  "backupStrategy" : {
    "period" : "16:00 GMT+08:00",
    "prefix" : "snapshot",
    "keepday" : 7,
    "frequency" : "DAY",
    "bucket" : "css-obs-backup",
    "basePath" : "css_repository/obs-path",
    "agency" : "css_obs_agency",
    "maxSnapshotBytesPerSeconds" : 40,
    "maxRestoreBytesPerSeconds" : 40
  },
  "httpsEnable" : true,
  "authorityEnable" : true,
  "adminPwd" : "{password}",
  "enterprise_project_id" : "0",
  "tags" : [ {
    "key" : "k1",
    "value" : "v1"
  }, {
    "key" : "k2",
```

```
"value" : "v2"  
  } ]  
}  
}
```

Example Responses

Status code: 200

Request succeeded.

- Example response for creating a pay-per-use cluster.

```
{  
  "cluster" : {  
    "id" : "ef683016-871e-48bc-bf93-74a29d60d214",  
    "name" : "ES-Test"  
  }  
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.1.2 Querying the Cluster List

Function

This API is used to query the cluster list and cluster status, including cluster nodes, Kibana public network access information, public network IP address, and IPv4 address and port number for private network access. After a cluster is created, check the cluster information and update the cluster in a timely manner.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters

Table 5-13 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Table 5-14 Query Parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Definition: The start value of the query. Constraints: N/A Value range: 1-1000 Default value: The default value is 1, indicating that the query starts from the first cluster.

Parameter	Mandatory	Type	Description
limit	No	Integer	Definition: Query the total number of clusters of different engine types. Constraints: N/A Value range: 1-1000 Default value: The default value is 10, indicating that 10 clusters are queried at a time.
datastoreType	No	String	Definition: Cluster engine type. Constraints: N/A Value range: • elasticsearch: Provides distributed search capabilities, log analytics, reporting, and semantic search based on open-source Elasticsearch. • logstash: Provides functions such as data ingestion, transformation, cleaning, and parsing based on open-source Logstash. • opensearch: Provides distributed search, log analytics, reporting, and semantic search based on open-source OpenSearch. It is the future version of CSS's Elasticsearch clusters. • When this parameter is left blank, all types of clusters are queried. Default value: The default value is empty.

Request Parameters

None

Response Parameters

Status code: 200

Table 5-15 Response body parameters

Parameter	Type	Description
totalSize	Integer	Definition: Number of clusters. If datastoreType is specified, the number of clusters of the specified type is displayed. Otherwise, the number of all clusters is displayed. Value range: N/A
clusters	Array of ClusterList objects	Definition: Cluster object list. Value range: N/A

Table 5-16 ClusterList

Parameter	Type	Description
datastore	Object	Definition: Cluster data search engine type. Value range: N/A
snapshotPolicy	Object	Definition: Information about an automatic snapshot creation policy. Value range: N/A
instances	Array of ClusterListInstances objects	Definition: Cluster node list. Value range: N/A
publicKibanaResp	Object	Definition: Public network access information for Kibana/Dashboards, including the IP address, whitelist, and bandwidth. Value range: N/A

Parameter	Type	Description
elbWhiteList	Object	Definition: Public network access control information. Value range: N/A
updated	String	Definition: Time when the cluster was last modified. Value range: The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
name	String	Definition: Cluster name. Value range: N/A
publicIp	String	Definition: Public IP address information. Value range: N/A
created	String	Definition: Cluster creation time. The returned cluster list is sorted by creation time in descending order. The latest cluster is displayed at the top. Value range: The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
id	String	Definition: Cluster ID. Value range: N/A.
status	String	Definition: Cluster status. Value range: • 100: creating • 200: available • 300: unavailable • 303: creation failed

Parameter	Type	Description
endpoint	String	Definition: IPv4 address and port number for accessing the cluster through a private network, for example, 10.10.20.2:9100 . Value range: N/A
vpcId	String	Definition: VPC ID. Value range: N/A
subnetId	String	Definition: Current subnet of a cluster. The subnet is used for adding nodes, including scaling, adding dedicated master/client nodes, and enabling VPC Endpoint. IP addresses will be allocated to newly added nodes from the new subnet. Value range: N/A
securityGroupId	String	Definition: Security group ID. Value range: N/A
bandwidthSize	Integer	Definition: Public network bandwidth. Unit (Mbit/s). Value range: N/A
httpsEnable	Boolean	Definition: Whether HTTPS access is enabled. Value range: • false: HTTPS access is disabled. • true: HTTPS access is enabled.

Parameter	Type	Description
authorityEnable	Boolean	Definition: Whether to enable security mode-based authentication for the cluster. Value range: • true: Enable security mode-based authentication. • false: Disable security mode-based authentication.
diskEncrypted	Boolean	Definition: Whether disk encryption is enabled. To enable it, submit a service ticket. Value range: • true: Disks are encrypted. • false: Disk are not encrypted.
backupAvailable	Boolean	Definition: Whether to enable cluster snapshots. Value range: • true: Cluster snapshots are enabled. • false: Cluster snapshots are disabled.
actionProgress	Object	Definition: Cluster behavior progress, which shows the progress of cluster creation and scaling in percentage. CREATING indicates the progress of creation. Value range: N/A

Parameter	Type	Description
actions	Array of strings	Definition: Current behavior of a cluster. Value range: • REBOOTING • GROWING • RESTORING: The cluster is being restored. • SNAPSHOTTING: A snapshot is being created. • BIND_FAILURE: failed to associate with an EIP. • SHRINKING: scaling down or in. • REPLACE_NODE: replacing nodes. • AZ_MIGRATION: switching AZs. • RESIZING_FLAVOR: changing node specifications. • MODE_CHANGING: changing the security mode. • SG_CHANGING: changing the security group. • UPGRADING: upgrading. • RESETTING_PASSWORD: The password is being reset. • CONFIGURED: The configuration has been updated and cluster restart is pending.
enterpriseProjectId	String	Definition: ID of the enterprise project that a cluster belongs. If the user of the cluster has not enabled the enterprise project, the setting of this parameter is not returned. Value range: N/A
tags	Array of ClusterListTags objects	Definition: Cluster tags. Value range: N/A

Parameter	Type	Description
failedReason	ClusterListFailedReasons object	Definition: Cluster failure cause. If the cluster status is available, this parameter is not returned. Value range: N/A
bandwidthResourceId	String	Definition: ID of the resource accessible from the public network. Value range: N/A
ipv6Endpoint	String	Definition: IPv6 address and port number accessed from the cluster's private network. Value range: N/A
cmkId	String	Definition: Encryption key ID used by the current cluster. Value range: N/A
currentSubnetIds	String	Definition: Subnets used by all nodes in the cluster. If there are multiple subnets (two at most), separate them with a comma (,). Value range: N/A
desc	String	Definition: Cluster description. Value range: It can contain a maximum of 128 characters.

Table 5-17 ClusterListDatastore

Parameter	Type	Description
type	String	Definition: Cluster engine type. Value range: - elasticsearch: Elasticsearch cluster. - logstash: Logstash cluster. - opensearch: OpenSearch cluster.
version	String	Definition: Cluster engine version. Value range: For details, see the supported versions in Before You Start .
supportSecurity-mode	Boolean	Definition: Whether the cluster supports the security mode. Value range: true: supported. false: not supported.
subVersion	String	Definition: Cluster image version. Value range: N/A
isEosCluster	Boolean	Definition: Whether the cluster image version has reached EOS. Value range: - true: yes. - false: no.

Table 5-18 SnapshotPolicyResp

Parameter	Type	Description
backupEnable	Boolean	Definition: Whether to enable automatic snapshots for the cluster. Value range: • true: Yes • false: No
bakPeriod	String	Definition: Snapshot creation time. Value range: N/A
bakFrequency	String	Definition: Snapshot creation interval. Value range: • DAY: every day • MON: every Monday • TUR: every Tuesday • WED: every Wednesday • THU: every Thursday • FRI: every Friday • SAT: every Saturday • SUN: every Sunday • HOUR: every hour
bakKeepDay	Integer	Definition: Number of retained snapshots. Value range: N/A

Table 5-19 ClusterListInstances

Parameter	Type	Description
status	String	Definition: Cluster node status. Value range: • 100: creating • 200: available • 300: unavailable • 303: creation failed
type	String	Definition: Type of the current cluster node. Value range: • ess: data node • ess-cold: cold data node • ess-client: client node • ess-master: master node • lgs: Logstash node.
id	String	Definition: Instance ID. Value range: N/A
name	String	Definition: Instance name. Value range: N/A
specCode	String	Definition: Node flavor name. Value range: N/A
azCode	String	Definition: AZ of the node. Value range: N/A
ip	String	Definition: Instance IP address. Value range: N/A

Parameter	Type	Description
volume	ClusterVolumeResp object	Definition: Instance disk information. Value range: N/A
resourceId	String	Definition: Instance resource ID. Value range: N/A
subnetId	String	Definition: Subnet ID of the current node. Value range: N/A

Table 5-20 ClusterVolumeResp

Parameter	Type	Description
type	String	Definition: Instance disk type. Value range: N/A
size	Integer	Definition: Instance disk size. Value range: N/A
resourceIds	Array of strings	Definition: Instance disk resource ID. Value range: N/A

Table 5-21 publicKibanaRespBody

Parameter	Type	Description
eipSize	Integer	Definition: The bandwidth size in Mbit/s. Value range: N/A

Parameter	Type	Description
elbWhiteListResp	kibanaElbWhiteListResp object	Definition: Kibana public network access control information. Value range: N/A
publicKibanaIp	String	Definition: IP address for accessing Kibana or Dashboards through the public network. Value range: N/A
bandwidthResourceId	String	Definition: Resource ID for Kibana/Dashboards public network access. Value range: N/A

Table 5-22 kibanaElbWhiteListResp

Parameter	Type	Description
enableWhiteList	Boolean	Definition: Whether to enable the whitelist that controls public network access for Kibana or OpenSearch Dashboards. Value range: • true: Enable the access control whitelist. • false: Disable the access control whitelist.
whiteList	String	Definition: Whitelist for Kibana public network access. Value range: N/A

Table 5-23 elbWhiteListResp

Parameter	Type	Description
enableWhiteList	Boolean	Definition: Whether to enable public network access control. If a whitelist is set, only IP addresses on the whitelist can access the cluster. If no whitelist is set, all IP addresses can access the cluster. Value range: • true: Public network access control is enabled. • false: Public network access control is disabled.
whiteList	String	Definition: Whitelist for public network access. Value range: IP address.

Table 5-24 ClusterListTags

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

Table 5-25 ClusterListFailedReasons

Parameter	Type	Description
errorCode	String	Definition: Error code. Value range: • See Error Codes.

Parameter	Type	Description
errorMsg	String	Definition: Error details. Value range: N/A

Example Requests

- Querying the First Two Clusters

- Method 1: GET https://{Endpoint}/v1.0/{project_id}/clusters?start=1&limit=2
- Method 2: GET https://{Endpoint}/v1.0/{project_id}/clusters?limit=2

- Querying the First 10 Clusters

- Method 1: GET https://{Endpoint}/v1.0/{project_id}/clusters?start=1&limit=10
- Method 2: GET https://{Endpoint}/v1.0/{project_id}/clusters?limit=10

Example Responses

Status code: 200

Request succeeded. This is only an example. The responses to the four requests are similar. In the example, there is only one cluster list.

```
{
  "totalSize": 1,
  "clusters": [
    {
      "datastore": {
        "type": "elasticsearch",
        "version": "7.10.2",
        "subVersion": "7.10.2_24.3.0_0827",
        "isEosCluster": false,
        "supportSecuritymode": false
      },
      "instances": [
        {
          "status": "200",
          "type": "ess",
          "id": "{INSTANCE_ID}",
          "name": "css-8bc5-ess-esn-1-1",
          "specCode": "ess.spec-4u8g",
          "azCode": "{AZ_CODE}",
          "volume": {
            "type": "ULTRAHIGH",
            "size": 40,
            "resourceIds": [
              "{RESOURCE_ID}"
            ]
          },
          "ip": "192.168.0.122",
          "resourceId": "{RESOURCE_ID}",
          "subnetId": "{SUBNET_ID}"
        }
      ],
      "publicKibanaResp": {
        "eipSize": 10,
        "publicKibanaIp": "100.95.152.28:9200",
        "bandwidthResourceId": "18bec13f-5cc1-4631-867f-33505d15be12"
      },
      "elbWhiteList": {
```

```
{
  "whiteList": "",
  "enableWhiteList": false
},
"updated": "2023-10-09T02:07:13",
"name": "css-8bc5",
"desc": "csscluster"
"snapshotPolicy": {
  "backupEnable": false,
  "bakPeriod": "",
  "bakFrequency": "",
  "bakKeepDay": 1
},
"publicIp": "100.85.222.202",
"created": "2023-10-09T02:07:13",
"id": "{CLUSTER_ID}",
"status": "200",
"endpoint": "192.168.0.122:9200",
"vpclId": "{VPC_ID}",
"subnetId": "{SUBNET_ID}",
"currentSubnetIds": "{SUBNET_ID}",
"securityGroupId": "{SECURITY_GROUP_ID}",
"bandwidthResourceId": "{BANDWIDTH_RESOURCE_ID}",
"bandwidthSize": 3,
"httpsEnable": true,
"authorityEnable": true,
"diskEncrypted": false,
"backupAvailable": false,
"actionProgress": {},
"actions": [],
"enterpriseProjectId": "0",
"tags": [],
"period": true
}
]
```

Status Codes

Status Code	Description
200	Request succeeded. This is only an example. The responses to the four requests are similar. In the example, there is only one cluster list.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.1.3 Querying Cluster Details

Function

This API is used to query and display details about a single cluster, including the cluster's search type, node list, network information, cluster name, cluster ID, and

cluster status, so that you can check whether the created cluster meets your expectation. You can also log in to the Huawei Cloud official website, access the CSS console, and obtain the cluster's private IP address, public IP address, version, and node information on the cluster overview page.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}

Table 5-26 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-27 Response body parameters

Parameter	Type	Description
datastore	ClusterDetailDatastore object	Definition: Cluster data search engine type. Value range: N/A
instances	Array of ClusterDetailInstances objects	Definition: List of node objects. Value range: N/A
publicKibanaResp	publicKibanaRespBody object	Definition: Public network access information for Kibana/Dashboards, including the IP address, whitelist, and bandwidth. Value range: N/A
elbWhiteList	elbWhiteListResp object	Definition: Public network access control information. Value range: N/A
updated	String	Definition: Time when the cluster was last modified. Value range: The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
name	String	Definition: Cluster name. Value range: N/A
publicip	String	Definition: Public IP address and port number. Value range: N/A

Parameter	Type	Description
snapshotPolicy	SnapshotPolicyResp object	Definition: Information about an automatic snapshot creation policy. Value range: N/A
created	String	Definition: Time when a cluster was created. Value range: The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
id	String	Definition: Cluster ID. Value range: N/A
status	String	Definition: Cluster status. Value range: 100: An operation, such as creation, is in progress. • 200: available • 300: unavailable • 303: creation failed
endpoint	String	Definition: IPv4 address and port number accessed from the cluster's private network. Value range: N/A
vpclId	String	Definition: VPC ID. Value range: N/A

Parameter	Type	Description
subnetId	String	Definition: Current subnet of a cluster. This subnet is used to add nodes, including scaling, adding dedicated master/client nodes, and enabling the VPC Endpoint service. IP addresses will be allocated to newly added nodes from the current subnet. Value range: N/A
securityGroupId	String	Definition: Security group ID. Value range: N/A
vpceIp	String	Definition: VPC endpoint IP address. Value range: N/A
vpceIpv6Address	String	Definition: IPv6 address of the VPC endpoint. Value range: N/A
bandwidthSize	Integer	Definition: Public network bandwidth, in Mbit/s. Value range: N/A
httpsEnable	Boolean	Parameter description: Communication encryption status. Options: • false: Communication encryption is not enabled. • true: Communication encryption is enabled.
diskEncrypted	Boolean	Parameter description: Whether disks are encrypted. Options: • true: Disks are encrypted. • false: Disk are not encrypted.

Parameter	Type	Description
authorityEnable	Boolean	Definition: Whether to enable security mode-based authentication for the cluster. Value range: • true: Enable security mode-based authentication. • false: Disable security mode-based authentication.
backupAvailable	Boolean	Definition: Whether snapshot is enabled. Value range: • true: The snapshot function is enabled. • false: The snapshot function is disabled.
actionProgress	Object	Definition: Cluster behavior progress, which shows the progress of cluster creation and scaling in percentage. Value range: N/A

Parameter	Type	Description
actions	Array of strings	Definition: List of ongoing tasks in the cluster. Value range: • REBOOTING • GROWING • RESTORING: The cluster is being restored. • SNAPSHOTTING: A snapshot is being created. • BIND_FAILURE: failed to associate with an EIP. • SHRINKING: scaling down or in. • REPLACE_NODE: replacing nodes. • AZ_MIGRATION: switching AZs. • RESIZING_FLAVOR: changing node specifications. • MODE_CHANGING: changing the security mode. • SG_CHANGING: changing the security group. • UPGRADING: upgrading. • RESETTING_PASSWORD: The password is being reset. • CONFIGURED: The configuration has been updated and cluster restart is pending.
enterpriseProjectId	String	Definition: ID of the enterprise project that a cluster belongs. If the user of the cluster has not enabled the enterprise project, the setting of this parameter is not returned. Value range: N/A
tags	Array of ClusterDetailTags objects	Definition: Cluster tags. Value range: N/A

Parameter	Type	Description
failedReason	ClusterDetailFailedReasons object	Definition: Failure cause of a cluster change task. If the cluster status is available, this parameter is not returned. Value range: N/A
bandwidthResourceId	String	Definition: ID of the resource accessible from the public network. Value range: N/A
ipv6Endpoint	String	Definition: IPv6 address and port number accessed from the cluster's private network. Value range: N/A
supportVpcep	Boolean	Definition: Whether the VPCEP service can be enabled for the current cluster. Value range: • true: supported. • false: not supported.
cmkId	String	Definition: Encryption key ID used by the current cluster. Value range: N/A
currentSubnetIds	String	Definition: Subnets used by all nodes in the cluster. If there are multiple subnets (two at most), separate them with a comma (,). Value range: N/A

Parameter	Type	Description
desc	String	Definition: Cluster description. Value range: It can contain a maximum of 128 characters.

Table 5-28 ClusterDetailDatastore

Parameter	Type	Description
type	String	Definition: Engine type. Value range: - elasticsearch: Provides distributed search capabilities, log analytics, reporting, and semantic search based on open-source Elasticsearch. - logstash: provides distributed search, log analytics, reporting, and semantic search based on open-source OpenSearch. - opensearch: Provides functions such as data ingestion, transformation, cleaning, and parsing based on open-source Logstash.
version	String	Definition: Version of the CSS cluster engine. For details, see the supported versions in Before You Start . Value range: N/A
subVersion	String	Definition: Cluster image version. Value range: N/A
isEosCluster	Boolean	Definition: Whether the cluster image version has reached EOS. Value range: - true: yes. - false: no.

Parameter	Type	Description
supportSecurity-mode	Boolean	Definition: Whether the cluster supports the security mode. Value range: • true: supported. • false: not supported.

Table 5-29 ClusterDetailInstances

Parameter	Type	Description
status	String	Definition: Node status. Value range: • 100: creating • 200: available • 300: unavailable • 303: creation failed
resourceId	String	Definition: Instance resource ID. Value range: N/A
type	String	Definition: Node type. Value range: • ess: data node • ess-cold: cold data node • ess-client: client node • ess-master: master node • lgs: Logstash node
id	String	Definition: Instance ID. Value range: N/A
name	String	Definition: Instance name. Value range: N/A

Parameter	Type	Description
specCode	String	Definition: Node flavor name. Value range: N/A
azCode	String	Definition: AZ of the node. Value range: N/A
ip	String	Definition: Instance IP address. Value range: N/A
volume	ShowClusterVolumeRsp object	Definition: Instance disk information. Value range: N/A
subnetId	String	Definition: Subnet ID of the current node. Value range: N/A

Table 5-30 ShowClusterVolumeRsp

Parameter	Type	Description
type	String	Definition: Instance disk type. Value range: N/A
size	Integer	Definition: Instance disk size. Value range: N/A
resourceIds	Array of strings	Definition: Instance disk resource ID. Value range: N/A

Table 5-31 publicKibanaRespBody

Parameter	Type	Description
eipSize	Integer	Definition: The bandwidth size in Mbit/s. Value range: N/A
elbWhiteListResp	kibanaElbWhiteListResp object	Definition: Kibana public network access control information. Value range: N/A
publicKibanaIp	String	Definition: IP address for accessing Kibana or Dashboards through the public network. Value range: N/A
bandwidthResourceId	String	Definition: Resource ID for Kibana/Dashboards public network access. Value range: N/A

Table 5-32 kibanaElbWhiteListResp

Parameter	Type	Description
enableWhiteList	Boolean	Definition: Whether to enable the whitelist that controls public network access for Kibana or OpenSearch Dashboards. Value range: • true: Enable the access control whitelist. • false: Disable the access control whitelist.
whiteList	String	Definition: Whitelist for Kibana public network access. Value range: N/A

Table 5-33 elbWhiteListResp

Parameter	Type	Description
enableWhiteList	Boolean	Definition: Whether to enable public network access control. If a whitelist is set, only IP addresses on the whitelist can access the cluster. If no whitelist is set, all IP addresses can access the cluster. Value range: • true: Public network access control is enabled. • false: Public network access control is disabled.
whiteList	String	Definition: Whitelist for public network access. Value range: IP address.

Table 5-34 SnapshotPolicyResp

Parameter	Type	Description
backupEnable	Boolean	Definition: Whether to enable automatic snapshots for the cluster. Value range: • true: Yes • false: No
bakPeriod	String	Definition: Snapshot creation time. Value range: N/A

Parameter	Type	Description
bakFrequency	String	Definition: Snapshot creation interval. Value range: • DAY: every day • MON: every Monday • TUR: every Tuesday • WED: every Wednesday • THU: every Thursday • FRI: every Friday • SAT: every Saturday • SUN: every Sunday • HOUR: every hour
bakKeepDay	Integer	Definition: Number of retained snapshots. Value range: N/A

Table 5-35 ClusterDetailTags

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

Table 5-36 ClusterDetailFailedReasons

Parameter	Type	Description
errorCode	String	Parameter description: Error code. Options: • CSS.6000: A cluster fails to be created. • CSS.6001: A cluster fails to be scaled out. • CSS.6002: A cluster fails to be restarted. • CSS.6004: A node fails to be created in a cluster. • CSS.6005: A service fails to be initialized.
errorMsg	String	Definition: Error details. Value range: N/A

Example Requests

Whether the details of a cluster are obtained

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/5c2bab0d-ec5f-491f-9418-b8da38a26760
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "datastore": {
    "type": "elasticsearch",
    "version": "7.10.2",
    "subVersion": "7.10.2_24.3.0_0827",
    "isEosCluster": false,
    "supportSecuritymode": false
  },
  "instances": [
    {
      "status": "200",
      "type": "ess",
      "id": "{INSTANCE_ID}",
      "name": "css-8bc5-ess-esn-1-1",
      "specCode": "ess.spec-4u8g",
      "azCode": "{AZ_CODE}",
      "volume": {
        "type": "ULTRAHIGH",
        "size": 40,
        "resourceIds": [
```

```
    "{RESOURCE_ID}"
  ]
},
"ip": "192.168.0.122",
"resourceId": "{RESOURCE_ID}",
"subnetId": "{SUBNET_ID}" //Subnet ID used by the node
}
],
"publicKibanaResp": {
  "eipSize": 10,
  "publicKibanaIp": "100.95.152.28",
  "elbWhiteListResp": null,
  "bandwidthResourceId": "18bec13f-5cc1-4631-867f-33505d15be12"
},
"elbWhiteList": {
  "whiteList": "",
  "enableWhiteList": false
},
"updated": "2023-10-09T02:07:13",
"name": "css-8bc5",
"desc": "csscluster"
"snapshotPolicy": {
  "backupEnable": false,
  "bakPeriod": "",
  "bakFrequency": "",
  "bakKeepDay": 1
},
"publicIp": "100.85.222.202",
"created": "2023-10-09T02:07:13",
"id": "{CLUSTER_ID}",
"status": "200",
"endpoint": "192.168.0.122:9200",
"vpcId": "{VPC_ID}", //VPC ID used by the cluster
"subnetId": "{SUBNET_ID}", //Subnet ID used by all nodes in the cluster. Typically, a cluster uses only one
subnet.
"securityGroupId": "{SECURITY_GROUP_ID}",
"bandwidthResourceId": "{BANDWIDTH_RESOURCE_ID}",
"bandwidthSize": 3,
"httpsEnable": true,
"authorityEnable": true,
"diskEncrypted": false,
"backupAvailable": false,
"actionProgress": {},
"actions": [],
"enterpriseProjectId": "0",
"tags": [],
"period": true,
"supportVpcep": true,
"currentSubnetIds": "{SUBNET_ID}" //Default subnet ID of the cluster, which is used to assign new IP
addresses to newly added nodes during a cluster scale-out operation when a VPC endpoint is created.
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.1.4 Deleting a Cluster

Function

This API is used to delete a cluster. All resources of the deleted cluster, including customer data, will be released. If you want to retain the data in a customer cluster, create a snapshot before deleting the cluster.

NOTE

This API can also be used to unsubscribe from yearly/monthly clusters. Clusters frozen for public security reasons cannot be deleted. Deleting a cluster will also clear its data. Exercise caution.

Calling Method

For details, see [Calling APIs](#).

URI

DELETE /v1.0/{project_id}/clusters/{cluster_id}

Table 5-37 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster you want to delete. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Delete a cluster.

DELETE https://{Endpoint}/v1.0/{project_id}/clusters/2a197c4d-5467-4003-931d-83ec49939cf

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.1.5 Modifying the Name or Description of a Cluster

Function

This API is used to change the cluster name or description.

Constraints

The name change here does not affect the cluster's underlying name.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/changenname

Table 5-38 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: Cluster ID. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-39 Request body parameters

Parameter	Mandatory	Type	Description
display_name	Yes	String	Definition: New cluster name. Constraints: display_name and desc cannot both be null at the same time. Value range: The cluster name must start with a letter and can contain 4 to 32 characters. Only letters, digits, periods (.), hyphens (-), and underscores (_) are allowed. Default value: N/A
desc	No	String	Definition: Cluster description. Constraints: The value can contain up to 128 characters. display_name and desc cannot be null at the same time. Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Change the display name of the current cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/changename
{
  "display_name": "ES-Test-new",
```

```
"desc" : "test1"  
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.1.6 Changing the Cluster Administrator Password

Function

This API is used to change the admin password for a security-mode cluster.

A security-mode cluster is a cluster for which the security mode is enabled. This parameter is available only for Elasticsearch and OpenSearch clusters.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/password/reset

Table 5-40 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose password you want to change. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-41 Request body parameters

Parameter	Mandatory	Type	Description
newpassword	Yes	String	Definition: New password of the cluster admin account. Constraints: This API is available only when authorityEnable is set to true during cluster creation. Value range: • Must contain 8 to 32 characters. • Must contain at least three of the following character types: letters, digits, and special characters ~!@#\$%^&*()-_+=~\ []{};,:<.>/? • Do not use the administrator name, or the administrator name spelled backwards, as the password. Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Change the administrator password of the current cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/password/reset
{
  "newpassword" : "xxxxxx"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.1.7 Restarting a Cluster (Deprecated)

Function

This API is used to restart the data nodes of an Elasticsearch or OpenSearch cluster. Services will be interrupted during the restart. To restart all nodes or other types of nodes in a cluster, you are advised to use [Restarting a Cluster V2](#).

Constraints

If a cluster is available, ensure that it has stopped processing service requests, such as importing and searching for data. Otherwise, data may be lost when the cluster is restarted.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/restart

Table 5-42 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster you want to restart. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Restart the data nodes of a cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/47e49a5e-8ced-4d0d-ae15-2af62ac468e3/restart
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.1.8 Query the disk usage of a cluster.

Function

Query the disk usage of a cluster.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/volume

Table 5-43 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details, see Obtaining a Project ID and Name . Constraints: N/A Value range: For details, see Obtaining a Project ID and Name . Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster to be updated. For details about how to obtain the cluster ID, see Obtaining a Cluster ID . Constraints: N/A Value range: To obtain the value, see Obtaining the Cluster ID . Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-44 Response body parameters

Parameter	Type	Description
diskInfoList	Array of DiskInfo objects	Definition: Disk usage of each node. Value range: N/A

Table 5-45 DiskInfo

Parameter	Type	Description
id	String	Definition: Node ID. Value range: N/A

Parameter	Type	Description
name	String	Definition: Node name. Value range: N/A
diskCapacity	String	Definition: Total node storage capacity. Value range: N/A
diskUsed	String	Definition: Used node storage capacity. Value range: N/A
percentage	String	Definition: Node storage usage. Value range: N/A

Example Requests

Query the disk usage of each node in a cluster.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/{cluster_id}/volume
```

Example Responses

Status code: 200

Returned value for querying the disk usage of each node in a cluster.

```
{
  "diskInfoList": [ {
    "percentage": "5.16",
    "diskCapacity": "97.86",
    "diskUsed": "5.05",
    "id": "700915c3-15da-455f-9a6e-8b853b8665ac",
    "name": "css-9062-ess-esn-1-1"
  } ]
}
```

Status Codes

Status Code	Description
200	Returned value for querying the disk usage of each node in a cluster.

Error Codes

See [Error Codes](#).

5.1.9 Downloading a Security Certificate

Function

The certificate obtained through this API can be used for HTTPS-based secure cluster connections, where encryption and user authentication occur.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/cer/download

Table 5-46 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-47 Response body parameters

Parameter	Type	Description
file	String	Definition: File stream Default value: N/A

Example Requests

Download the security certificate.

```
GET https://{Endpoint}/v1.0/{project_id}/cer/download
```

Example Responses

Status code: 200

Request succeeded.

```
"Certificate:xxx"
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.1.10 Upgrading a Cluster Kernel

Function

This API is used to upgrade your cluster to incorporate feature enhancements or fixes to known issues.

The following upgrade types are supported:

- During a same-version upgrade, kernel patches are updated for a cluster. The cluster is upgraded to the latest image of the current version to fix known

issues or optimize performance. For example, if the current cluster version is 7.10.2(7.10.2_24.3.3_0102), upon a same-version upgrade, the cluster will be upgraded to the latest image 7.10.2(7.10.2_24.3.4_0109) of version 7.10.2. (The version numbers used here are examples only.)

- Cross-version upgrade: Upgrade the cluster version. The cluster will be upgraded to the latest image of the target version to enhance functions or incorporate versions. For example, if the current cluster version is 7.6.2(7.6.2_24.3.3_1224), upon a cross-version upgrade, the cluster will be upgraded to the latest image 7.10.2(7.10.2_24.3.4_0109) of version 7.10.2. (The version numbers used here are examples only.)
- Cross-engine upgrade: Change an Elasticsearch cluster to an OpenSearch cluster. Cross-engine upgrade means to upgrade an Elasticsearch cluster to the latest image of the target OpenSearch version. For example, if the Elasticsearch cluster version is 7.10.2(7.10.2_24.3.3_0102), upon a cross-engine upgrade, the Elasticsearch cluster will be upgraded to the latest image 1.3.6(1.3.6_24.3.4_0109) of OpenSearch 1.3.6. (The version numbers used here are examples only.) This feature is not supported currently.

Constraints

- Only Elasticsearch and OpenSearch clusters support this option.
- A maximum of 20 clusters can be upgraded at the same time.
- During an upgrade, the data on the upgraded nodes needs to be migrated to other nodes. The timeout for data migration per node is 48 hours. Upgrade will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/inst-type/{inst_type}/image/upgrade

Table 5-48 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be upgraded. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
inst_type	Yes	String	Definition: Types of nodes to be upgraded. Constraints: N/A Value range: Currently, only all is supported, indicating all nodes. Default value: N/A

Request Parameters

Table 5-49 Request body parameters

Parameter	Mandatory	Type	Description
target_image_id	Yes	String	Definition: ID of the target image version. For how to obtain this value, see Obtaining the Target Image ID . Constraints: N/A Value range: N/A Default value: N/A
upgrade_type	Yes	String	Definition: Upgrade type. Constraints: N/A Value range: • same: same-version upgrade. • cross: cross-version upgrade. • cross-engine: cross-engine upgrade. Default value: N/A

Parameter	Mandatory	Type	Description
indices_backup_check	Yes	Boolean	Definition: Whether to check full index snapshots. Snapshots help prevent potential data loss caused by upgrade failures. Constraints: CSS cannot check the content or backup times of snapshots. You should manually check existing snapshots. If any of them is over one month old, create the latest snapshot. Value range: • true: Verify the load. • false: Do not verify the load. Default value: true
agency	Yes	String	Definition: Agency name. After a node is upgraded, NICs need to be reattached to it. To do that, you must have the permission to access VPC resources. By configuring an IAM agency, you can authorize CSS to access its VPC resources through an associated account. Constraints: N/A Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
cluster_load_check	No	Boolean	Definition: Whether to check the cluster load. and reduce the likelihood of a cluster upgrade failure caused by an overload. The check items are as follows: • nodes.thread_pool.search.queue < 1000: Check that the maximum number of requests in the search queue is less than 1000. • nodes.thread_pool.write.queue < 200: Check that the maximum number of requests in the write queue is less than 200. • nodes.process.cpu.percent < 90: Check that the maximum CPU usage of cluster nodes is less than 90%. • nodes.os.cpu.load_average/ Number of vCPUs < 80%: Check that the number of running processes plus the number of processes waiting for CPUs is less than 80/ %of the total number of vCPUs. • If any of the results is abnormal, wait until the load drops or actively optimize it before performing the upgrade. Constraints: N/A Value range: • true: The platform evaluates the cluster load. • false: The platform does not evaluate the cluster load. Default value: true

Parameter	Mandatory	Type	Description
batch_size	No	Integer	Definition: Set data migration concurrency control. Increasing the data migration concurrency can accelerate the upgrade process, but faster data migration leads to higher I/O usage. A higher concurrency is likely to lead to higher cluster load, which may impact cluster performance. Constraints: N/A Value range: You are advised to retain the default value 1. The value should not exceed half of the number of data nodes. Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Perform a same-version upgrade.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/inst-type/all/image/upgrade
{
  "target_image_id" : "{target_image_id}",
  "upgrade_type" : "same",
  "indices_backup_check" : true,
  "agency" : "css-test-agency",
  "cluster_load_check" : true
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should not repeat the request without modifications.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.1.11 Obtaining the ID of the Target Image to Upgrade To

Function

This API is used to obtain the IDs of target images the current cluster can be upgraded to. The upgrade types are as follows:

- During a same-version upgrade, kernel patches are updated for a cluster. The cluster is upgraded to the latest image of the current version to fix known issues or optimize performance. For example, if the current cluster version is 7.10.2(7.10.2_24.3.3_0102), upon a same-version upgrade, the cluster will be upgraded to the latest image 7.10.2(7.10.2_24.3.4_0109) of version 7.10.2. (The version numbers used here are examples only.)
- Cross-version upgrade: Upgrade the cluster version. The cluster will be upgraded to the latest image of the target version to enhance functions or incorporate versions. For example, if the current cluster version is 7.6.2(7.6.2_24.3.3_1224), upon a cross-version upgrade, the cluster will be upgraded to the latest image 7.10.2(7.10.2_24.3.4_0109) of version 7.10.2. (The version numbers used here are examples only.)
- Cross-engine upgrade: Change an Elasticsearch cluster to an OpenSearch cluster. Cross-engine upgrade means to upgrade an Elasticsearch cluster to the latest image of the target OpenSearch version. For example, if the Elasticsearch cluster version is 7.10.2(7.10.2_24.3.3_0102), upon a cross-engine upgrade, the Elasticsearch cluster will be upgraded to the latest image 1.3.6(1.3.6_24.3.4_0109) of OpenSearch 1.3.6. (The version numbers used here are examples only.) This feature is not supported currently.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/target/{upgrade_type}/images

Table 5-50 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be upgraded. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
upgrade_type	Yes	String	Definition: Cluster/image upgrade type. Constraints: N/A Value range: • same: same-version upgrade. • cross: cross-version upgrade. • cross-engine: cross-engine upgrade. Default value: N/A

Table 5-51 Query Parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Definition: Start value of the query. The default value is 0. Constraints: N/A Value range: 0-1000 Default value: 0
limit	No	Integer	Definition: Number of images to be queried. The default value is 10 , indicating that 10 images are queried at a time. Constraints: N/A Value range: 1-1000 Default value: 10

Request Parameters

None

Response Parameters

Status code: 200

Table 5-52 Response body parameters

Parameter	Type	Description
needUploadUpgradePlugin	Boolean	Definition: Whether to upload the plug-in of the target version. Value range: true: Upload the plug-ins of the target version. false: Do not upload the plug-ins of the target version.

Parameter	Type	Description
imageInfoList	Array of GetTargetImageIdDetail objects	Definition: Image details. Value range: N/A
totalSize	Integer	Definition: Number of target images. Value range: N/A

Table 5-53 GetTargetImageIdDetail

Parameter	Type	Description
id	String	Definition: ID of an image that can be upgraded. Value range: N/A
displayName	String	Definition: Name of an image that can be upgraded. Value range: N/A
imageDesc	String	Definition: Image description. Value range: N/A

Parameter	Type	Description
datastoreType	String	Definition: Cluster engine type. Value range: - elasticsearch: Cloud Search Service (CSS) provides distributed search capabilities, log analytics, reporting, and semantic search based on open-source Elasticsearch. - logstash: Provides functions such as data ingestion, transformation, cleaning, and parsing based on open-source Logstash. - opensearch: Provides distributed search, log analytics, reporting, and semantic search based on open-source OpenSearch. It is the future version of CSS's Elasticsearch clusters.
datastoreVersion	String	Definition: Image engine version. Value range: N/A
priority	Integer	Definition: Priority of the target image. A larger value indicates a higher priority. Value range: N/A

Example Requests

Obtain the target image for a cross-version upgrade of the cluster.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/target/cross/images
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "needUploadUpgradePlugin" : false,
  "imageInfoList" : [ ],
  "totalSize" : 0
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should not repeat the request without modifications.
403	Request rejected. The server has received the request and understood it, but the server refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.1.12 Obtaining Details About a Cluster Upgrade or AZ Switchover

Function

By deploying a CSS cluster across multiple AZs, you can increase the cluster's availability, lower the likelihood of data loss, and minimize service downtime. You can deploy a cluster across two or three AZs within a region. The upgrade may take a long time. This API can be used to check phase-by-phase information during a cluster upgrade or AZ switchover.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/upgrade/detail

Table 5-54 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be upgraded. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-55 Query Parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Definition: The start value of the query. The default value is 1, indicating that the query starts from the first task. Constraints: N/A Value range: 1-1000 Default value: 1

Parameter	Mandatory	Type	Description
limit	No	Integer	Definition: Number of AZs to be queried. The default value is 10 , indicating that 10 AZs are queried at a time. Constraints: N/A Value range: 1-1000 Default value: 10
action_mode	No	String	Definition: Query details about a cluster upgrade or AZ change operation. Constraints: N/A Value range: ● UPGRADING: query cluster upgrade details. ● AZ_MIGRATION: query AZ change details. Default value: UPGRADING

Request Parameters

None

Response Parameters

Status code: 200

Table 5-56 Response body parameters

Parameter	Type	Description
totalSize	Integer	Definition: Total number of cluster upgrade or AZ change tasks. Value range: N/A

Parameter	Type	Description
detailList	Array of GetUpgradeDetailInfo objects	Definition: Details of cluster upgrade or AZ change tasks. Value range: N/A

Table 5-57 GetUpgradeDetailInfo

Parameter	Type	Description
id	String	Definition: ID of the upgrade or AZ change task. Value range: N/A
startTime	String	Definition: Start time of an upgrade or AZ change task. Value range: N/A
endTime	String	Definition: End time of an upgrade or AZ change task. Value range: N/A
status	String	Definition: Status of an upgrade or AZ change task. Value range: • RUNNING: The task is being executed. • SUCCESS: The task is successful. • FAILED: The job failed. • PARTIAL_FAILED: The task failed partially.

Parameter	Type	Description
agencyName	String	Definition: Agency name. You can create an agency to allow CSS to call other cloud services. Value range: N/A
imageInfo	GetTargetImageIdDetail object	Definition: Image details. Value range: N/A
totalNodes	String	Definition: Names of all nodes that need to be upgraded or switched to another AZ. Value range: N/A
completedNodes	String	Definition: Nodes where upgrade or AZ change tasks have been completed. Value range: N/A
currentNodeName	String	Definition: Name of the node that is being upgraded or switched to another AZ. Value range: N/A
executeTimes	String	Definition: Number of times the upgrade or AZ change task was retried. Value range: N/A
migrateParam	String	Definition: This parameter is available only when action_mode is set to AZ_MIGRATION . It indicates the request parameter of the AZ change task. Value range: N/A

Parameter	Type	Description
finalAzInfoMap	String	Definition: This parameter is available only when action_mode is set to AZ_MIGRATION . It indicates the AZ details of each node after AZ switchover. Value range: N/A
currentNodeDetail	Array of CurrentNodeDetail objects	Definition: Task details of the node that is being upgraded or switched to another AZ. Value range: N/A
batchSize	Integer	Definition: Number of nodes where data is being migrated concurrently. Value range: N/A
currentBatchNodes	String	Definition: Nodes where data is being migrated. Value range: N/A

Table 5-58 GetTargetImageIdDetail

Parameter	Type	Description
id	String	Definition: ID of an image that can be upgraded. Value range: N/A
displayName	String	Definition: Name of an image that can be upgraded. Value range: N/A

Parameter	Type	Description
imageDesc	String	Definition: Image description. Value range: N/A
datastoreType	String	Definition: Cluster engine type. Value range: • elasticsearch: Cloud Search Service (CSS) provides distributed search capabilities, log analytics, reporting, and semantic search based on open-source Elasticsearch. • logstash: Provides functions such as data ingestion, transformation, cleaning, and parsing based on open-source Logstash. • opensearch: Provides distributed search, log analytics, reporting, and semantic search based on open-source OpenSearch. It is the future version of CSS's Elasticsearch clusters.
datastoreVersion	String	Definition: Image engine version. Value range: N/A
priority	Integer	Definition: Priority of the target image. A larger value indicates a higher priority. Value range: N/A

Table 5-59 CurrentNodeDetail

Parameter	Type	Description
order	Integer	Definition: SN of the upgrade or AZ switchover task. Value range: N/A

Parameter	Type	Description
name	String	Definition: Name of the upgrade or AZ switchover task. Value range: N/A
status	String	Definition: Status of the upgrade or AZ change task of the current cluster. Value range: • RUNNING: running • SUCCESS • FAILED • WAITING
desc	String	Definition: Description of the upgrade or AZ change task of the current cluster. Value range: N/A
beginTime	String	Definition: Start time of the upgrade or AZ change task of the current cluster. Value range: N/A
endTime	String	Definition: End time of the upgrade or AZ change task of the current cluster. Value range: N/A

Example Requests

- Query cluster upgrade details.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/upgrade/detail
```

- Query details about a cluster AZ change task.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/upgrade/detail?action_mode=AZ_MIGRATION
```

Example Responses

Status code: 200

Request succeeded.

- Example response for querying cluster upgrade details.

```
{
  "totalSize" : 1,
  "detailList" : [ {
    "id" : "b7ac4c5b-3bda-4feb-a303-eb80f4bce986",
    "startTime" : "2023-01-05T02:23:39",
    "endTime" : "",
    "status" : "RUNNING",
    "agencyName" : "css_test_agency",
    "imageInfo" : {
      "id" : "439b5d30-5968-45df-b088-d030a858522d",
      "displayName" : "7.10.2_22.5.1_1230",
      "imageDesc" : "The latest image of version 7.10.2 will be officially released on December 30, 2022. the stability of the cluster has been optimized. It is recommended that you upgrade to this version.",
      "datastoreType" : "elasticsearch",
      "datastoreVersion" : "7.10.2",
      "priority" : 16
    },
    "totalNodes" : "css-test-ess-esn-1-1,css-test-ess-esn-2-1,css-test-ess-esn-3-1",
    "completedNodes" : "",
    "currentNodeName" : "css-test-ess-esn-1-1",
    "executeTimes" : "1",
    "currentNodeDetail" : [ {
      "order" : 0,
      "name" : "Data migration",
      "status" : "SUCCESS",
      "desc" : "Data is migrated from a node to other nodes in the cluster by running the exclude command.",
      "beginTime" : "2023-01-05T02:23:42",
      "endTime" : "2023-01-05T02:29:51"
    }, {
      "order" : 12,
      "name" : "Task status update",
      "status" : "WAITING",
      "desc" : "The cluster task status is refreshed. If there are no nodes to be replaced, the task is marked as completed. Otherwise, another node starts to be replaced.",
      "beginTime" : "",
      "endTime" : ""
    } ]
  } ]
}
```

- Example response for querying cluster AZ change details.

```
{
  "totalSize" : 1,
  "detailList" : [ {
    "id" : "8ebe958b-b8c6-4939-b5a7-85aee9468888",
    "startTime" : "2022-12-29T08:32:29",
    "endTime" : "",
    "status" : "RUNNING",
    "agencyName" : "css_test_agency",
    "totalNodes" : "css-test-ess-esn-2-1,css-test-ess-esn-1-1,css-test-ess-esn-3-1",
    "completedNodes" : "css-test-ess-esn-2-1",
    "currentNodeName" : "css-test-ess-esn-1-1",
    "executeTimes" : "1",
    "migrateParam" : "{\"instType\":\"ess\",\"migrateType\":\"az_migrate\",\"sourceAz\":\"cn-north-4a\",\"targetAz\":\"cn-north-4c\"}",
    "finalAzInfoMap" : "{\"cn-north-4c\":{\"css-test-ess-esn-2-1,css-test-ess-esn-1-1,css-test-ess-esn-3-1\"}}",
    "currentNodeDetail" : [ {
      "order" : 0,
      "name" : "Data migration",
      "status" : "SUCCESS",
      "desc" : "Data is migrated from a node to other nodes in the cluster by running the exclude command.",
      "beginTime" : "2022-12-29T08:41:05",

```

```
"endTime" : "2022-12-29T08:41:23"
}, {
  "order" : 12,
  "name" : "Task status update",
  "status" : "RUNNING",
  "desc" : "The cluster task status is refreshed. If there are no nodes to be replaced, the task is marked as completed. Otherwise, another node starts to be replaced.",
  "beginTime" : "2023-01-04T06:53:42",
  "endTime" : ""
}]
}]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should not repeat the request without modifications.
403	Request rejected. The server has received the request and understood it, but the server refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.1.13 Retrying or Terminating a Failed Upgrade or AZ Switchover Task

Function

The upgrade or AZ change may take a long time and may fail due to network issues. This API can be used to retry or stop a failed task.

Constraints

- To call this API, the task status must be failed.
- During a cross-version upgrade, if Task Status is Failed, you can terminate the upgrade task only if all nodes have not been upgraded.
- During an AZ switchover, if Task Status is Failed, you can terminate the upgrade task only if none of the nodes has been switched.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/upgrade/{action_id}/retry

Table 5-60 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster for which you want to retry the failed task. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
action_id	Yes	String	Definition: ID of the task to be retried. For details about how to obtain the cluster ID, see Obtaining Details About a Cluster Upgrade or AZ Switchover . Constraints: ID of a failed task. Value range: N/A Default value: N/A

Table 5-61 Query Parameters

Parameter	Mandatory	Type	Description
retry_mode	No	String	Definition: Select a retry mode. When this parameter is set, the upgrade task will be terminated. When it is not, the upgrade task will be retried. Constraints: Currently, its value can only be abort . Value range: abort: Terminates the upgrade task. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

- Example request for retrying a cluster upgrade or AZ switchover task.
PUT https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/upgrade/bcdb711c-a7f0-4007-b8ee-9f13c05f8326/retry
- Example request for terminating a cluster upgrade or AZ switchover task.
PUT https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/upgrade/bcdb711c-a7f0-4007-b8ee-9f13c05f8326/retry?retry_mode=abort

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.

Status Code	Description
400	Invalid request. The client should not repeat the request without modifications.
403	Request rejected. The server has received the request and understood it, but the server refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.1.14 Changing the Security Group

Function

This API is used to change the security group after a cluster is created.

Constraints

Make sure the security group allows port 9200 in the inbound direction. Incorrect security group configuration may make the cluster inaccessible. Exercise caution when performing this operation.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/sg/change

Table 5-62 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose security group you want to change. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-63 Request body parameters

Parameter	Mandatory	Type	Description
security_group_ids	Yes	String	Definition: ID of the target security group. Constraints: N/A Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Change the security group of a cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/sg/change
{
  "security_group_ids" : "b1038649-1f77-4ae9-b64d-9af56e422652"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should not repeat the request without modifications.
403	Request rejected. The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.1.15 Changing the Subnet of a Cluster

Function

This API is used to bind a new subnet to newly added nodes after cluster creation, or in the case of a cluster subnet change.

NOTE

By default, subnets with the same VPC are connected. Ensure that the new subnet can reach your service system. Note that a cluster with automatic IPv6 address creation enabled can only be switched to a new subnet that uses IPv6 addresses.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/subnet/change

Table 5-64 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose subnet is to be changed. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-65 Request body parameters

Parameter	Mandatory	Type	Description
subnet_id	Yes	String	Definition: Subnet ID. Constraints: The value must be standard UUID format: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx. Value range: The value must be standard UUID format: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx. Default value: N/A

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Modify the subnet of a cluster.

```
POST /v1.0/6204a5bd270343b5885144cf9c8c158d/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/subnet/change
{
  "subnet_id" : "b1234567-1f77-4ae9-b64d-9af56e123456"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request and then retry.

Status Code	Description
403	Request rejected. The server has received the request and understood it, but refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.1.16 Automatically create an agency.

Function

If the built-in CSS agency does not exist, the system automatically creates an agency and grants the permissions required by CSS to it.

If the built-in CSS agency exists, the system removes high-risk permissions and sets the minimal permissions.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/agency/create

Table 5-66 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Request Parameters

Table 5-67 Request body parameters

Parameter	Mandatory	Type	Description
domain_id	Yes	String	Definition: Account ID. For details, see Querying Projects Based on Specified Criteria. Constraints: N/A. Value range: N/A. Default value: N/A.
domain_name	Yes	String	Definition: Account name. For details, see Authentication. Constraints: N/A. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
type	Yes	String	Definition: Type of the agency to be created. Constraints: N/A. Value range: • obs: agency permissions required for creating snapshots and log backups. • vpc: agency permissions required for version upgrade, AZ change, scale-in, and node replacement. • smn: agency permissions required for using the alerting plug-in. • elb: agency permissions required for using load balancing. Default value: N/A.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Example of the request for creating an agency used to create snapshots and log backups.

```
POST https://{Endpoint}/v1.0/{project_id}/agency/create
```

```
{
  "domain_id" : "id",
  "domain_name" : "name",
  "type" : "obs"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. Modify the request and then try again.
403	The request is rejected. The server has received the request and understood it, but the server refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.2 Specifications Query

5.2.1 Obtaining the Instance Specifications List

Function

This API is used to obtain the node specifications, including the node specifications version, cluster data search engine, number of vCPUs per node, and node memory size. You may use this API to check node specifications after a cluster is created.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/es-flavors

Table 5-68 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-69 Response body parameters

Parameter	Type	Description
versions	Array of EsflavorsVersionsResp objects	Definition: Flavor version list. Value range: N/A

Table 5-70 EsflavorsVersionsResp

Parameter	Type	Description
version	String	Definition: Elasticsearch engine version. For details, see the supported versions in Before You Start . Value range: N/A

Parameter	Type	Description
flavors	Array of EsflavorsVersionsFlavorsResp objects	Definition: Flavor list. Value range: N/A
type	String	Definition: Instance type. Value range: • ess: data node • ess-cold: cold data node • ess-master: master node. • ess-client: client node • lgs: Logstash node

Table 5-71 EsflavorsVersionsFlavorsResp

Parameter	Type	Description
cpu	Integer	Definition: Number of CPU cores of an instance. Value range: N/A
ram	Integer	Definition: Memory size of an instance. Unit: GB. Value range: N/A
name	String	Definition: Flavor name. Value range: N/A
region	String	Definition: Available region. Value range: N/A
diskrange	String	Definition: Disk capacity range of an instance. Value range: N/A

Parameter	Type	Description
availableAZ	String	Definition: AZ. Value range: N/A
flavor_id	String	Definition: Flavor ID. Value range: N/A

Example Requests

Obtain the specifications of all nodes in the project space.

```
GET https://{Endpoint}/v1.0/{project_id}/es-flavors
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "versions": [ {
    "flavors": [ {
      "cpu": 4,
      "ram": 32,
      "name": "ess.spec-4u32g",
      "region": "xx-xxx-xx",
      "diskrange": "40,2560",
      "availableAZ": "xx-xxx-xx,xx-xxx-xx",
      "flavor_id": "2d8daf1b-873f-4c2e-a7b9-2f9cbcf2f213"
    }, {
      "cpu": 8,
      "ram": 64,
      "name": "ess.spec-8u64g",
      "region": "xx-xxx-xx",
      "diskrange": "80,5120",
      "availableAZ": "xx-xxx-xx,xx-xxx-xx",
      "flavor_id": "b3d33ec6-d58a-40f0-aa51-4f671ce64b2a"
    }, {
      "cpu": 16,
      "ram": 128,
      "name": "ess.spec-16u128g",
      "region": "xx-xxx-xx",
      "diskrange": "160,10240",
      "availableAZ": "xx-xxx-xx,xx-xxx-xx",
      "flavor_id": "f74419ca-bc91-4558-b4e2-90eeefb37c6e"
    } ],
    "type": "ess",
    "version": "x.x.x"
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.2.2 Querying Available Node Specifications for a Cluster

Function

This API is used to query available alternative node specifications for a specified cluster.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/resize-flavors

Table 5-72 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: For details, see Obtaining a Project ID and Name . Default value: N/A

Table 5-73 Query Parameters

Parameter	Mandatory	Type	Description
clusterId	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
type	No	String	Definition: Cluster node types Constraints: N/A Value range: ● ess: data node. ● ess-cold: cold data node. ● ess-client: client node. ● ess-master: master node. ● lgs: Logstash node. Default value: Elasticsearch/OpenSearch cluster: ess Logstash cluster: lgs

Request Parameters

None

Response Parameters

Status code: 200

Table 5-74 Response body parameters

Parameter	Type	Description
id	String	Definition: Engine type ID. Value range: N/A
dbname	String	Definition: Engine name. Value range: N/A
versions	Array of ResizeFlavorRsp-VersionBody objects	Definition: Version list. Value range: N/A

Table 5-75 ResizeFlavorRspVersionBody

Parameter	Type	Description
id	String	Definition: Engine version ID. Value range: N/A
name	String	Definition: Engine version name. Value range: N/A
flavors	Array of ResizeFlavor objects	Definition: Target flavor of the change. Value range: N/A

Table 5-76 ResizeFlavor

Parameter	Type	Description
str_id	String	Definition: Flavor ID. Value range: N/A
cpu	Integer	Definition: Number of vCPUs available with an instance. Value range: N/A
ram	Integer	Definition: Memory size of an instance. Unit: GB. Value range: N/A
name	String	Definition: Flavor name. Value range: N/A
region	String	Definition: Regions where the node flavor is available. Value range: N/A
diskrange	String	Definition: Disk capacity of an instance, in GB. Value range: N/A
typename	String	Definition: Node type. Value range: • ess: data node • ess-cold: cold data node • ess-client: client node • ess-master: master node • lgs: Logstash node

Parameter	Type	Description
condOperationStatus	String	Definition: Flavor sales status. Value range: • normal: The flavor is in normal commercial use. • sellout: The flavor has been sold out.
localdisk	String	Definition: Whether the node uses local disks. Value range: • false: non-local disk flavor. • true: local disk flavor.
edge	Boolean	Definition: Whether this is a node flavor for edge deployments. Value range: • false: central AZ flavor. • true: edge AZ flavor.

Example Requests

Querying Available Node Specifications for a Cluster

```
GET /v1.0/{project_id}/resize-flavors?clusterId=xxx
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "id": "cf7e2c8f-096c-4fcf-b174-1ebe060679fb",
  "dbname": "elasticsearch",
  "versions": [ {
    "id": "01f53413-0a58-4b0c-848a-f625846bae23",
    "name": "7.10.2",
    "flavors": [ {
      "cpu": 2,
      "ram": 8,
      "name": "ess.spec-2u8g",
      "region": "cn-north-7",
      "typename": "ess",
      "diskrange": "40,800",
      "condOperationStatus": "normal",
      "localdisk": "false",
      "edge": false,
      "support": true,
      "str_id": "b58dfcfb-5f73-4f05-8c06-5171b12fd618="
    } ]
  } ]
}
```

```
}]  
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. Modify the request and then try again.
403	The request is rejected. The server has received the request and understood it, but the server refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.2.3 Querying Details About a Flavor

Function

Query node specification details.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/flavors/{flavor_id}

Table 5-77 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: For details, see Obtaining a Project ID and Name . Default value: N/A
flavor_id	Yes	String	Definition: Flavor ID. Constraints: N/A Value range: ID returned by the node specifications list API. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-78 Response body parameters

Parameter	Type	Description
str_id	String	Definition: Flavor ID. Value range: N/A

Parameter	Type	Description
name	String	Definition: Flavor name. Value range: N/A
condOperationStatus	String	Definition: This parameter takes effect region-wide. If an AZ is not configured in the condOperationAz parameter, the value of this parameter is used by default. Value range: • normal: The flavor is in normal commercial use. • sellout: The flavor has been sold out.
condOperationAz	String	Definition: This parameter takes effect AZ-wide. If an AZ is not configured in this parameter, the value of the condOperationAz parameter is used by default. The value format of this parameter is az(xx) . The value in parentheses is the flavor status of the AZ. If the parentheses are left blank, the configuration is invalid. The value range for this parameter is the same as that for condOperationStatus . Value range: N/A
flavor_detail	Array of ShowFlavorDetailRspFlavorDetail objects	Definition: Flavor attributes. Value range: N/A

Table 5-79 ShowFlavorDetailRspFlavorDetail

Parameter	Type	Description
name	String	Definition: Attribute name. Value range: - cpu: flavor CPU attribute. - mem: flavor memory attribute. - diskrange: flavor disk capacity range attribute. - flavorTypeCn: flavor classification attribute in Chinese. - flavorTypeEn: flavor classification attribute in English.
value	String	Definition: Attribute value. Value range: N/A

Example Requests

Sample request for querying node flavor details.

```
GET https://{Endpoint}/v1.0/{project_id}/flavors/{flavor_id}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "name" : "ess.spec-2u8g",
  "condOperationStatus" : "normal",
  "condOperationAz" : "cn-north-7a(sellout),cn-north-7b(normal),cn-north-7c(normal)",
  "str_id" : "b58dfcfb-5f73-4f05-8c06-5171b12fd618",
  "flavor_detail" : [ {
    "name" : "cpu",
    "value" : "2"
  }, {
    "name" : "diskrange",
    "value" : "40,800"
  }, {
    "name" : "flavorTypeCn",
    "value" : "General computing"
  }, {
    "name" : "flavorTypeEn",
    "value" : "General computing"
  }, {
    "name" : "mem",
    "value" : "8"
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. Modify the request and then try again.
403	The request is rejected. The server has received the request and understood it, but the server refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.2.4 Query available node specifications by engine.

Function

Query available node specifications by engine.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/datastore/{datastore_id}/flavors

Table 5-80 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
datastore_id	Yes	String	Definition: Engine type ID. Constraints: N/A Value range: Elasticsearch: cf7e2c8f-096c-4fcf-b174-1ebe060679fb. Opensearch: 07ec9f86-ec2f-49e7-8913-373003aedef32. Logstash: 575276bb-87e5-4e18-8e1e-e748d8ad3a06. Default value: N/A

Table 5-81 Query Parameters

Parameter	Mandatory	Type	Description
datastore_version_id	No	String	Definition: Engine type ID. Constraints: N/A Value range: Elasticsearch 7.10.2: 01f53413-0a58-4b0c-848a-f625846bae23. Opensearch 2.19.0: 11a9df5c-711f-496c-866d-a4521c179671. Logstash 7.10.0: f5609cf0-3514-49ef-87db-a3df2858a46f. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-82 Response body parameters

Parameter	Type	Description
id	String	Definition: Engine type ID. Value range: N/A
dbname	String	Definition: Engine name. Value range: N/A
versions	Array of FlavorRespVersionBody objects	Definition: Engine version. Value range: N/A
modelList	ModelList object	Definition: Model list. Value range: N/A

Table 5-83 FlavorRespVersionBody

Parameter	Type	Description
id	String	Definition: Version ID. Value range: N/A
name	String	Definition: Version name. Value range: N/A
flavors	Array of Flavor objects	Definition: Specification information. Value range: N/A

Table 5-84 Flavor

Parameter	Type	Description
cpu	Integer	Definition: Number of vCPUs available with an instance. Value range: N/A
ram	Integer	Definition: Memory size of an instance. Unit: GB. Value range: N/A
name	String	Definition: Flavor name. Value range: N/A
region	String	Definition: Regions where the node flavor is available. Value range: N/A
typename	String	Definition: Name of a node type Value range: N/A
diskrange	String	Definition: Instance disk capacity range, in GB. Value range: N/A
condOperationStatus	String	Definition: This parameter takes effect region-wide. If an AZ is not configured in the condOperationAz parameter, the value of this parameter is used by default. Value range: • normal: The flavor is in normal commercial use. • sellout: The flavor has been sold out.

Parameter	Type	Description
condOperationAz	String	Definition: This parameter takes effect AZ-wide. If an AZ is not configured in this parameter, the value of the condOperationAz parameter is used by default. The value format of this parameter is az(xx) . The value in parentheses is the flavor status of the AZ. If the parentheses are left blank, the configuration is invalid. The value range for this parameter is the same as that for condOperationStatus . Value range: N/A
localdisk	String	Definition: Whether the node uses local disks. Value range: • true: local disks • false: non-local disks
flavorTypeCn	String	Definition: Flavor categories in Chinese. Value range: N/A
flavorTypeEn	String	Definition: Flavor categories in English. Value range: N/A
edge	Boolean	Definition: Whether this is a node flavor for edge deployments. Value range: • true: This is a node flavor for edge deployments. • false: This is a general flavor.
str_id	String	Definition: Flavor ID. Value range: N/A

Parameter	Type	Description
isAllowHttps	Boolean	Definition: Whether the node type supports HTTPS access. Value range: N/A

Table 5-85 ModelList

Parameter	Type	Description
totalSize	Integer	Definition: Model quantity Value range: N/A
models	Array of Model objects	Definition: Model list Value range: N/A

Table 5-86 Model

Parameter	Type	Description
id	String	Definition: Model ID Value range: N/A
name	String	Definition: Model name Value range: N/A
datastore_type	String	Definition: Model type Value range: N/A

Parameter	Type	Description
datastore_version	String	Definition: Model version Value range: N/A
is_text_model	String	Definition: Whether it is a text model Value range: N/A
model_version_id	String	Definition: Model version ID Value range: N/A
desc	String	Definition: Model description Value range: N/A
language	String	Definition: Model language Value range: N/A

Example Requests

None

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. Modify the request and then try again.

Status Code	Description
403	The request is rejected. The server has received the request and understood it, but the server refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.3 Tag Management

5.3.1 Querying All Tags

Function

This API is used to query all tags of a user in the current project, including the keys and values. This facilitates tag management by group.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/{resource_type}/tags

Table 5-87 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID. Default value: N/A

Parameter	Mandatory	Type	Description
resource_type	Yes	String	Definition: Resource type. Currently, its value can only be css-cluster . Constraints: N/A Value range: css-cluster: Indicates the cluster type. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-88 Response body parameters

Parameter	Type	Description
tags	Array of ShowAllTagsTagsResp objects	Definition: List of cluster tags. Value range: N/A

Table 5-89 ShowAllTagsTagsResp

Parameter	Type	Description
key	String	Definition: Tag key. Value range: N/A
values	Array of strings	Definition: Tag values. Value range: N/A

Example Requests

Query all cluster tags in a project.

```
GET https://{Endpoint}/v1.0/{project_id}/css-cluster/tags
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "tags": [ {
    "key": "key1",
    "values": [ "value1", "value2" ]
  }, {
    "key": "key2",
    "values": [ "value1", "value2" ]
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should modify the request instead of re-initiating it.
404	The requested resource could not be found. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.3.2 Querying Tags of a Specified Cluster

Function

Tags are cluster identifiers. You can add tags to clusters to identify and manage cluster resources. You can add tags to a cluster when creating a cluster or add them on the details page of a created cluster. This API is used to query the tags of a specified cluster. It facilitates tag management.

Calling Method

For details, see [Calling APIs](#).

URI

```
GET /v1.0/{project_id}/{resource_type}/{cluster_id}/tags
```

Table 5-90 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
resource_type	Yes	String	Definition: Resource type. Currently, its value can only be css-cluster . Constraints: N/A Value range: css-cluster: Indicates the cluster type. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-91 Response body parameters

Parameter	Type	Description
tags	Array of ShowTagsTagsResp objects	Definition: List of cluster tags. Value range: N/A

Table 5-92 ShowTagsTagsResp

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

Example Requests

Query the tags of a cluster.

```
GET https://{Endpoint}/v1.0/{project_id}/css-cluster/4f3deec3-efa8-4598-bf91-560aad1377a3/tags
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "tags": [ {
    "key": "key1",
    "value": "value1"
  }, {
    "key": "key2",
    "value": "value3"
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should modify the request instead of re-initiating it.
404	The requested resource could not be found. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.3.3 Adding Tags to a Cluster

Function

Tags are cluster identifiers. You can add tags to clusters to identify and manage cluster resources.

This API is used to add tags to a cluster.

Constraints

Each cluster can have a maximum of 20 tags.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/{resource_type}/{cluster_id}/tags

Table 5-93 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster that you want to add tags to. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
resource_type	Yes	String	Definition: Resource type. Currently, its value can only be css-cluster . Constraints: N/A Value range: css-cluster: Indicates the cluster type. Default value: N/A

Request Parameters

Table 5-94 Request body parameters

Parameter	Mandatory	Type	Description
tag	Yes	Tag object	Definition: Tag object you want to add. Constraints: N/A Value range: N/A Default value: N/A

Table 5-95 Tag

Parameter	Mandatory	Type	Description
key	Yes	String	Definition: Tag key. Constraints: • The key cannot be empty. Its length is 1 to 128 characters. • The value can contain UTF-8 letters, digits, spaces, and special characters _ . : = + - @. • Tags starting with <i>sys</i> are system tags and cannot be used by tenants. • Do not use spaces to start or end the key. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
value	Yes	String	Definition: Tag value. Constraints: - The value contains 0 to 255 characters. - The value can contain UTF-8 letters, digits, spaces, and special characters _./\=-+-@ Value range: N/A Default value: N/A

Response Parameters

Status code: 204

Request succeeded.

None

Example Requests

Add tags to a cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/css-cluster/4f3deec3-efa8-4598-bf91-560aad1377a3/tags
{
  "tag": {
    "key": "K1",
    "value": "V1"
  }
}
```

Example Responses

None

Status Codes

Status Code	Description
204	Request succeeded.
400	Invalid request. The client should modify the request instead of re-initiating it.

Status Code	Description
409	The request could not be completed due to a conflict with the current state of the resource. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.3.4 Deleting a Cluster Tag

Function

Tags are cluster identifiers. You can add tags to clusters to identify and manage cluster resources.

This API is used to delete cluster tags that are no longer needed.

Calling Method

For details, see [Calling APIs](#).

URI

DELETE /v1.0/{project_id}/{resource_type}/{cluster_id}/tags/{key}

Table 5-96 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster that you want to delete tags from. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
resource_type	Yes	String	Definition: Resource type. Currently, its value can only be css-cluster . Constraints: N/A Value range: css-cluster: Indicates the cluster type. Default value: N/A
key	Yes	String	Definition: Name of the tag you want to delete Constraints: If the tag name is in Chinese, encode the tag name using URL before calling the API. Value range: N/A Default value: N/A

Request Parameters

None

Response Parameters

Status code: 204

Request succeeded.

None

Example Requests

Delete the tags of a cluster.

```
DELETE https://{Endpoint}/v1.0/{project_id}/css-cluster/2a197c4d-5467-4003-931d-83ec49939cf/tags/K1
```

Example Responses

None

Status Codes

Status Code	Description
204	Request succeeded.
400	Invalid request. The client should modify the request instead of re-initiating it.
404	The requested resource could not be found. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.3.5 Adding or Deleting Cluster Tags in Batches

Function

This API is used to add or delete tags in batches for Elasticsearch, Logstash, and OpenSearch clusters.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/{resource_type}/{cluster_id}/tags/action

Table 5-97 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster for which you want to add or delete tags in batches. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
resource_type	Yes	String	Definition: Cluster resource type. Currently, its value can only be css-cluster . Constraints: N/A Value range: css-cluster: Indicates the cluster type. Default value: N/A

Request Parameters

Table 5-98 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Definition: Cluster tag operation type. This attribute is used to identify the required operation type. Constraints: N/A Value range: ● create: Tags are added in batches. ● delete: Tags are deleted in batches. Default value: N/A
tags	Yes	Array of Tag objects	Definition: List of cluster tags. Constraints: N/A Value range: N/A Default value: N/A

Table 5-99 Tag

Parameter	Mandatory	Type	Description
key	Yes	String	Definition: Tag key. Constraints: • The key cannot be empty. Its length is 1 to 128 characters. • The value can contain UTF-8 letters, digits, spaces, and special characters _ . : = + - @. • Tags starting with <i>sys</i> are system tags and cannot be used by tenants. • Do not use spaces to start or end the key. Value range: N/A Default value: N/A
value	Yes	String	Definition: Tag value. Constraints: • The value contains 0 to 255 characters. • The value can contain UTF-8 letters, digits, spaces, and special characters _.:/=+-@ Value range: N/A Default value: N/A

Response Parameters

Status code: 204

Request succeeded.

None

Example Requests

Create cluster tags in batches.

```
POST https://{Endpoint}/v1.0/{project_id}/css-cluster/4f3deec3-efa8-4598-bf91-560aad1377a3/tags/action
{
  "action": "create",
  "tags": [ {
    "key": "K1",
    "value": "V1"
  }, {
    "key": "K2",
    "value": "V2"
  } ]
}
```

Example Responses

None

Status Codes

Status Code	Description
204	Request succeeded.
400	Invalid request. Do not retry the request before modification.
409	The request could not be completed due to a conflict with the current state of the resource. This status code indicates that the resource that the client attempts to create already exists, or the request fails to be processed because of the update of the conflict request.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.4 Modify Configuration

5.4.1 Adding Instances and Expanding Instance Storage Capacity

Function

If a cluster struggles to maintain performance in the face of rapid data growth or sustained high memory usage, you can add more nodes to it or increase the storage capacity of its existing nodes.

Constraints

- The node storage capacity of a cluster can only be increased—not decreased. Select an appropriate node storage capacity based on the data volume and projected data growth.
- For a pay-per-use cluster, you cannot increase the number of nodes and node storage capacity at the same time. For a yearly/monthly cluster, however, you cannot do that.
- The storage capacity of the master and client nodes in a cluster cannot be expanded.
- The storage capacity of data nodes that use local disks cannot be expanded.
- For a Logstash cluster, you can only add more nodes, but cannot increase the storage capacity of its nodes.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/role_extend

Table 5-100 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster you want to scale out. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-101 Request body parameters

Parameter	Mandatory	Type	Description
grow	Yes	Array of RoleExtendG rowReq objects	Definition: A detailed description of the cluster scale-out request. Constraints: N/A

Table 5-102 RoleExtendGrowReq

Parameter	Mandatory	Type	Description
type	Yes	String	Definition: Type of the node to be added. Constraints: You can only add instances, rather than increase storage capacity, on nodes of the ess-master, ess-client, and lgs types. Value range: • ess: data node • ess-master: master node • ess-client: client node • ess-cold: cold data node • lgs: Logstash nodes in Logstash clusters You can select multiple node types, but cannot select the same node for different roles. NOTE lgs cannot be combined with other options. Default value: N/A

Parameter	Mandatory	Type	Description
nodesize	Yes	Integer	Definition: Number of instances to be added. Constraints: After scaling, the number of nodes of each type must meet the following requirements: • Data nodes: If the cluster does not have master nodes, the value ranges from 1 to 32. Otherwise, the value ranges from 1 to 200. • Master nodes: 3, 5, 7, or 9 (must be an odd number from 3 to 9) • Client nodes: 1 to 64. • Cold data nodes: 1 to 32. • Logstash nodes: 1 to 100. Value range: See Constraints. Default value: N/A

Parameter	Mandatory	Type	Description
disksize	Yes	Integer	Definition: Storage capacity to be added, in GB. Constraints: The sum of the original instance storage capacity plus the expanded instance storage capacity cannot exceed the default upper limit of storage capacity set during cluster creation. If there is no need to add more instances, set this parameter to 0. With a yearly/monthly cluster, you cannot change the number of nodes and disk capacity at the same time. Value range: • The scale-out step of ESS and ESS-cold nodes is 20. • Storage capacity of ess-master and ess-client nodes cannot be expanded. • Logstash does not support storage capacity expansion. • Multiple constraints are put in different rows in random order. Default value: N/A

Response Parameters

Status code: 200

Table 5-103 Response body parameters

Parameter	Type	Description
id	String	Definition: Cluster ID. Value range: N/A

Example Requests

Change the number of instances and storage capacity of the current cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/role_extend
{
  "grow" : [ {
    "type" : "ess-master",
    "nodesize" : 2,
    "disksize" : 0
  }, {
    "type" : "ess",
    "nodesize" : 0,
    "disksize" : 40
  }, {
    "type" : "ess-client",
    "nodesize" : 1,
    "disksize" : 0
  }, {
    "type" : "ess-cold",
    "nodesize" : 1,
    "disksize" : 0
  } ],
  "is_auto_pay" : 1
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "id" : "4f3deec3-efa8-4598-bf91-560aad1377a3"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.4.2 Adding Dedicated Master or Client Nodes or Cold Data Nodes

Function

This API is used to add dedicated master or client nodes or cold data nodes to an existing cluster that previously does not have such nodes. (When planning a cluster, you cannot always accurately predict future changes in data volumes. Add dedicated master or client nodes or cold data nodes is an effective way to scale up a cluster.)

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/type/{type}/independent

Table 5-104 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster that requires dedicated master or client nodes or cold data nodes. For details about how to obtain the cluster ID, see Obtaining the Cluster ID. Constraints: N/A Value range: Cluster ID. Default value: N/A
type	Yes	String	Definition: Types of dedicated nodes to add. Constraints: N/A Value range: • ess-master: master node • ess-client: client node • ess-cold: cold data node Default value: N/A

Request Parameters

Table 5-105 Request body parameters

Parameter	Mandatory	Type	Description
type	Yes	IndependentTypeReq object	Definition: Request body parameters for dedicated master, client, or cold data nodes. Constraints: N/A Value range: N/A Default value: N/A

Table 5-106 IndependentTypeReq

Parameter	Mandatory	Type	Description
flavor_ref	Yes	String	Definition: Flavor ID. Constraints: N/A Value range: You can obtain the value of this parameter by calling the API for Obtaining the Instance Specifications List . Select the flavor ID suitable for your cluster version. Default value: N/A

Parameter	Mandatory	Type	Description
node_size	Yes	Integer	Definition: Number of dedicated nodes. Constraints: N/A Value range: • If the node type is ess-master, the number of nodes must be an odd number in the range 3 to 9. • If the node type is ess-client, the number of nodes must be in the range 1 to 64. • If the node type is ess-cold, indicating to add dedicated cold data nodes, the number of nodes must be in the range 1 to 32. Default value: N/A
volume_type	No	String	Definition: Node storage type. Constraints: This parameter cannot be set when flavor_ref is set to a local disk flavor. Value range: • COMMON: common I/O • HIGH: high I/O • ULTRAHIGH: ultra-high I/O • ESSD: ultra-fast SSD Default value: If flavor_ref is not set to a local disk flavor, the default value is ULTRAHIGH.

Parameter	Mandatory	Type	Description
volume_size	No	Integer	Definition: Node storage capacity. Constraints: • This parameter cannot be set when flavor_ref is set to a local disk flavor. • The value must be greater than 0 and a common multiple of 4 and 10, in GB. • Adding dedicated cold data nodes: 100 GB or the minimum disk capacity supported by the selected node flavor, whichever is larger. • Adding dedicated master or client nodes: The default size is 40 GB and cannot be changed. Value range: You can obtain the disk size from the diskrange attribute in Obtaining the Instance Specifications List. Default value: If flavor_ref is not set to a local disk flavor: NOTE Adding dedicated cold data nodes: The disk size should be greater than 100 GB.

Response Parameters

Status code: 200

Table 5-107 Response body parameters

Parameter	Type	Description
id	String	Definition: Cluster ID. Value range: N/A

Example Requests

Add dedicated master or client nodes or cold data nodes.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/type/ess-client/independent

{
  "type": {
    "flavor_ref": "d9dc06ae-b9c4-4ef4-acd8-953ef4205e27",
    "node_size": 3,
    "volume_type": "COMMON",
    "volume_size": 40
  }
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "id": "320afa24-ff2a-4f44-8460-6ba95e512ad4"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
403	Request rejected. The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.
500	The server has received the request but could not understand it.

Error Codes

See [Error Codes](#).

5.4.3 Changing Node Specifications (Discarded)

Function

This API is used to modify the node specifications of a cluster. It can only change the specifications of ess nodes (data nodes). If you need to change other types of nodes as well, use the API for [Changing the Specifications of Specified Nodes](#).

NOTE

All mission-critical data has been backed up before a disk change. This is to prevent data loss.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/flavor

Table 5-108 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose specifications you want to change. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-109 Request body parameters

Parameter	Mandatory	Type	Description
newFlavorId	Yes	String	Definition: ID of the new node flavor or disk type. It is determined by operationType. Constraints: N/A Value range: • When operationType is set to vm, newFlavorId indicates the node flavor ID after the change. This parameter is obtained by calling the Obtaining the Instance Specifications List API. The API selects a flavor_id based on the values of the name attribute. • When operationType is set to volume, newFlavorId indicates the new disk type after the change. The following disk types are supported: • COMMON: common I/O • HIGH: high I/O • ULTRAHIGH: ultra-high I/O • ESSD: ultra-fast SSD NOTE The old and new node specifications must be under the same Elasticsearch engine version. Default value: N/A

Parameter	Mandatory	Type	Description
operationType	No	String	Definition: Operation type for specifications change. Constraints: N/A Value range: • vm: Change the node flavor. • volume: Change the disk type. Default value: vm
needCheckReplica	No	Boolean	Definition: Whether to check replicas when changing cluster specifications. Constraints: N/A Value range: • <i>*true</i>: Enable replica verification. • false: Disable replica verification. Default value: true
needCheckClusterStatus	No	Boolean	Definition: Whether to check the cluster status when changing cluster specifications. Constraints: N/A Value range: • true: Enable cluster load check. • false: Cluster status verification is ignored. Default value: true

Parameter	Mandatory	Type	Description
clusterLoadCheck	No	Boolean	Definition: Whether to check the cluster load when changing cluster specifications. Constraints: N/A Value range: • true: enable cluster load check. • false: Skip cluster load check. Default value: true

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

- ID of the new flavor.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/flavor
```

```
{
  "needCheckReplica" : false,
  "newFlavorId" : "35b060a4-f152-48ce-8773-36559ceb81f2",
  "isAutoPay" : 1,
  "needCheckClusterStatus" : true
}
```

- Change the disk type.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/flavor
```

```
{
  "needCheckReplica" : false,
  "newFlavorId" : "ULTRAHIGH",
  "operationType" : "volume",
  "isAutoPay" : 1,
  "clusterLoadCheck" : true
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.4.4 Scaling In a Cluster by Removing a Specific Node

Function

If a cluster has excess capacity due to off-peak traffic or reduced data volume, you can reduce its nodes to cut costs.

Constraints

- During a scale-in, the data on the to-be-removed nodes needs to be migrated to the remaining nodes. The timeout for data migration per node is 48 hours. Scale-in will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- Yearly/Monthly clusters do not support the removal of specified nodes by calling this API.
- All mission-critical data has been backed up before a cluster scale-in. This is to prevent data loss.
- Only Elasticsearch and OpenSearch clusters support this option.
- During a scale-in, the data on the to-be-removed nodes needs to be migrated to the remaining nodes. The timeout for data migration per node is 48 hours. Scale-in will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- For a cluster without master nodes, scale-in is allowed only if the number of data nodes plus cold data nodes is at least three. During a scale-in, you can only remove less than half of the data nodes plus cold data nodes, and the remaining number must be greater than the maximum number of index replicas.

- For a cluster with master nodes, scale-in is allowed only if the cluster has at least two data nodes. During a scale-in, you can only remove less than half of the cluster's master nodes.
- After the scale-in, the cluster nodes' disk usage must be less than 80%.
- After the scale-in, there has to be at least one of each type of node in each AZ. For a cross-AZ cluster, the difference between the numbers of the same-type nodes in different AZs cannot exceed 1.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/node/offline

Table 5-110 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be scaled in. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-111 Request body parameters

Parameter	Mandatory	Type	Description
migrate_data	No	String	Definition: Whether to migrate data. Value range: • true: Migrate data. During data migration, the system migrates all data from the to-be-removed nodes to the remaining nodes, and removes these nodes upon completion of the data migration. • false: Do not migrate data. If the data on the to-be-removed nodes has replicas on other nodes, data migration can be skipped and the cluster change can be completed faster. Default value: true
shrink_nodes	Yes	Array of strings	Definition: ID of the node to be removed. Constraints: N/A Value range: Obtain the ID attribute in instances by referring to Querying Cluster Details . Default value: N/A

Parameter	Mandatory	Type	Description
agency_name	No	String	Definition: Agency name that grants the current account the permission to access and use OBS. To store snapshots to an OBS bucket, you must have the required OBS access permissions. Constraints: VPC permissions required by the agency: "vpc:subnets:get","vpc:ports:*". This parameter is mandatory when the new IAM plane is connected. This parameter is optional when the old IAM plane is connected. Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Scale in a cluster by removing specified nodes.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/node/offline
{
  "shrink_nodes": [ "2077bdf3-b90d-412e-b460-635b9b159c11" ],
  "migrate_data": "true"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.4.5 Change the specifications of a specified node type.

Function

If the workloads on a cluster's data plane change, you can scale the cluster vertically by changing its node specifications or node storage type.

- Cluster types that support node specifications change: Elasticsearch, OpenSearch, and Logstash.
- Cluster types that support node storage type (disk type) change: Elasticsearch and OpenSearch. All mission-critical data must be backed up before a node storage type (disk type) change. This helps to prevent data loss.

You can specify the following node types when calling this API:

- `ess`: data node (only for Elasticsearch and OpenSearch).
- `ess-cold`: cold data node (only for Elasticsearch and OpenSearch).
- `ess-client`: client node (only for Elasticsearch and OpenSearch).
- `ess-master`: master node (only for Elasticsearch and OpenSearch).
- `lgs`: Logstash node (only for Logstash).

Constraints

- Make sure the cluster status is Available and there are no ongoing tasks.
- The node specifications and storage type cannot be changed for nodes that use local disks.
- The node specifications and storage type cannot be changed at the same time.

- The node storage type can be changed only for data nodes and cold data nodes.
- When you change the node storage type, data needs to be migrated between different nodes. The timeout for data migration per node is 48 hours. Upgrade will fail if this timeout expires. When the cluster has large quantities of data, you are advised to manually adjust the data migration rate and avoid performing the migration during peak hours.
- For a cluster without master nodes, the node storage type can be changed only if the number of data nodes plus cold data nodes is at least three.
- For a cluster with master nodes, this operation is allowed only if the cluster has at least two data nodes.
- During a node storage type change, there is always one node that is unavailable. To ensure service continuity, there has to be at least two nodes of each type for the cluster in each of its AZs, and the total number of data nodes and cold data nodes is greater than the maximum number of index replicas plus 1.
- During a node specifications changes, nodes are brought offline in order to make the changes. To ensure service continuity, make sure all shards have replicas.
- Make sure the disk usage is always less than 80/ %during the change.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/{types}/flavor

Table 5-112 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster whose specifications you want to change. For details about how to obtain the cluster ID, see Obtaining the Cluster ID. Constraints: N/A Value range: Cluster ID. Default value: N/A
types	Yes	String	Definition: Type of the cluster node you want to modify. Constraints: N/A Value range: • ess: data node • ess-cold: cold data node • ess-client: client node • ess-master: master node • lgs: Logstash node. Default value: N/A

Request Parameters

Table 5-113 Request body parameters

Parameter	Mandatory	Type	Description
new_flavor_id	Yes	String	Definition: ID of the new node flavor or node storage type (disk type). It is determined by operation_type. Constraints: The value range depends on operation_type. Value range: - When operation_type is set to vm, newFlavorId indicates the node flavor ID after the change. This parameter is obtained by calling the Obtaining the Instance Specifications List API. The API selects a flavor_id based on the values of the name attribute. - When operation_type is set to volume, new_flavor_id indicates the node storage type (disk type). The options are as follows: ● COMMON: common I/O ● HIGH: high I/O ● ULTRAHIGH: ultra-high I/O ● ESSD: ultra-fast SSD Default value: N/A

Parameter	Mandatory	Type	Description
operation_type	No	String	Definition: Operation type. Constraints: N/A Value range: • vm: Change the node flavor. • volume: Change the disk type. Default value: vm
need_check_replica	No	Boolean	Definition: Whether to verify replicas. Constraints: N/A Value range: • <i>*true</i>: Enable replica verification. • false: Disable replica verification. Default value: true
need_check_cluster_status	No	Boolean	Definition: Whether to check the cluster status. The value can be true or false. The default value is true . Constraints: N/A Value range: • true: Enable cluster load check. • false: Cluster status verification is ignored. Default value: true

Parameter	Mandatory	Type	Description
cluster_load_check	No	Boolean	Definition: Whether to check the cluster load. The value can be true or false . The default value is true . Constraints: N/A Value range: • true: enable cluster load check. • false: Skip cluster load check. Default value: true

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Change the disk type of data nodes for a yearly/monthly Elasticsearch cluster. Before the change, disable index replica verification, and enable cluster status and cluster load check. The change in fees will apply automatically.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ess/flavor
```

```
{
  "need_check_replica" : false,
  "new_flavor_id" : "HIGH",
  "is_auto_pay" : 1,
  "need_check_cluster_status" : true,
  "operation_type" : "volume",
  "cluster_load_check" : true
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.4.6 Scaling In Nodes of a Specific Type

Function

This API is used to remove instances of specific types and reduce instance storage capacity in a cluster. Yearly/Monthly clusters do not support the removal of specified node types by calling this API.

NOTE

All mission-critical data has been backed up before a cluster scale-in. This is to prevent data loss.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/extend/{project_id}/clusters/{cluster_id}/role/shrink

Table 5-114 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be scaled in. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-115 Request body parameters

Parameter	Mandatory	Type	Description
shrink	Yes	Array of ShrinkNodeReq objects	Parameter description: Type and quantity of nodes you want to scale in.

Parameter	Mandatory	Type	Description
agency_name	No	String	Definition: Agency name. You can create an agency to allow CSS to call other cloud services. Constraints: VPC permissions required by the agency: "vpc:subnets:get","vpc:ports:*". This parameter is mandatory when the new IAM plane is connected. This parameter is optional when the old IAM plane is connected. Value range: N/A Default value: N/A
operation_type	No	String	Definition: Operation type. Constraints: N/A Value range: vm: Reduce nodes. Default value: vm
cluster_load_check	No	Boolean	Definition: Whether to check the cluster load. Constraints: N/A Value range: true: enable cluster load check. true: skip cluster load check. Default value: true

Table 5-116 ShrinkNodeReq

Parameter	Mandatory	Type	Description
reducedNodeNum	Yes	Integer	Definition: Number of nodes you want to reduce. Constraints: • After the scale-in, there has to be at least one node of each type in each AZ. • In a cross-AZ cluster, the difference between the numbers of same-type nodes in different AZs cannot exceed 1. • For a cluster with no Master nodes, the number of removed data nodes (including cold data nodes and other types of nodes) in a scale-in must be fewer than half of the number of the original data nodes, and the number of remaining data nodes after a scale-in must be greater than the maximum number of index replicas. For a cluster with master nodes, the number of removed master nodes in a scale-in must be fewer than half of the number of the original master nodes. After scale-in, there has to be an odd number of master nodes, and there has to be at least three of them. Value range: See Constraints. Default value: N/A

Parameter	Mandatory	Type	Description
type	Yes	String	Definition: Node type. Constraints: N/A [If operation_type is set to volume , this parameter cannot be set to ess-client or ess-master .] (tag:non_helpcenter) Value range: • ess: data node • ess-cold: cold data node • ess-client: client node • ess-master: master node • lgs: Logstash node Default value: N/A

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Scale in a cluster by specifying the node type and the number of nodes to be removed.

```
POST https://{Endpoint}/v1.0/extend/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/role/shrink
{
  "shrink": [ {
    "type": "ess",
    "reducedNodeNum": 1
  } ]
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.4.7 Replacing a Node

Function

This API is used to replace a failed node.

If a node in an Elasticsearch cluster is faulty, you can replace it to restore services.

The node replacement process is as follows:

1. Migrate data from the node that needs to be replaced to other available nodes.
2. Rebuild the node using the original ID, IP address, specifications, and AZ.
3. Add the new node into the cluster. The system automatically triggers a shard reallocation, migrating some of the shards

to the new node

This process does not interrupt services because data is migrated from the replaced node to other available nodes.

NOTE

All mission-critical data has been backed up before a node replacement. This is to prevent data loss.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/instance/{instance_id}/replace

Table 5-117 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster where nodes are to be replaced. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
instance_id	Yes	String	Definition: ID of the node to be replaced. Constraints: N/A Value range: Node ID. Obtain the ID attribute in instances by referring to Querying Cluster Details . Default value: N/A

Table 5-118 Query Parameters

Parameter	Mandatory	Type	Description
migrateData	No	String	Definition: Whether to migrate data. Constraints: N/A Value range: • true: Migrate data. • false: Do not migrate data. Default value: true
agency	No	String	Definition: Agency name. To store snapshots to an OBS bucket, you must have the required OBS access permissions. Select an IAM agency to grant the current account the permission to access and use OBS. Constraints: • VPC permissions required by the agency: "vpc:subnets:get", "vpc:ports.*". • This parameter is mandatory when the new IAM plane is connected. This parameter is optional when the old IAM plane is connected. • The value can contain only letters (a-z and A-Z), digits (0-9), hyphens (-), and underscores (_). It cannot contain Chinese characters or other special characters. Value range: N/A Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Replace a node.

```
PUT https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/instance/43e63449-339c-4280-a6e9-da36b0685995/replace?migrateData=true
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should modify the request instead of re-initiating it.
404	The requested resource could not be found. The client should modify the request instead of re-initiating it.

Error Codes

See [Error Codes](#).

5.4.8 Changing the Security Mode

Function

This API is used to configure the security mode of a cluster. The security mode settings for clusters include:

- Security mode + HTTP: strikes a balance between performance and security. Make sure the cluster is deployed in a secure environment. Do not expose the cluster's network interface to the public network.
- Security mode + HTTPS: prioritizes security and in-transit data encryption, and allows public network access.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/mode/change

Table 5-119 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the target cluster. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-120 Request body parameters

Parameter	Mandatory	Type	Description
authority_enable	Yes	Boolean	Definition: Whether to enable the security mode for the cluster. Constraints: Only Elasticsearch 6.5.4 security-mode clusters and later support this parameter. Value range: • true: enable. • false: disable. Default value: true
admin_pwd	No	String	Definition: Cluster password in security mode. Constraints: Only Elasticsearch 6.5.4 security-mode clusters and later support this parameter. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
https_enable	Yes	Boolean	Definition: Whether to enable HTTPS for the cluster. Constraints: - Elasticsearch: Only Elasticsearch 6.5.4 and later versions support security-mode clusters. - HTTPS access can be disabled only for OpenSearch 1.3.6 and 2.19.0 security-mode clusters. For other versions, HTTPS access is forcibly enabled and cannot be disabled. Value range: true: Enable HTTPS. false: Disable HTTPS. Default value: true

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Change the security mode of the current cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/mode/change
{
  "authority_enable" : true,
  "admin_pwd" : "admin@1234",
  "https_enable" : true
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.4.9 Scaling Out a Cluster (Discarded)

Function

This API is used to add data nodes to an Elasticsearch or OpenSearch cluster. You are advised to use [Adding Instances and Expanding Instance Storage Capacity](#).

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/extend

Table 5-121 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster you want to scale out. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-122 Request body parameters

Parameter	Mandatory	Type	Description
grow	Yes	ExtendClusterGrowReq object	Definition: Detailed description about the cluster scale-out request. Constraints: N/A Value range: N/A Default value: N/A

Table 5-123 ExtendClusterGrowReq

Parameter	Mandatory	Type	Description
modifySize	Yes	Integer	Definition: Number of instances to be scaled out. The total number of existing instances and newly added instances in a cluster cannot exceed 32. Constraints: N/A Value range: 1-32 Default value: N/A

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Change the number of instances in the current cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/extend
{
  "grow": {
    "modifySize": 4
  },
  "isAutoPay": 1
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.

Status Code	Description
404	The requested resource could not be found. Modify the request before retry.

Error Codes

See [Error Codes](#).

5.4.10 Changing AZs for a Cluster Instance

Function

An availability zone (AZ) is a physical region where resources use independent power supplies and networks. AZs in the same region can communicate with each other through internal networks but are physically isolated.

Multi-AZ deployment is a high-availability solution provided by CSS. You can deploy a cluster across two or three AZs in the same region. This helps prevent data loss and lower the risk of downtime.

This API is used to change AZs by specifying node types.

Constraints

- To ensure service continuity, the total number of data nodes plus cold data nodes in the cluster must be at least three.
- During the change, nodes are brought offline one by one and then new nodes are created. Ensure that the total disk capacity of the remaining nodes is sufficient for handling all the data of the cluster, and that the nodes' disk usage stays below 80%.
- To make sure that all index shards in the cluster can be allocated to the remaining data nodes and cold data nodes, the total number of data nodes and cold data nodes must be greater than the maximum number of index replicas plus 1.
- Before a change completes, some nodes may have already been moved to a new AZ. In this case, the AZs before and after the change are both displayed. After the change succeeds, the new AZs and their nodes will be displayed properly.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/inst-type/{inst_type}/azmigrate

Table 5-124 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose AZs are to be changed. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
inst_type	Yes	String	Definition: Types of nodes whose AZs are to be changed. Constraints: All mission-critical data has been backed up before an AZ switchover. This is to prevent data loss. Value range: • all: all node types • ess: data node • ess-cold: cold data node • ess-client: client node • ess-master: master node Default value: N/A

Request Parameters

Table 5-125 Request body parameters

Parameter	Mandatory	Type	Description
source_az	Yes	String	Definition: AZ where the node is located. Constraints: N/A Value range: N/A Default value: N/A
target_az	Yes	String	Definition: Destination AZ of the node. Constraints: N/A Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
migrate_type	Yes	String	Definition: AZ migration method. Constraints: • To complete this AZ change, the system will try to move as few nodes as possible to rebuild the cluster. During this process, the YML configuration file of the nodes that are not modified is also updated. You need to restart the cluster to make the change take effect. • If you migrate from one AZ to another, there is no need to restart the cluster. If multiple AZs are migrated, you must restart the cluster to make the change take effect. Value range: • multi_az_change: Add one or two AZs for a single-AZ cluster, or add one AZ for a dual-AZ cluster. • az_migrate: Migrate data from one AZ to another. Default value: N/A

Parameter	Mandatory	Type	Description
agency	Yes	String	Definition: Agency name. After a node is rebuilt, NICs need to be reattached to the new node. This means you need to have VPC permissions. Select an IAM agency to grant the current account the permission to access and use VPC. Constraints: The value can contain only letters (a-z and A-Z), digits (0-9), hyphens (-), and underscores (_). It cannot contain Chinese characters or other special characters. Value range: N/A Default value: N/A
indices_backup_check	No	Boolean	Definition: Whether to check full index snapshots. You are advised to select Check full index snapshot. This ensures that all data has been backed up, so that in case the task fails, the data can be restored using this snapshot. Constraints: N/A Value range: • true: Check for full index snapshots. • false: Do not check for full index snapshots. Default value: true

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Change AZs for the current cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/inst-type/all/azmigrate
```

```
{
  "source_az": "cn-north-4c",
  "target_az": "cn-north-4a",
  "migrate_type": "az_migrate",
  "agency": "css-test-agency"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. The client should not repeat the request without modifications.
403	The request is rejected. The server has received the request and understood it, but the server is refusing to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.5 Word Dictionary Management

5.5.1 Querying the Status of a Custom Word Dictionary

Function

You can configure custom word dictionaries to support word segmentation. This gives the search engine enhanced performance when searching by keywords such as company names, for example, Huawei, and buzzwords from social media. You can also search text data based on a synonym dictionary. CSS uses the IK and synonym analyzers. The IK analyzer uses a main word dictionary and a stop word dictionary. The synonym analyzer uses a synonym word dictionary. The IK analyzer uses the ik_max_word and ik_smart word segmentation policies. The synonym

analyzer uses the ik_synonym word segmentation policy. This API is used to query the loading status of a custom word dictionary.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/thesaurus

Table 5-126 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose word dictionary status you want to query For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-127 Response body parameters

Parameter	Type	Description
status	String	Definition: Custom word dictionary loading status. Value range: • Init indicates initialization is in progress. • Loaded indicates that the loading succeeded. • Loading indicates that the loading is in progress. • Failed indicates that the loading failed. • NeedReboot indicates that a cluster restart is required. • Undefined indicates an undefined state.
bucket	String	Definition: The OBS bucket that stores the word dictionary file that is last updated. Value range: N/A
mainObj	String	Definition: Main word dictionary file. Value range: N/A
stopObj	String	Definition: Stop word dictionary file. Value range: N/A
synonymObj	String	Definition: Synonym dictionary file. Value range: N/A
staticMainObj	String	Definition: Static main word dictionary file. Value range: N/A

Parameter	Type	Description
staticStopObj	String	Definition: Static main word dictionary file. Value range: N/A
extraMainObj	String	Definition: Extra main word dictionary file. Value range: N/A
extraStopObj	String	Definition: Extra stop word dictionary. Value range: N/A
updateTime	String	Definition: Latest update time of a dictionary. Value range: N/A
updateDetails	String	Definition: Word dictionary update details. Value range: N/A
clusterId	String	Definition: ID of the cluster where a custom word dictionary you want to configure. Value range: N/A
id	String	Definition: ID of the word dictionary. Value range: N/A

Example Requests

Query the loading status of a custom word dictionary.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/thesaurus
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "status" : "Loaded",
  "id" : "e766bd5a-79b0-4d1a-8402-fdeb017a36d1",
  "bucket" : "test-bucket",
  "mainObj" : "word/main.txt",
  "stopObj" : "word/stop.txt",
  "synonymObj" : "word/synonym.txt",
  "staticMainObj" : "Unused",
  "staticStopObj" : "Unused",
  "extraMainObj" : "Unused",
  "extraStopObj" : "Unused",
  "updateTime" : 1521184757000,
  "updateDetails" : "allinstancesareloadedsuccessfully.",
  "clusterId" : "ea244205-d641-45d9-9dcb-ab2236bcd07e"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
500	The server is able to receive but unable to understand the request.

Error Codes

See [Error Codes](#).

5.5.2 Loading Custom Word Dictionaries

Function

You can configure custom word dictionaries to support word segmentation. This gives the search engine enhanced performance when searching by keywords such as company names, for example, Huawei, and buzzwords from social media. You can also search text data based on a synonym dictionary. CSS uses the IK and synonym analyzers. The IK analyzer uses a main word dictionary and a stop word dictionary. The synonym analyzer uses a synonym word dictionary. The IK analyzer uses the `ik_max_word` and `ik_smart` word segmentation policies. The synonym analyzer uses the `ik_synonym` word segmentation policy. This API is used to load a custom word dictionary stored in OBS. When the preset word dictionaries are inadequate for tokenization, you can use custom word dictionaries.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/thesaurus

Table 5-128 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster where a custom word dictionary you want to configure. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-129 Request body parameters

Parameter	Mandatory	Type	Description
bucket_name	Yes	String	Definition: OBS bucket where the word dictionary file is stored. Constraints: The storage class of the bucket must be standard or infrequently accessed. Archive storage is not supported. Value range: N/A Default value: N/A
main_object	No	String	Definition: Main word dictionary file. Constraints: • Must be a text file encoded in UTF-8 without BOM. Each line contains one word. The maximum file size is 100 MB. • Modify the parameters of at least one of the seven word dictionaries. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
stop_object	No	String	Definition: Stop word dictionary file. Constraints: • Must be a text file encoded in UTF-8 without BOM. Each line contains one word. The maximum file size is 100 MB. • Modify the parameters of at least one of the seven word dictionaries. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
synonym_object	No	String	Definition: Synonym dictionary file. Constraints: • Must be a text file encoded in UTF-8 without BOM. Each line contains one word. The maximum file size is 100 MB. • Modify the parameters of at least one of the seven word dictionaries. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
static_main_object	No	String	Definition: Static main word dictionary file. Constraints: • Must be a text file encoded in UTF-8 without BOM. Each line contains one group of related words. The maximum file size is 100 MB. • Modify the parameters of at least one of the seven word dictionaries. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Only new clusters created after this word dictionary function was brought online are supported. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
static_stop_object	No	String	Definition: Static stop word dictionary file. Constraints: • Must be a text file encoded in UTF-8 without BOM. Each line contains one group of related words. The maximum file size is 100 MB. • Modify the parameters of at least one of the seven word dictionaries. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Only new clusters created after this word dictionary function was brought online are supported. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
extra_main_object	No	String	Definition: Extra main word dictionary file. Constraints: • Must be a text file encoded in UTF-8 without BOM. Each line contains one group of related words. The maximum file size is 100 MB. • Modify the parameters of at least one of the seven word dictionaries. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Only new clusters created after this word dictionary function was brought online are supported. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
extra_stop_object	No	String	Definition: Extra stop word dictionary file. Constraints: • Must be a text file encoded in UTF-8 without BOM. Each line contains one group of related words. The maximum file size is 100 MB. • Modify the parameters of at least one of the seven word dictionaries. Note: Passing an empty "" character string will clear the word dictionary. Passing nothing or null will leave the word dictionary unchanged. Only new clusters created after this word dictionary function was brought online are supported. Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Enable and configure the word dictionary.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/thesaurus
```

```
{
  "bucket_name" : "test-bucket",
  "main_object" : "word/main.txt",
  "stop_object" : "word/stop.txt",
  "synonym_object" : "word/synonym.txt",
  "static_main_object" : "word/staticMain.txt",
  "static_stop_object" : "word/staticStop.txt",
  "extra_main_object" : "word/extraMain.txt",
  "extra_stop_object" : "word/extraStop.txt"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
403	Request rejected. The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.
500	The server is able to receive the request but unable to understand the request.

Error Codes

See [Error Codes](#).

5.5.3 Disabling a Word Dictionary

Function

You can configure custom word dictionaries to support word segmentation. This gives the search engine enhanced performance when searching by keywords such as company names, for example, Huawei, and buzzwords from social media. You can also search text data based on a synonym dictionary.

This API is used to clear custom word dictionaries (main word dictionary, stop word dictionary, and synonym dictionary).

Calling Method

For details, see [Calling APIs](#).

URI

DELETE /v1.0/{project_id}/clusters/{cluster_id}/thesaurus

Table 5-130 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose custom word dictionary you want to delete. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Delete a custom word dictionary.

```
DELETE https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/thesaurus
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
403	Request rejected. The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.
500	The server is able to receive but unable to understand the request.

Error Codes

See [Error Codes](#).

5.6 Kibana/Dashboards Public Network Access

5.6.1 Enabling Kibana/Dashboards Public Network Access

Function

This API is used to enable Kibana/Dashboards public network access.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/publickibana/open

Table 5-131 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose Kibana public network access you want to enable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-132 Request body parameters

Parameter	Mandatory	Type	Description
eip_size	Yes	Integer	Bandwidth. Unit: Mbit/s
elb_white_list	No	StartKibanaPublicReqElbWhitelist object	ELB whitelist.

Table 5-133 StartKibanaPublicReqElbWhitelist

Parameter	Mandatory	Type	Description
enable_white_list	Yes	Boolean	Whether to enable the whitelist function. • true: The whitelist is enabled. • false: The whitelist is disabled.
white_list	Yes	String	Whitelist.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Enable public network access for Kibana or OpenSearch Dashboards.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/publickibana/open
{
  "eip_size" : 5,
  "elb_white_list" : {
    "enable_white_list" : true,
    "white_list" : "192.168.0.xx"
  },
  "is_auto_pay" : 1
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.

Status Code	Description
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.6.2 Disabling Kibana/Dashboards Public Network Access

Function

This API is used to disable public network access to Kibana or OpenSearch Dashboards. [Kibana/Dashboards public network access cannot be disabled for yearly/monthly clusters by calling this API.] (tag:cmcc,ctc,hws,hws_hk)

NOTE

If you disable Kibana public network access and then re-enable it, the public IP address for accessing Kibana may change. Exercise caution.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/publickibana/close

Table 5-134 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose Kibana/Dashboards public network access you want to disable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-135 Request body parameters

Parameter	Mandatory	Type	Description
eip_size	No	Integer	Bandwidth. Unit: Mbit/s
elb_white_list	No	StartKibanaPublicReqElbWhitelist object	ELB whitelist.

Table 5-136 StartKibanaPublicReqElbWhitelist

Parameter	Mandatory	Type	Description
enable_white_list	Yes	Boolean	Whether to enable the whitelist function. • true: The whitelist is enabled. • false: The whitelist is disabled.
white_list	Yes	String	Whitelist.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Disable public network access to Kibana or OpenSearch Dashboards.

```
PUT https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/publickibana/close
{
  "eip_size" : 5,
  "elb_white_list" : {
    "enable_white_list" : true,
    "white_list" : "192.168.0.xx"
  }
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.

Status Code	Description
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.6.3 Modify Kibana/Dashboards Public Network Bandwidth

Function

This API is used to modify the public network bandwidth of Kibana or OpenSearch Dashboards.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/publickibana/bandwidth

Table 5-137 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster whose Kibana or Dashboards public network bandwidth you want to modify. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-138 Request body parameters

Parameter	Mandatory	Type	Description
band_width	Yes	UpdatePublicKibanaBandwidthReqBandWidth object	Bandwidth.

Table 5-139 UpdatePublicKibanaBandwidthReqBandWidth

Parameter	Mandatory	Type	Description
size	Yes	Integer	New bandwidth size.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Modify the Kibana/Dashboards public network bandwidth.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/publickibana/bandwidth
```

```
{
  "band_width" : {
    "size" : 5
  },
  "is_auto_pay" : 1
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.6.4 Modifying Kibana/Dashboards Public Network Access

Function

This API is used to modify the Kibana/Dashboards access permission by modifying the Kibana/Dashboards whitelist.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/publickibana/whitelist/update

Table 5-140 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose Kibana/Dashboards access permission you want to modify. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-141 Request body parameters

Parameter	Mandatory	Type	Description
white_list	Yes	String	Whitelisted IP addresses or network segments that can access Kibana. Separate them with commas (,), and each IP address or network segment must be unique.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

The latest whitelist.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/publickibana/whitelist/update
```

```
{
  "white_list" : "192.168.0.21"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.6.5 Disabling Kibana/Dashboards Public Network Access Control

Function

This API is used to disable public network access control for Kibana or OpenSearch Dashboards.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/publickibana/whitelist/close

Table 5-142 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose Kibana public network access control you want to disable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Disable public network access to Kibana/Dashboards.

```
PUT https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/publickibana/whitelist/close
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7 Log Management

5.7.1 Enabling Log Backup or Log Collection

Function

This API is used to enable log backup or log ingestion for a cluster. You can enable log backup and log ingestion separately.

- Log backup: Cluster logs are periodically synchronized to OBS buckets. You can download them for in-depth analysis at any time.
- Log ingestion: Cluster logs are transmitted in real time to the current cluster or another cluster on the same network. (Relevant version and network compatibility requirements must be met.) You can use Kibana or other visualization tools for log search and analysis.

Constraints

- Log ingestion is supported only when the Elasticsearch cluster version is 7.10.2 and the image version is no earlier than 7.10.2_24.2.0_x.x.x.

- Log ingestion is supported only when the OpenSearch cluster version is 1.3.6 or 2.19.0 and the image version is no earlier than x.x.x_24.2.0_x.x.x.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/logs/open

Table 5-143 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster for which log backup or ingestion is to be enabled. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-144 Query Parameters

Parameter	Mandatory	Type	Description
action	No	String	Definition: Configure log backup or log ingestion for a cluster. Constraints: N/A Value range: - base_log_collect: Enables log backup. - real_time_log_collect: Enables log ingestion. Default value: base_log_collect

Request Parameters

Table 5-145 Request body parameters

Parameter	Mandatory	Type	Description
agency	No	String	Definition: Agency name that authorizes CSS to store snapshots in specified OBS buckets. Constraints: This parameter is mandatory when action is set to base_log_collect. Value range: The value can contain only letters (a-z and A-Z), digits (0-9), hyphens (-), and underscores (_). It cannot contain Chinese characters or other special characters. Default value: N/A

Parameter	Mandatory	Type	Description
log_base_path	No	String	Definition: Storage path of backed up logs in the OBS bucket. The service will back up logs to the \${Cluster name} directory under the corresponding backup path. Constraints: This parameter is mandatory when action is set to base_log_collect. Value range: • The backup path cannot contain the following characters: \:*?"<> '{}. • Additionally, it cannot: • Start with a slash (/). • Start or end with a period (.) • Contain more than two consecutive slashes (/) or periods (.) • Exceed 512 characters. Default value: N/A
log_bucket	No	String	Definition: OBS bucket name for log backup. Constraints: This parameter is mandatory when action is set to base_log_collect. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
auto_enable	No	Boolean	Definition: Whether to enable automatic backup of cluster logs. Constraints: N/A Value range: • true: Enable automatic backup of cluster logs. • false: Disable automatic backup of cluster logs. Default value: false
period	No	String	Definition: Cluster log backup start time. Constraints: This parameter is mandatory when action is set to base_log_collect and auto_enable is set to true. Value range: Format: Greenwich Mean Time (GMT). Example: 00:00 GMT+08:00 Default value: N/A

Parameter	Mandatory	Type	Description
index_prefix	No	String	Definition: Index prefix used for log ingestion. The log index is named using the format "index prefix+Log ingestion date". Constraints: This parameter is mandatory when action is set to real_time_log_collect. Value range: The index name can contain 1 to 128 characters. Only digits, lowercase letters, underscores (_), and hyphens (-) are allowed. Default value: N/A
keep_days	No	Integer	Definition: Log retention period. Ingested logs will be deleted upon expiration of this period. Constraints: This parameter is mandatory when action is set to real_time_log_collect. Value range: 1-3650 Default value: N/A

Parameter	Mandatory	Type	Description
target_cluster_id	No	String	Definition: ID of the destination cluster for log storage. Constraints: - This parameter is mandatory when action is set to real_time_log_collect. - The destination cluster can be the current cluster or another cluster. If you select another cluster, you need to use the Log Collection Connectivity Test API to verify whether the cluster can be used as the destination cluster. If the status code is 200, it can be used as the destination cluster. Otherwise, it cannot. Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

- Enable the log backup function.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/logs/open
{
  "agency": "css_obs_agency",
  "log_base_path": "css_repository/logs",
  "log_bucket": "000-words",
  "auto_enable": true,
  "period": "00:00 GMT+08:00"
}
```

- Enable real-time log collection function.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/logs/open?
action=real_time_log_collect
{
```

```
"index_prefix" : "css_log",  
"keep_days" : 30,  
"target_cluster_id" : "4f3deec3-efa8-4598-bf91-560aad1377a3"  
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.2 Disabling the Log Backup or Collection Function

Function

This API is used to disable log backup or log ingestion for a cluster. When log backup or log ingestion is no longer required, you can disable it to reduce storage costs.

Constraints

- Disabling log backup does not automatically delete existing log backups. Instead, you need to manually delete them on the OBS console.
- Disabling log ingestion does not automatically remove ingested log data. Rather, it will be deleted by the system upon expiration of its retention period. You can also manually delete it from the target cluster before the retention period expires.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/logs/close

Table 5-146 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster for which log backup or ingestion is to be disabled. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-147 Query Parameters

Parameter	Mandatory	Type	Description
action	No	String	Definition: Disables log backup or log collection. Constraints: You can disable log backup or log collection only when it is enabled. For details about how to enable log backup or log ingestion, see Enabling Logs . Value range: - base_log_collect: Disables log backup. - real_time_log_collect: Stops log ingestion. Default value: base_log_collect

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

- Disable log backup.
PUT https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/logs/close? action=base_log_collect
- Disable log ingestion.
PUT https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/logs/close? action=real_time_log_collect

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.3 Querying the Log Backup Task List

Function

This API is used to query the log backup task list of a cluster. To be able to use this API, the target cluster must have log backup enabled, and log backup tasks were triggered earlier. This allows you to keep track of operations that were performed on the cluster.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/logs/records

Table 5-148 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster you want to query. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-149 Query Parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Definition: Start value for log query. The default value is 1, indicating that the query starts from the first log record. Constraints: N/A Value range: 1-1000 Default value: 1

Parameter	Mandatory	Type	Description
limit	No	Integer	Definition: Number of log records to be queried. The default value is 10, indicating that 10 log records are queried at a time. Constraints: N/A Value range: 1-1000 Default value: 10
status	No	String	Definition: Status of the log tasks to be queried. If this parameter is not specified, all log task records of the cluster are queried. Constraints: N/A Value range: • FAIL: failed log tasks. • NOT_FAIL: ongoing or successful log tasks. • If this parameter is not specified, all tasks are queried. Default value: Empty string: Tasks of all states are queried.

Request Parameters

None

Response Parameters

Status code: 200

Table 5-150 Response body parameters

Parameter	Type	Description
clusterLogRecord	Array of clusterLogRecord objects	Parameter description: List of cluster log backup tasks. Options: N/A
totalSize	Integer	Definition: Total number of log records. Value range: N/A

Table 5-151 clusterLogRecord

Parameter	Type	Description
id	String	Definition: Cluster log task ID, which is generated based on the system UUID. Value range: N/A
clusterId	String	Parameter description: Cluster ID. Options: N/A
createAt	String	Definition: Cluster log creation time. Value range: Format: Unix timestamp.
logPath	String	Parameter description: Storage path of backed up logs in the OBS bucket. Options: N/A
status	String	Definition: Cluster log task status. Value range: • RUNNING: The backup is ongoing. • SUCCESS: The backup is successful. • FAIL: The backup failed.

Parameter	Type	Description
finishedAt	Long	Definition: End time of cluster log creation. If it is still ongoing, the value is null. Value range: Format: Unix timestamp.
jobTypes	String	Definition: Type of a log creation task. Value range: - Manual: Manual backup. - Auto: Automatic backup.
failedMsg	String	Definition: Error returned for log creation. If the task status is not failed, this parameter is null. Value range: N/A
jobId	String	Definition: ID of a log creation task. Value range: N/A

Example Requests

Querying the First 10 Log Backup Tasks

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/logs/records
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "clusterLogRecord" : [ {
    "id" : "d455a541-597e-4846-a6be-baad0ea361b1",
    "clusterId" : "4213d908-f5dc-4633-8401-cfd7175fca0c",
    "createAt" : 1656042837000,
    "logPath" : "css-backup-1610678043608/css/log",
    "status" : "RUNNING",
    "finishedAt" : null,
    "jobTypes" : "Manual",
    "failedMsg" : null,
    "jobId" : "2c9080df7c171342017c5e0884f8011c"
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.4 Querying Basic Log Configurations

Function

This API is used to query the basic configuration of cluster log backup or real-time log collection. If the log backup or log collection function has been enabled in the cluster, you can use this API to query the basic configurations of log backup and log collection.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/logs/settings

Table 5-152 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster you want to query. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-153 Query Parameters

Parameter	Mandatory	Type	Description
action	No	String	Parameter description: Queries the log backup or log collection settings of a cluster. Constraints: N/A Options: - base_log_collect: Queries cluster log backup settings. - real_time_log_collect: Queries cluster log collection settings. Default value: base_log_collect

Request Parameters

None

Response Parameters

Status code: 200

Table 5-154 Response body parameters

Parameter	Type	Description
logConfiguration	logConfiguration object	Parameter description: Log backup configuration. Options: N/A
realTimeLogCollectRecord	realTimeLogCollect object	Parameter description: Real-time log ingestion configuration. Options: N/A

Table 5-155 logConfiguration

Parameter	Type	Description
id	String	Parameter description: Log backup ID, which is generated using the system UUID. Options: N/A
clusterId	String	Parameter description: Cluster ID. Options: N/A
obsBucket	String	Parameter description: Name of the OBS bucket for storing logs. Options: N/A

Parameter	Type	Description
agency	String	Definition: Agency name. You can create an agency to allow CSS to call other cloud services. An agency name cannot contain special characters or Chinese characters. Value range: Only a-z, A-Z, 0-9, hyphens (-) and underscores (_) are allowed.
updateAt	Long	Definition: Update time. Value range: The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
basePath	String	Parameter description: Storage path of backed up logs in the OBS bucket. Options: N/A
autoEnable	Boolean	Definition: Whether to enable automatic backup. Value range: • true: Enable automatic backup. • false: Disable automatic backup.
period	String	Definition: Start time of automatic log backup. Value range: Format: Greenwich Mean Time (GMT). When autoEnable is set to false , the value of this parameter is null .
logSwitch	Boolean	Definition: Whether to enable logging. Value range: • true: Enable logging. • false: Disable logging.

Table 5-156 realTimeLogCollect

Parameter	Type	Description
id	String	Parameter description: Log collection ID, which is generated using the system UUID. Options: N/A
clusterId	String	Definition: Cluster ID Value range: N/A
indexPrefix	String	Parameter description: Prefix of the index for saving logs. Options: N/A
keepDays	Integer	Parameter description: Log retention duration. Options: N/A
targetClusterId	String	Parameter description: ID of the target cluster where logs are saved. Options: N/A

Parameter	Type	Description
status	String	Definition: Status of a real-time log ingestion task. Value range: • 100: A real-time log ingestion task is being created. • 150: Real-time log ingestion is available. • 200: The real-time log ingestion task is activated. • 300: Real-time log ingestion failed. • 302: The real-time log ingestion task failed to be deleted. • 303: The real-time log ingestion task failed to be created. • 304: The real-time log ingestion task is being disabled. • 400: The real-time log ingestion task is disabled.
createAt	Long	Parameter description: Start time of a real-time log collection task. Options: N/A
updateAt	Long	Parameter description: Update time of a real-time log collection task. Options: N/A

Example Requests

- Query the cluster's log backup settings.
GET https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/logs/settings?action=base_log_collect
- Query the cluster's log ingestion settings.
GET https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/logs/settings?action=real_time_log_collect

Example Responses

Status code: 200

Request succeeded.

- Example response to a request to query the cluster's log backup settings.

```
{
  "logConfiguration" : {
    "id" : "00032118-aff5-40e8-b19a-dd4bb576e572",
    "clusterId" : "e3201ceb-1a3e-49f3-bb2f-23a816440b20",
    "obsBucket" : "css-autobk-notdel-cn-north-7",
    "agency" : "css_obs_agency",
    "updateAt" : 1639624882000,
    "basePath" : "css/log",
    "autoEnable" : false,
    "period" : null,
    "logSwitch" : false
  },
  "realTimeLogCollectRecord" : null
}
```

- Example response to a request to query the cluster's log ingestion settings.

```
{
  "logConfiguration" : null,
  "realTimeLogCollectRecord" : {
    "id" : "17939b7b-5a93-4ca5-8d3c-b9f2d0e715b4",
    "clusterId" : "e3201ceb-1a3e-49f3-bb2f-23a816440b20",
    "keepDays" : 30,
    "updateAt" : 1717666418870,
    "createAt" : 1717666405897,
    "targetClusterId" : "8c19644b-f2ef-44fc-abef-230a4c578ce7",
    "indexPrefix" : "aaa",
    "status" : "200"
  }
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.5 Modifying Basic Settings for Log Backup or Log Collection

Function

This API is used to modify basic configurations for cluster log backup or real-time log collection. If the log backup or log collection function has been enabled in the cluster, you can use this API to modify the basic configurations of log backup and log collection.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/logs/settings

Table 5-157 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose basic log configurations you want to modify. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-158 Query Parameters

Parameter	Mandatory	Type	Description
action	No	String	Parameter description: Modifies the settings of log backup or log collection. Constraints: N/A Options: - base_log_collect: Modifies log backup settings. - real_time_log_collect: Modifies log collection settings. Default value: base_log_collect

Request Parameters

Table 5-159 Request body parameters

Parameter	Mandatory	Type	Description
agency	No	String	Definition: Agency name that authorizes CSS to store snapshots in specified OBS buckets. Constraints: This parameter is mandatory when action is set to base_log_collect. Value range: The value can contain only letters (a-z and A-Z), digits (0-9), hyphens (-), and underscores (_). It cannot contain Chinese characters or other special characters. Default value: N/A

Parameter	Mandatory	Type	Description
log_base_path	No	String	Definition: Storage path of backed up logs in the OBS bucket. The service will back up logs to the \${Cluster name} directory under the corresponding backup path. Constraints: This parameter is mandatory when action is set to base_log_collect. Value range: • The backup path cannot contain the following characters: \:*?"<> '{}. • Additionally, it cannot: • Start with a slash (/). • Start or end with a period (.). • Contain more than two consecutive slashes (/) or periods (.). • Exceed 512 characters. Default value: N/A
log_bucket	No	String	Definition: OBS bucket name for log backup. Constraints: This parameter is mandatory when action is set to base_log_collect. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
auto_enable	No	Boolean	Definition: Whether to enable automatic backup of cluster logs. Constraints: N/A Value range: • true: Enable automatic backup of cluster logs. • false: Disable automatic backup of cluster logs. Default value: false
period	No	String	Definition: Cluster log backup start time. Constraints: Format: Greenwich Mean Time (GMT). Value range: N/A Default value: N/A
index_prefix	No	String	Definition: Index prefix for storing logs. Constraints: This parameter is mandatory when action is set to real_time_log_collect. Value range: The index name can contain 1 to 128 characters. Only digits, lowercase letters, underscores (_), and hyphens (-) are allowed. Default value: N/A

Parameter	Mandatory	Type	Description
keep_days	No	Integer	Parameter description: Number of days for log retention. Constraints: This parameter is mandatory when action is set to real_time_log_collect. Options: 1-3650 Default value: N/A
target_cluster_id	No	String	Definition: ID of the destination cluster for log storage. Constraints: • This parameter is mandatory when action is set to real_time_log_collect. • The destination cluster can be the current cluster or another cluster. If logs are saved to another cluster, you need to use the connectivity testing API to test connectivity to the destination cluster. If status code 200 is returned, the destination cluster can be reached. Otherwise, you need to select another cluster. Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

- Modify log backup settings.

POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/logs/settings

```
{
  "agency" : "css_obs_agency",
  "log_base_path" : "css_repository/logs",
  "log_Bucket" : "000-words",
  "auto_enable" : true,
  "period" : "00:00 GMT+08:00"
}
```

- Modify real-time log ingestion settings.

POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/logs/settings?action=real_time_log_collect

```
{
  "index_prefix" : "css_log",
  "keep_days" : 30,
  "target_cluster_id" : "4f3deec3-efa8-4598-bf91-560aad1377a3",
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.6 Enabling the Automatic Log Backup Policy

Function

CSS provides log query and log backup, enabling you to easily manage and analyze logs to efficiently locate faults, optimize performance, and enhance system security.

- **Log query:** On the log management page of the CSS management console, you can query the latest 10,000 log records by node type and other criteria. A maximum of 100 results are displayed, enabling you to quickly locate issues.
- **Log backup:** Cluster logs are periodically synchronized to OBS buckets. You can download them for in-depth analysis at any time. You can configure custom log backup policies by specifying backup schedules and storage locations. The system backs up all critical logs, including run logs and deprecation logs. They provide comprehensive data for auditing and troubleshooting purposes.

This API is used to enable an automatic log backup policy. If log backup is enabled for a cluster, you can use this API to configure an automatic backup policy to periodically back up cluster logs to a specified OBS bucket and directory.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/logs/policy/update

Table 5-160 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster for which automatic log backup is to be enabled. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-161 Request body parameters

Parameter	Mandatory	Type	Description
period	Yes	String	Definition: Log backup start time. Format: Greenwich Mean Time (GMT). Constraints: N/A Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Enable the automatic log backup policy.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/logs/policy/update
{
```

```
"period" : "00:00 GMT+08:00"  
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.7 Disabling the Automatic Log Backup Policy

Function

CSS provides log query and log backup, enabling you to easily manage and analyze logs to efficiently locate faults, optimize performance, and enhance system security.

- Log query: On the log management page of the CSS management console, you can query the latest 10,000 log records by node type and other criteria. A maximum of 100 results are displayed, enabling you to quickly locate issues.
- Log backup: Cluster logs are periodically synchronized to OBS buckets. You can download them for in-depth analysis at any time. You can configure custom log backup policies by specifying backup schedules and storage locations. The system backs up all critical logs, including run logs and deprecation logs. They provide comprehensive data for auditing and troubleshooting purposes.

This API is used to disable an automatic log backup policy. If an automatic log backup policy is enabled for a cluster, you can use this API to disable it.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/logs/policy/close

Table 5-162 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose automatic log backup policy you want to disable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Disable an automatic log backup policy.

PUT https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/logs/policy/close

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.8 Backing Up Logs

Function

This API is used to back up logs. If log backup is enabled for a cluster, you can call this API to back up the cluster's run logs to the specified OBS bucket. This ensures traceability, security, and reliability of cluster data.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/logs/collect

Table 5-163 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose logs you want to back up. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Back up logs.
POST https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/logs/collect

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.9 Searching for Logs

Function

CSS supports log backup and query, helping you efficiently locate issues. You can periodically back up cluster logs to OBS buckets for long-term storage, from which you download log data for troubleshooting or auditing purposes. You can also query cluster logs through an API. You can search for specific run logs by specifying a time period, node names, and log levels. This helps you quickly locate faults or performance bottlenecks. This API is used to query cluster logs by specifying the node name, log level, and log type.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/logs/search

Table 5-164 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster you want to query. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-165 Request body parameters

Parameter	Mandatory	Type	Description
instance_name	Yes	String	Definition: Node name. Obtain the name attribute in instances by referring to Querying Cluster Details . Constraints: N/A Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
level	No	String	Definition: Queried log level. Constraints: • When log_type is set to deprecation, only ALL is supported. • This parameter is invalid when a Logstash cluster is queried. By default, all log levels are queried for a Logstash cluster. Value range: • DEBUG: Queries DEBUG logs. • INFO: Queries INFO logs. • WARN: Queries WARN logs. • ERROR: Queries ERROR logs. • ALL: all log levels, including DEBUG, INFO, WARN, ERROR, and TRACE.

Parameter	Mandatory	Type	Description
			NOTE • Separate multiple log levels using vertical bars (), for example, WARN ERROR, indicating that WARN and ERROR logs will be queried. Do not use this method with ALL. • To query trace logs, you need to first enable TRACE logs, and then query all log levels using ALL. Then, filter out TRACE logs. • To enable TRACE logs for an Elasticsearch cluster on Kibana, run the following command: • To enable TRACE logs for an OpenSearch cluster on Dashboards, run the following command: <pre>PUT _cluster/settings { "transient": { "logger.org.elasticsearch.transport. TransportService.tracer": "trace", "transport.tracer.include": " ", "http.tracer.include": "", "logger.org.elasticsearch.http.Http Tracer": "trace" } } PUT _cluster/settings { "transient": { "logger.org.opensearch.transport. TransportService.tracer": "trace", "transport.tracer.include": " ", "http.tracer.include": "", "logger.org.opensearch.http.HttpT racer": "trace" } } Default value: ALL</pre>

Parameter	Mandatory	Type	Description
log_type	Yes	String	Definition: Queried log type. Constraints: N/A Value range: • deprecation: Queries deprecation logs. (Deprecation operation logs can be queried only when all log levels are selected.) • indexingSlow: Queries slow indexing logs. • searchSlow: Queries slow query logs. • instance: Queries run logs. Default value: N/A
limit	No	Integer	Definition: Number of log records returned. By default, 100 log records are returned. A maximum of 10,000 log records can be returned, and the log size cannot exceed 1 MB. Constraints: N/A Value range: 1-10000 Default value: 100

Parameter	Mandatory	Type	Description
time_index	No	String	Definition: Returns logs generated before a specified time. Constraints: The time format is yyyy-MM-ddTHH:mm:ss,SSS. Value range: N/A Default value: N/A
keyword	No	String	Definition: Keyword used for filtering. Constraints: N/A Value range: Only letters, digits, underscores (_), hyphens (-), periods (.), spaces, and square brackets are allowed. Maximum length: 64 characters. Default value: N/A

Response Parameters

Status code: 200

Table 5-166 Response body parameters

Parameter	Type	Description
logList	Array of logList objects	Parameter description: Log list. Options: N/A

Parameter	Type	Description
type	String	Definition: Queried log type. Value range: - deprecation: Queries deprecation logs. - indexingSlow: Queries slow indexing logs. - searchSlow: Queries slow query logs. - instance: Queries run logs.
completed	Boolean	Definition: Whether all log files have been queried. Value range: - true: All log files have been queried. There are no more results. - false: Only some of the log files have been queried. The query result is returned because the number of requested records has been reached or the log size has reached 1 MB.

Table 5-167 logList

Parameter	Type	Description
content	String	Parameter description: Log content. Options: N/A
date	String	Definition: Date Value range: N/A

Parameter	Type	Description
level	String	Definition: Queried log level. Value range: • DEBUG: DEBUG logs. • INFO: INFO logs. • WARN: WARN logs. • ERROR: ERROR logs. • TRACE: TRACE logs.

Example Requests

Query logs by node name, log level, and log type.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/logs/search
{
  "instance_name": "css-4312-ess-esn-1-1",
  "level": "INFO",
  "log_type": "instance",
  "limit": 10,
  "time_index": "2025-08-25T01:07:19,082",
  "keyword": "received"
}
```

Example Responses

Status code: 200
Request succeeded.

```
{
  "logList": [ {
    "content": "xxxxxx",
    "date": "2021-10-08T03:55:54,718",
    "level": "INFO"
  } ],
  "type": "instance"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.

Status Code	Description
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.7.10 Log Ingestion Connectivity Testing

Function

This API is used to test the connectivity between two clusters. CSS clusters provide a log collection function, allowing users to collect clusters' run logs in real time and save them to a specified index of a specified CSS cluster. The target cluster can be the current cluster or another cluster. If logs are to be saved to another cluster, you need to use the connectivity testing API to test connectivity to the target cluster. If the target cluster cannot be reached, you need to select another cluster.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/logs/connectivity

Table 5-168 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be tested. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-169 Request body parameters

Parameter	Mandatory	Type	Description
target_cluster_id	Yes	String	Definition: Target cluster ID. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: Only cluster IDs are supported. Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Test connectivity to the target cluster.

```
POST /v1.0/6204a5bd270343b5885144cf9c8c158d/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/logs/connectivity
{
  "target_cluster_id": "4f3deec3-efa8-4598-bf91-560aad1377a4"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	The request could not be understood by the server due to malformed syntax. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.8 Public Network Access

5.8.1 Enabling Public Network Access

Function

Public network access is supported only when Security Mode and HTTPS Access are enabled for a cluster. When Public IP Address is enabled, a public IP address is automatically assigned, which will enable access to the security cluster from the public network. Additionally, you can configure access control from the public network by IP addresses or IP address ranges.

To enable public network access for Elasticsearch or OpenSearch clusters, a shared load balancer is typically used for load balancing. If your workloads require quicker access, you are advised to use a dedicated load balancer to connect to your clusters. For details about its configuration, see section "Configuring a Dedicated Load Balancer for an Elasticsearch Cluster."

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/public/open

Table 5-170 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster whose public network access you want to enable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-171 Request body parameters

Parameter	Mandatory	Type	Description
eip	Yes	BindPublicNetworkEip object	Definition: EIP for public network access. Constraints: N/A Value range: N/A Default value: N/A
white_list	No	String	Definition: Public network access control whitelist. Constraints: Separate the whitelisted CIDR blocks or IP addresses with commas (,), and make sure each of them is unique. Value range: N/A Default value: N/A

Table 5-172 BindPublicReqEip

Parameter	Mandatory	Type	Description
band_width	Yes	BindPublicReqEipBandWidth object	Definition: Public network bandwidth. Constraints: N/A Value range: N/A Default value: N/A

Table 5-173 BindPublicReqEipBandWidth

Parameter	Mandatory	Type	Description
size	Yes	Integer	Definition: Public network bandwidth, in Mbit/s. Constraints: N/A Value range: 1-200 Default value: N/A

Response Parameters

Status code: 200

Table 5-174 Response body parameters

Parameter	Type	Description
action	String	Definition: A setting required to enable public network access. Value range: The fixed value is bindZone , indicating that the binding is successful.

Example Requests

Example Responses

Status code: 200

Request succeeded.

```
{  
  "action" : "bindZone"  
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.8.2 Disabling Public Network Access

Function

This API is used to disable public network access. Public network access cannot be disabled for yearly/monthly clusters by calling this API.

NOTE

After public network access is disabled, users can no longer access the cluster via a public IP address. If you disable public network access and then re-enable it, the public IP address of the cluster may change. Exercise caution.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/public/close

Table 5-175 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose public network access you want to disable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-176 Request body parameters

Parameter	Mandatory	Type	Description
eip	No	UnBindPublicReqEipReq object	EIP information.

Table 5-177 UnBindPublicReqEipReq

Parameter	Mandatory	Type	Description
bandWidth	No	BindPublicReqEipBandWidth object	Definition: Public network bandwidth. Constraints: N/A Value range: N/A Default value: N/A

Table 5-178 BindPublicReqEipBandWidth

Parameter	Mandatory	Type	Description
size	Yes	Integer	Definition: Public network bandwidth, in Mbit/s. Constraints: N/A Value range: 1-200 Default value: N/A

Response Parameters

Status code: 200

Table 5-179 Response body parameters

Parameter	Type	Description
action	String	Operations. The fixed value is unbindZone , indicating that the unbinding is successful.

Example Requests

Disable public network access.

```
PUT https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/public/close
```

```
{
  "eip": {
    "bandWidth": {
      "size": 5
    }
  }
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "action": "unbindZone"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.8.3 Modifying Public Network Access Bandwidth

Function

This API is used to modify the public network access bandwidth.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/public/bandwidth

Table 5-180 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose public network access bandwidth you want to modify. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-181 Request body parameters

Parameter	Mandatory	Type	Description
band_width	Yes	BindPublicReqEipBandWidth object	Definition: Public network bandwidth. Constraints: N/A Value range: N/A Default value: N/A

Table 5-182 BindPublicReqEipBandWidth

Parameter	Mandatory	Type	Description
size	Yes	Integer	Definition: Public network bandwidth, in Mbit/s. Constraints: N/A Value range: 1-200 Default value: N/A

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Modify public network access bandwidth.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/public/
bandwidth

{
  "band_width" : {
    "size" : 5
  },
  "is_auto_pay" : 1
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.

Status Code	Description
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.8.4 Enabling the Public Network Access Control Whitelist

Function

This API is used to enable the public network access control whitelist.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/public/whitelist/update

Table 5-183 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster whose public network access control whitelist you want to enable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-184 Request body parameters

Parameter	Mandatory	Type	Description
white_list	Yes	String	IP address of the user for whom the whitelist is enabled.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Enable the public network access control whitelist.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/public/whitelist/update
{
  "white_list" : "192.168.0.xx"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.8.5 Disabling the Public Network Access Control Whitelist

Function

This API is used to disable the public network access control whitelist.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/public/whitelist/close

Table 5-185 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose public network access control whitelist you want to disable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Disable the public network access control whitelist.

```
PUT https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/public/whitelist/close
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.9 Snapshot Management

5.9.1 Enabling or Modifying the Cluster Snapshot Function

Function

This API is used to enable or modify the cluster snapshot function. Cluster snapshots enhance data protection and restoration capabilities for clusters. By storing a snapshot in an OBS bucket, you save a point-in-time copy of the cluster's data. By restoring this snapshot, you can restore the cluster to a previous state.

When enabling or modifying the cluster snapshot function, you can enable or disable automatic snapshot creation.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/index_snapshot/setting

Table 5-186 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: Cluster ID. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-187 Request body parameters

Parameter	Mandatory	Type	Description
bucket	Yes	String	Definition: Name of the OBS bucket that stores snapshots. Constraints: • The storage class must be Standard. • The region must be the same as the region where the cluster is located. Value range: N/A Default value: N/A
agency	Yes	String	Definition: Agency name. Select an IAM agency to grant the current account the permission to access and use OBS. To store snapshots to an OBS bucket, you must have the required OBS access permissions. Constraints: The agency name can contain only letters (case-sensitive), digits, underscores (_), and hyphens (-). Otherwise, the backup will fail. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
base_path	No	String	Definition: Storage path of the snapshot in the OBS bucket. Constraints: N/A Value range: • The backup path cannot contain the following characters: \:*?"<> '{}. Additionally, it cannot: • Start with a slash (/). • Start or end with a period (.). • Contain more than two consecutive slashes (/) or periods (.). • Exceed 512 characters. Default value: N/A

Parameter	Mandatory	Type	Description
max_snapshot_bytes_per_seconds	No	String	Definition: This parameter sets the maximum backup speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Constraints: N/A Value range: The following values and formats are allowed: • Number + Unit • 0 • -1 The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). Setting this parameter to 0mb, 0, -1 means there is no speed limit. An overly high backup speed may lead to excessive resource usage, which may impact system stability. Default value: The default value is 40 MB. If this parameter is left blank, the default value is used.

Parameter	Mandatory	Type	Description
max_restore_bytes_per_seconds	No	String	Definition: This parameter sets the maximum recovery speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Constraints: For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter. If Maximum recovery speed (per second) is lower than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the former is reached. If Maximum recovery speed (per second) is higher than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the latter is reached. Value range: The following values and formats are allowed: • Number + Unit • 0 • -1 The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). Setting this parameter to 0mb, 0, -1 means there is no speed limit. (However, for OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed

Parameter	Mandatory	Type	Description
			is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter.) An overly high recovery speed may lead to excessive resource usage, which may impact system stability. Default value: For Elasticsearch clusters of 7.6.2 or earlier, the default setting is 40mb. For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the default setting is no limit, but the recovery speed is still limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter. If this parameter is left blank, the default setting is used.
enable	No	String	Definition: Whether to enable automatic snapshot creation. Constraints: N/A Value range: • true: Enable automatic snapshot creation. • false: Disable automatic snapshot creation. Default value: false

Parameter	Mandatory	Type	Description
indices	No	String	Definition: Name of the index to be backed up. - You can specify indexes you want to back up. To specify multiple indexes, use commas (,) to separate them, for example, index1,index2,index3. - You can use an asterisk (*) <i>to match multiple indexes</i>. For example, <i>index</i> indicates that all indexes with the prefix index will be restored. Constraints: This parameter is valid only when enable is set to true. Value range: The value can contain 0 to 1,024 characters. It cannot contain uppercase letters, spaces, or the following special characters: "< >/? Default value: The default value is empty, indicating that all indexes will be backed up.

Parameter	Mandatory	Type	Description
prefix	No	String	Definition: Prefix of the name of an automatically created snapshot. A snapshot name consists of the snapshot name prefix and timestamp, for example, snapshot-1566921603720. Constraints: This parameter is mandatory when enable is set to true. Value range: Enter up to 32 characters and start with a lowercase letter. Lowercase letters, digits, hyphens (-), and underscores (_) are allowed. Default value: N/A
period	No	String	Definition: Time when a snapshot is created every day. Constraints: NOTE When frequency is set to HOUR, you do not need to specify period. By default, snapshot creation starts from the next hour after the settings are successful. Value range: Snapshots can only be created on the hour. The time format is the time followed by the time zone, specifically, HH:mm z. In the format, HH:mm refers to the hour time and z refers to the time zone. For example, 00:00 GMT+08:00 and 01:00 GMT+08:00. Default value: 00:00 GMT+08:00

Parameter	Mandatory	Type	Description
keepday	No	Integer	Definition: Customize the number of snapshots to be retained. Expired snapshots will be automatically deleted on the half hour. The deletion policy applies only to automated snapshots that are executed at the same frequency as the current automatic snapshot creation policy. Constraints: • This parameter is mandatory when enable is set to true. • The number of retained snapshots depends on the execution frequency and indexes set in the automatic snapshot creation policy. If the execution interval is short or the index data volume is large, the number of retained automatic snapshots may not reach the preset value. Value range: 1–90 Default value: N/A

Parameter	Mandatory	Type	Description
frequency	No	String	Definition: Frequency of automatically creating snapshots. Constraints: The number of retained snapshots depends on the execution frequency and indexes set in the automatic snapshot creation policy. If the execution interval is short or the index data volume is large, the number of retained automatic snapshots may not reach the preset value. Value range: • HOUR: Execute once every hour on the hour. • DAY: Execute once every day. • SUN, MON, TUE, WED, THU, FRI, and SAT: Execute the task at the specified day of every week. For example, SUN indicates that the task is executed once every Sunday. Default value: DAY

Parameter	Mandatory	Type	Description
delete_auto	No	String	Definition: Whether to clear all the automatically created snapshots when disabling automatic snapshot creation. Constraints: This parameter takes effect only when enable is set to false. Value range: true: Delete all automatically created snapshots when automatic snapshot creation is disabled. false: Do not delete snapshots that were automatically created. Default value: false

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Enable the snapshot function but do not enable automatic snapshot creation.

POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/index_snapshot/setting

```
{
  "bucket" : "test-bucket",
  "agency" : "usearch",
  "base_path" : "css_repository/Es-name",
  "max_snapshot_bytes_per_seconds" : 0,
  "max_restore_bytes_per_seconds" : 0
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
406	The server could not fulfill the request according to the content characteristics of the request.
412	The server did not meet one of the preconditions contained in the request.
504	A gateway timeout error occurred.

Error Codes

See [Error Codes](#).

5.9.2 Manually Creating a Snapshot

Function

CSS allows you to use snapshots to back up and restore Elasticsearch cluster data. By storing a snapshot in an OBS bucket, you save a point-in-time copy of the cluster's data. By restoring this snapshot, you can restore the cluster to a previous state. There are two ways to create snapshots to back up a CSS cluster: automatic and manual.

- Automatic snapshot creation: Snapshots are created periodically based on a preset time policy, for example, daily or weekly, to ensure continuous data protection. This reduces manual effort and improves backup reliability and efficiency.
- Manual snapshot creation: You can manually create snapshots when necessary, for example, before performing a mission-critical operation (such as a cluster upgrade), so you can use these snapshots to quickly restore the cluster to a previous state in case anything goes wrong. Manual snapshots provide more flexibility.

This API is used to manually create a snapshot.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/index_snapshot

Table 5-188 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster for which you want to create a snapshot. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-189 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Definition: Snapshot name. Constraints: N/A Value range: Snapshot name. Enter 4 to 64 characters. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed. The value must start with a letter. Default value: N/A
description	No	String	Definition: Snapshot description. Constraints: N/A Value range: The value contains 0 to 256 characters and cannot contain !<>=&"" Default value: N/A

Parameter	Mandatory	Type	Description
indices	No	String	Definition: Name of the index to be backed up. - You can specify indexes you want to back up. To specify multiple indexes, use commas (,) to separate them, for example, index1,index2,index3. - You can use an asterisk (*) <i>to match multiple indexes</i>. For example, <i>index</i> indicates that all indexes with the prefix index will be restored. Constraints: This parameter is valid only when enable is set to true . Value range: The value can contain 0 to 1,024 characters. It cannot contain uppercase letters, spaces, or the following special characters: "< >/?". Default value: The default value is empty, indicating that all indexes will be backed up.

Response Parameters

Status code: 201

Table 5-190 Response body parameters

Parameter	Type	Description
backup	backupRsp object	Definition: The snapshot information. Value range: N/A

Table 5-191 backupRsp

Parameter	Type	Description
id	String	Definition: Snapshot ID. Value range: N/A
name	String	Definition: Snapshot name. Value range: Snapshot name. Enter 4 to 64 characters. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed. The value must start with a letter.

Example Requests

Create a snapshot.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/index_snapshot
{
  "name" : "snapshot_001",
  "indices" : "myindex1myindex2"
}
```

Example Responses

Status code: 201

Resource created.

```
{
  "backup" : {
    "id" : "9dc4f5c9-33c0-45c7-9378-ae35ae350682",
    "name" : "snapshot_101"
  }
}
```

Status Codes

Status Code	Description
201	Resource created.
400	Invalid request. Modify the request before retry.

Status Code	Description
403	Request rejected. The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.
406	The server could not fulfill the request according to the content characteristics of the request.
500	The server has received the request but could not understand it.
501	The server does not support the function required to fulfill the request.

Error Codes

See [Error Codes](#).

5.9.3 Restoring a Snapshot

Function

This API is used to restore a specified snapshot. Use a snapshot to restore a point-in-time copy of a cluster to the original cluster or to another eligible cluster.

Constraints

- Before snapshot restoration is completed, the indexes being restored to the destination cluster cannot be queried.
- If the two clusters have indexes with the same names but different shard structures, such indexes cannot be restored.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/index_snapshot/{snapshot_id}/restore

Table 5-192 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: Cluster ID. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
snapshot_id	Yes	String	Definition: The snapshot ID. For details, see Querying Snapshots . Constraints: The snapshot status is Available . Value range: N/A Default value: N/A

Request Parameters

Table 5-193 Request body parameters

Parameter	Mandatory	Type	Description
target_cluster	Yes	String	Definition: ID of the destination cluster to which the snapshot is restored. Constraints: • If the current cluster is available, snapshots are restored to the current cluster by default. Alternatively, you can choose to restore them to another specified cluster. • The destination cluster must be in the same region as the current cluster, and must have a later version than the current cluster. Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
indices	No	String	Definition: Specify the name of the index you want to restore. Constraints: The value is a string of 0 to 1024 characters that cannot contain uppercase letters, spaces, or the following special characters: "< >/?. • When restoring an index whose name is prefixed with .kibana, the index name must be specified. • The .opendistro_security index cannot be restored. Value range: • You can use an asterisk (*) <i>to match multiple indexes</i>. For example, index indicates that all indexes with the prefix index will be restored. When an asterisk (*) is used for index matching, the .opendistro_security index and any system indexes whose name is prefixed with .kibana are filtered out by default. • You can restore indexes by specifying their names, for example, index1,index2,index3. Default value: By default, this parameter is left blank. That is, no index name is specified, and all indexes will be restored.

Parameter	Mandatory	Type	Description
rename_pattern	No	String	Definition: Rule for defining the indexes to be restored. Constraints: - Indexes that match this rule will be restored. The filtering condition must be a regular expression. - The renamePattern and renameReplacement parameters must be set at the same time in order for them to take effect. Value range: The value can contain 0 to 1,024 characters. Uppercase letters, spaces, and the following special characters are not allowed: "\< >/?" Default value: N/A
rename_replacement	No	String	Definition: Rule for renaming an index. Constraints: The renamePattern and renameReplacement parameters must be set at the same time in order for them to take effect. Value range: The value can contain 0 to 1,024 characters. Uppercase letters, spaces, and the following special characters are not allowed: "\< >/? For example, restored_index_\$1 indicates adding the restored_ prefix to the names of all the restored indexes. Default value: N/A

Parameter	Mandatory	Type	Description
replace_exist_indices	No	Boolean	Definition: Whether to overwrite same-name indexes in the destination cluster. Restoring a snapshot will overwrite existing data in the destination cluster. When there are same-name indexes in the destination cluster, you need to select this option in order to restore same-name, same-shard structure indexes. Indexes with a different shard structure cannot be restored. Constraints: N/A Value range: • false: Do not replace existing indexes. • true: Replace existing indexes. Default value: false

Response Parameters

Status code: 201

Resource created.

None

Example Requests

Restore a snapshot.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/index_snapshot/653c086f-ebe7-4383-b902-c031b354f1dd/restore
```

```
{
  "target_cluster" : "ea244205-d641-45d9-9dcab2236bcd07e",
  "indices" : "myindex1,myindex2",
  "rename_pattern" : "index",
  "rename_replacement" : "restored_index_$1"
}
```

Example Responses

None

Status Codes

Status Code	Description
201	Resource created.
400	Invalid request. Modify the request before retry.
403	Request rejected. The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.9.4 Deleting a Snapshot

Function

Delete snapshots that are no longer needed to reclaim storage spaces.

This API is used to delete a snapshot.

NOTE

After a snapshot is deleted, its data cannot be restored. Exercise caution.

Constraints

After a snapshot is deleted, its data cannot be restored. Exercise caution.

Calling Method

For details, see [Calling APIs](#).

URI

DELETE /v1.0/{project_id}/clusters/{cluster_id}/index_snapshot/{snapshot_id}

Table 5-194 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to which the to-be-deleted snapshot belongs. For details about how to obtain the cluster ID, see Obtaining a Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
snapshot_id	Yes	String	Definition: ID of the snapshot you want to delete. For details, see Querying Snapshots . Constraints: N/A Value range: N/A Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Delete a snapshot.

```
DELETE https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/index_snapshot/29a2254e-947f-4463-b65a-5f0b17515fae
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
403	Request rejected. The server has received the request and understood it, but refused to respond to it. The client should not repeat the request without modifications.

Error Codes

See [Error Codes](#).

5.9.5 Querying the Automatic Snapshot Creation Policy

Function

CSS allows you to use snapshots to back up and restore Elasticsearch cluster data. By storing a snapshot in an OBS bucket, you save a point-in-time copy of the cluster's data. By restoring this snapshot, you can restore the cluster to a previous state. There are two ways to create snapshots to back up a CSS cluster: automatic and manual. This API is used to query the automatic snapshot creation policy.

- Automatic snapshot creation: Snapshots are automatically created daily or weekly based on a preset time schedule, ensuring ongoing data protection. By configuring an automatic snapshot policy, you reduce manual operations while improving backup reliability and efficiency.

- Manual snapshot creation: You create snapshots manually for special occasions, for example, prior to a high-risk operation (such as a cluster upgrade). This ensures you can restore data using snapshots in case anything should go wrong. Manual snapshots provide additional flexibility.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/index_snapshot/policy

Table 5-195 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster where snapshots are to be automatically created. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-196 Response body parameters

Parameter	Type	Description
keepday	Integer	Definition: Customize the number of snapshots to be retained. Expired snapshots will be automatically deleted on the half hour. The deletion policy applies only to automated snapshots that are executed at the same frequency as the current automatic snapshot creation policy. Value range: 1-90
period	String	Definition: Time when a snapshot is created every day. Value range: Snapshots can only be created on the hour. The time format is the time followed by the time zone, specifically, HH:mm z . In the format, HH:mm refers to the hour time and z refers to the time zone. For example, 00:00 GMT+08:00 and 01:00 GMT+08:00.
frequency	String	Definition: Frequency of automatically creating snapshots. Value range: • HOUR: Execute once every hour on the hour. • DAY: Execute once every day. • SUN, MON, TUE, WED, THU, FRI, and SAT: Execute the task at the specified day of every week. For example, SUN indicates that the task is executed once every Sunday.

Parameter	Type	Description
prefix	String	Definition: Prefix of the name of an automatically created snapshot. A snapshot name consists of the snapshot name prefix and timestamp, for example, snapshot-1566921603720. Value range: Enter up to 32 characters and start with a lowercase letter. Lowercase letters, digits, hyphens (-), and underscores (_) are allowed.
bucket	String	Definition: Name of the OBS bucket that stores snapshots. Value range: N/A
basePath	String	Definition: Storage path of the snapshot in the OBS bucket. Value range: • The backup path cannot contain the following characters: \:*?"<> '{}. • Additionally, it cannot: • Start with a slash (/). • Start or end with a period (.). • Contain more than two consecutive slashes (/) or periods (.). • Exceed 512 characters.
agency	String	Definition: Agency name. Select an IAM agency to grant the current account the permission to access and use OBS. To store snapshots to an OBS bucket, you must have the required OBS access permissions. Value range: N/A

Parameter	Type	Description
enable	String	Definition: Whether to enable automatic snapshot creation. Value range: • true: Enable automatic snapshot creation. • false: Disable automatic snapshot creation.
indices	String	Definition: Name of the index to be backed up. • You can specify indexes you want to back up. To specify multiple indexes, use commas (,) to separate them, for example, index1,index2,index3. • You can use an asterisk () <i>to match multiple indexes</i>. For example, <i>index</i> indicates that all indexes with the prefix index will be restored. Value range: The value can contain 0 to 1,024 characters. It cannot contain uppercase letters, spaces, or the following special characters: "< >/?"
snapshotCmkId	String	Definition: Snapshot key ID. Value range: N/A

Parameter	Type	Description
maxSnapshotBytesPerSeconds	String	Definition: This parameter sets the maximum backup speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual backup speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Value range: The following values and formats are allowed: • Number + Unit • 0 • -1 The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). Setting this parameter to 0mb, 0, -1 means there is no speed limit. An overly high backup speed may lead to excessive resource usage, which may impact system stability.

Parameter	Type	Description
maxRestoreBytes-PerSeconds	String	Definition: This parameter sets the maximum recovery speed per node (bytes per second). When it is exceeded, flow control is triggered to prevent excessive resource usage and ensure system stability. The actual recovery speed may not reach the configured value, as it depends on many factors, such as OBS performance and disk I/O. Constraints: For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter. If Maximum recovery speed (per second) is lower than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the former is reached. If Maximum recovery speed (per second) is higher than <code>indices.recovery.max_bytes_per_sec</code>, flow control is triggered when the latter is reached. Value range: The following values and formats are allowed: • Number + Unit • 0 • -1 The number ranges from 0 to 9999. The unit can be k, kb, m, mb, g, gb, t, tb, p, pb, or b (case-insensitive). Setting this parameter to 0mb, 0, -1 means there is no speed limit. (However, for OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the recovery speed is also limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter.) An overly high recovery speed may lead to excessive resource usage, which may impact system stability. Default value:

Parameter	Type	Description
		For Elasticsearch clusters of 7.6.2 or earlier, the default setting is 40mb. For OpenSearch clusters and Elasticsearch clusters later than 7.6.2, the default setting is no limit, but the recovery speed is still limited by the <code>indices.recovery.max_bytes_per_sec</code> parameter. If this parameter is left blank, the default setting is used.

Example Requests

Query the automatic snapshot creation policy.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/index_snapshot/policy
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "keepday" : 2,
  "period" : "16:00 GMT+08:00",
  "frequency" : "DAY",
  "prefix" : "snapshot",
  "bucket" : "es-backup",
  "basePath" : "css_repository/tests",
  "agency" : "usearch",
  "enable" : "true",
  "indices" : "*",
  "snapshotCmkId" : null,
  "maxSnapshotBytesPerSeconds" : "40MB",
  "maxRestoreBytesPerSeconds" : "0MB"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
406	The server could not fulfill the request according to the content characteristics of the request.

Error Codes

See [Error Codes](#).

5.9.6 Configuring the Automatic Snapshot Creation Policy

Function

This API is used to configure automatic snapshot creation policies. When automatic snapshot creation is enabled, snapshots are automatically created daily or weekly based on a preset time schedule, ensuring ongoing data protection. By configuring an automatic snapshot policy, you reduce manual operations while improving backup reliability and efficiency.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/index_snapshot/policy

Table 5-197 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster for which you want to enable automatic snapshot creation. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-198 Request body parameters

Parameter	Mandatory	Type	Description
indices	No	String	Definition: Name of the index to be backed up. Constraints: • If no index is specified, snapshots will be created for all indexes in the cluster. • Or you can back up indexes by specifying their names, for example, index1,index2,index3. • The value can contain a maximum of 1024 characters and cannot contain uppercase letters, spaces, or "< >/?". You can use an asterisk (<i>) to match multiple indexes. For example, index indicates that all indexes with the prefix index will be backed up.</i> Value range: The value can contain 0 to 1,024 characters. It cannot contain uppercase letters, spaces, or the following special characters: "< >/?". Default value: The default value is empty, indicating that all indexes will be backed up.

Parameter	Mandatory	Type	Description
prefix	No	String	Definition: Prefix of the name of an automatically created snapshot. A snapshot name consists of the snapshot name prefix and timestamp, for example, snapshot-1566921603720. Constraints: This parameter is mandatory when enable is set to true . Value range: Enter up to 32 characters and start with a lowercase letter. Lowercase letters, digits, hyphens (-), and underscores (_) are allowed. Default value: N/A
period	No	String	Definition: Time when a snapshot is created every day. Constraints: NOTE When frequency is set to HOUR, you do not need to specify period. By default, snapshot creation starts from the next hour after the settings are successful. Value range: Snapshots can only be created on the hour. The time format is the time followed by the time zone, specifically, HH:mm z . In the format, HH:mm refers to the hour time and z refers to the time zone. For example, 00:00 GMT+08:00 and 01:00 GMT+08:00. Default value: 00:00 GMT+08:00

Parameter	Mandatory	Type	Description
keepday	No	Integer	Definition: Customize the number of snapshots to be retained. Expired snapshots will be automatically deleted on the half hour. The deletion policy applies only to automated snapshots that are executed at the same frequency as the current automatic snapshot creation policy. Constraints: • This parameter is mandatory when enable is set to true. • The number of retained snapshots depends on the execution frequency and indexes set in the automatic snapshot creation policy. If the execution interval is short or the index data volume is large, the number of retained automatic snapshots may not reach the preset value. Value range: 1–90 Default value: N/A
enable	Yes	String	Definition: Whether to enable automatic snapshot creation. Constraints: N/A Value range: • true: Enable automatic snapshot creation. • false: Disable automatic snapshot creation. Default value: false

Parameter	Mandatory	Type	Description
delete_auto	No	String	Definition: Whether to clear all the automatically created snapshots when disabling automatic snapshot creation. Constraints: If automatic snapshot creation is disabled without deleting automatically created snapshots, all such snapshots in the cluster cannot be manually deleted later once automatic snapshot creation is enabled again. Instead, they can only be deleted by the system based on the snapshot lifecycle policy. Value range: • true: Delete all automatically created snapshots when automatic snapshot creation is disabled. • false: Do not delete snapshots that were automatically created. Default value: false

Parameter	Mandatory	Type	Description
frequency	No	String	Definition: Frequency of automatically creating snapshots. Constraints: The number of retained snapshots depends on the execution frequency and indexes set in the automatic snapshot creation policy. If the execution interval is short or the index data volume is large, the number of retained automatic snapshots may not reach the preset value. Value range: • HOUR: Execute once every hour on the hour. • DAY: Execute once every day. • SUN, MON, TUE, WED, THU, FRI, and SAT: Execute the task at the specified day of every week. For example, SUN indicates that the task is executed once every Sunday. Default value: DAY

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Configure the automatic snapshot creation policy.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/index_snapshot/policy
{
  "prefix" : "snapshot",
  "period" : "16:00 GMT+08:00",
  "keepday" : 7,
  "frequency" : "DAY",
  "enable" : "true",
}
```

```
"delete_auto" : "false"  
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
406	The server could not fulfill the request according to the content characteristics of the request.

Error Codes

See [Error Codes](#).

5.9.7 Querying a Snapshot List

Function

This API is used to query all snapshots of a cluster, including the snapshot creation time, data engine, snapshot ID, and snapshot status, helping you manage the lifecycle of cluster snapshots. This API enables you to quickly obtain the details of a specified snapshot, verify snapshot availability, and help you restore data, roll back versions, or analyze storage costs. This API also provides structured data support required for automatic monitoring and warning for the O&M team, significantly improving cluster O&M efficiency and stability.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/index_snapshots

Table 5-199 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose snapshots you want to query. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-200 Response body parameters

Parameter	Type	Description
backups	Array of ListSnapshotBackupsResp objects	Definition: Snapshot list. Value range: N/A

Table 5-201 ListSnapshotBackupsResp

Parameter	Type	Description
created	String	Definition: Snapshot creation time. Value range: The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
datastore	ListSnapshotBackupsDatastoreResponse object	Definition: Search engine. Value range: N/A
description	String	Definition: Snapshot description. Value range: N/A
id	String	Definition: Snapshot ID. Value range: N/A
clusterId	String	Definition: Cluster ID. Value range: N/A
clusterName	String	Definition: Cluster name. Value range: N/A
name	String	Definition: Snapshot name. Value range: N/A
status	String	Definition: Snapshot status. Value range: • **BUILDING**: creating. • COMPLETED: available. • FAILED: unavailable. • PART_FAILED: partially available.

Parameter	Type	Description
backupType	String	Parameter description: Snapshot type. Options: • 0: Automatic creation. • 1: Manual creation.
backupMethod	String	Parameter description: Snapshot creation mode. Options: • auto: automatic creation. • manual: manual creation.
backupFrequency	String	Definition: Cluster snapshot creation frequency. Value range: • hour • day • SUN: Sunday • MON: Monday • TUE: Tuesday • WED: Wednesday • THU: Thursday • FRI: Friday • SAT: Saturday
indices	String	Definition: Index you want to back up. Value range: N/A
totalShards	Integer	Definition: Total number of shards of the index you want to back up. Value range: N/A
failedShards	Integer	Definition: Number of shards that fail to be backed up. Value range: N/A

Parameter	Type	Description
restoreStatus	String	Definition: Snapshot restoration status. Value range: • restoring: restoration in progress. • success: successful restoration. • failed: restoration failed.
startTime	Long	Definition: Snapshot start timestamp. Value range: N/A
endTime	Long	Definition: Snapshot end timestamp. Value range: N/A
bucketName	String	Definition: Name of the bucket that stores snapshot data. Value range: N/A

Table 5-202 ListSnapshotBackupsDatastoreResp

Parameter	Type	Description
type	String	Definition: Cluster engine type. Value range: • elasticsearch: Elasticsearch cluster. • openSearch: OpenSearch cluster.
version	String	Definition: Elasticsearch engine version. For details, see the supported versions in Before You Start . Value range: N/A

Example Requests

Query the cluster snapshot list.

GET https://{Endpoint}/v1.0/{project_id}/clusters/37cb1075-c38e-4cd8-81df-442d52df3786/index_snapshots

Example Responses

Status code: 200

Request succeeded.

```
{
  "backups": [ {
    "created": "2018-03-07T07:34:47",
    "datastore": {
      "type": "elasticsearch",
      "version": "x.x.x"
    },
    "description": "",
    "id": "e29d99c1-3d19-4ea4-ae8d-f252df76cbe9",
    "clusterId": "37cb1075-c38e-4cd8-81df-442d52df3786",
    "clusterName": "Es-xfx",
    "name": "snapshot-002",
    "status": "COMPLETED",
    "updated": "2018-03-07T07:40:12",
    "backupType": "1",
    "backupMethod": "manual",
    "backupFrequency": "HOURLY",
    "indices": ".kibanawebsite2",
    "totalShards": 6,
    "failedShards": 0,
    "version": "x.x.x",
    "restoreStatus": "success",
    "startTime": 1520408087099,
    "endTime": 1520408412219,
    "bucketName": "obs-b8ed"
  }, {
    "created": "2018-03-06T15:42:37",
    "datastore": {
      "type": "elasticsearch",
      "version": "x.x.x"
    },
    "description": "",
    "id": "29a2254e-947f-4463-b65a-5f0b17515fae",
    "clusterId": "37cb1075-c38e-4cd8-81df-442d52df3786",
    "clusterName": "Es-xfx",
    "name": "snapshot-001",
    "status": "COMPLETED",
    "backupType": "1",
    "backupMethod": "manual",
    "indices": ".kibana",
    "totalShards": 1,
    "failedShards": 0,
    "restoreStatus": "none",
    "startTime": 1520350957275,
    "endTime": 1520351284357,
    "bucketName": "obs-b8ed"
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.

Status Code	Description
406	The server could not fulfill the request according to the content characteristics of the request.

Error Codes

See [Error Codes](#).

5.9.8 Disabling Cluster Snapshots

Function

This API is used to disable cluster snapshots. If cluster snapshots are no longer required, you can disable them to reduce storage costs.

Constraints

Disabling cluster snapshots does not automatically delete existing snapshots. Instead, you need to manually delete them on the OBS console.

Calling Method

For details, see [Calling APIs](#).

URI

DELETE /v1.0/{project_id}/clusters/{cluster_id}/index_snapshots

Table 5-203 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: Specifies the cluster ID. For details, see Obtaining a Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Disable cluster snapshots.

```
DELETE https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/index_snapshots
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
406	The server could not fulfill the request according to the content characteristics of the request.

Error Codes

See [Error Codes](#).

5.9.9 (Not Recommended) Automatically Setting Basic Configurations of a Cluster Snapshot

Function

This API is used to configure basic settings for automatic cluster snapshot creation, including automatically creating an IAM agency and the OBS bucket and file path used for storing snapshots. This API also enables cluster snapshots. Cluster snapshots enhance data protection and restoration capabilities for clusters.

Constraints

- This API automatically creates an OBS bucket and agency for the snapshot. If there are multiple clusters, an OBS bucket will be created for each cluster via this API. As a result, the OBS quota may be insufficient, and many OBS buckets are difficult to maintain. You are advised to use the API [Enabling or Modifying the Cluster Snapshot Function](#).
- When snapshots are enabled for a cluster, the API will use the previous configuration to ensure consistency. If snapshots are enabled for the first, the system automatically creates an OBS bucket and file path and preferentially uses an existing agency. If there are no agencies, the system automatically creates an agency named css-obs-agency.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/index_snapshot/auto_setting

Table 5-204 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster where snapshots you want to back up. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Automatically configure basic settings of a cluster snapshot.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/ea244205-d641-45d9-9dcb-ab2236bcd07e/index_snapshot/auto_setting
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
406	The server could not fulfill the request according to the content characteristics of the request.

Error Codes

See [Error Codes](#).

5.10 VPC Endpoint

5.10.1 Enabling the VPC Endpoint Service

Function

This API is used to enable the VPCEP service for a cluster.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/vpcepservice/open

Table 5-205 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster whose VPC endpoint you want to enable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-206 Request body parameters

Parameter	Mandatory	Type	Description
endpoint_with_dns_name	No	Boolean	Enable the VPC endpoint.
profession_vpc	No	Boolean	Create a professional VPC endpoint. • true: This option will be enabled. • false: This option will be disabled.
dualstack_enable	No	Boolean	Whether to enable the IPv4/IPv6 dual-stack network. The IPv4/IPv6 dual-stack network can be enabled only when a professional VPC endpoint is created and the VPC of the cluster supports IPv6. • true: This option will be enabled. • false: This option will be disabled.

Response Parameters

Status code: 200

Table 5-207 Response body parameters

Parameter	Type	Description
action	String	Parameter description: Operation. The fixed value is createVpcepservice , indicating that the VPC endpoint is enabled.

Example Requests

Enable the VPC endpoint service.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/vpcepservice/open
{
  "endpoint_with_dns_name" : true
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "action" : "createVpcepservice"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.10.2 Disabling the VPC Endpoint Service

Function

This API is used to disable the VPCEP service for a cluster.

 NOTE

After the VPCEP service is disabled, users can no longer access the cluster via the VPCEP IP address or a private domain name. If you disable the VPCEP service and then re-enable it, the VPCEP IP address or private domain name may change. Exercise caution.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/vpcepservice/close

Table 5-208 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose VPC endpoint you want to disable. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-209 Response body parameters

Parameter	Type	Description
action	String	Parameter description: Action. The fixed value is deleteVpcepservice , indicating that the VPC endpoint is disabled.

Example Requests

Disable the VPC endpoint service.

```
PUT https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/vpcepservice/close
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "action": "deleteVpcepservice"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.10.3 Obtaining an Endpoint Connection

Function

The VPC endpoint service enables secure and reliable access across VPCs through a dedicated gateway, without exposing the network information of servers. The VPC endpoint address can be an IPv4 address (automatically allocated when the VPC endpoint service is enabled), an IPv6 address (allocated when a professional VPC endpoint is created and dual-stack networking is enabled), or an internal domain name (allocated when DNS is enabled). This API is used to obtain the VPCEP connection of a cluster.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/vpcepservice/connections

Table 5-210 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-211 Query Parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Definition: Start value for querying VPC endpoint connections. Constraints: N/A Value range: 0-1000 Default value: 0
limit	No	Integer	Definition: The number of VPC endpoint connections to be queried. Constraints: N/A Value range: 1-1000 Default value: The default value is 10, indicating that 10 tasks are queried at a time.

Request Parameters

None

Response Parameters

Status code: 200

Table 5-212 Response body parameters

Parameter	Type	Description
connections	Array of connections objects	Definition: Cluster VPC endpoint connection information. Value range: N/A
vpcServiceName	String	Definition: Name of the VPC endpoint service. Value range: N/A
permissions	Array of permissions objects	Definition: Permissions of the cluster VPC endpoint whitelist. Value range: N/A
vpcepUpdateSwitch	Boolean	Definition: Whether to enable VPC endpoint. Value range: • true: The VPC endpoint is enabled. • false: The VPC endpoint is disabled.
total_count	Integer	Definition: Number of cluster VPC endpoints. Value range: N/A

Table 5-213 connections

Parameter	Type	Description
id	String	Definition: Cluster VPC endpoint ID. Value range: N/A

Parameter	Type	Description
status	String	Definition: Cluster VPC endpoint status. Value range: • accepted: The VPC endpoint connection is enabled. • rejected: The VPC endpoint connection is disabled.
maxSession	String	Definition: Maximum number of connections to a VPC endpoint. Value range: N/A
specificationName	String	Definition: Cluster VPC endpoint name. Value range: N/A
created_at	String	Definition: Creation time. Value range The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
update_at	String	Definition: Update time. The default value is null. Value range The format is CCYY-MM-DDThh:mm:ss (ISO 8601).
domain_id	String	Definition: Account ID of the owner. Value range: N/A
vpceIp	String	Definition: IPv4 address of a cluster VPC endpoint. Value range: N/A
vpceIpv6Address	String	Definition: IPv6 address of a cluster VPC endpoint. Value range: N/A

Parameter	Type	Description
vpcepDnsName	String	Definition: Private domain name of a cluster VPC endpoint. Value range: N/A

Table 5-214 permissions

Parameter	Type	Description
id	String	Definition: Unique ID of the permission. Value range: N/A
permission	String	Definition: Permission details for the VPCEP connection whitelist. Value range: N/A
permission_type	String	Definition: VPC endpoint permission type. Value range: domainId: user account ID.
created_at	String	Definition: Specifies when the VPC endpoint was created. Value range: The format is CCYY-MM-DDThh:mm:ss (ISO 8601).

Example Requests

Obtain the VPCEP connection of a cluster.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/vpcepservice/connections
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "connections" : [ {
    "id" : "54b7f739-31a1-40d0-862b-ac85b83ab2da",
    "status" : "accepted",
    "maxSession" : "3000",
    "specificationName" : "default",
    "vpceIp" : "192.168.0.122",
    "created_at" : "2024-06-11T09:36:24Z",
    "domain_id" : "db9b76a0d9ae431f8e85e89da2ca867c"
  }, {
    "id" : "e88ad0bc-c2c7-419c-bd9b-a961111f0a42",
    "status" : "accepted",
    "maxSession" : "3000",
    "specificationName" : "default",
    "vpceIp" : "192.168.0.133",
    "created_at" : "2023-10-12T07:33:16Z",
    "domain_id" : "db9b76a0d9ae431f8e85e89da2ca867c"
  } ],
  "vpceUpdateSwitch" : false,
  "total_count" : 2,
  "vpceServiceName" : "cn-north-4.css-op-no-delete.cf072729-b076-49db-83d3-020dc9f589bd",
  "permissions" : [ {
    "id" : "432a6429-f377-4168-8b24-feb5885af08c",
    "permission" : "iam:domain::db9b76a0d9ae431f8e85e89da2ca867c",
    "permission_type" : "domainId",
    "created_at" : "2023-10-12T07:33:11Z"
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.10.4 Updating an Endpoint Connection

Function

This API is used to update the VPCEP connection of a cluster.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/vpcepservice/connections

Table 5-215 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose VPC endpoint you want to update. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-216 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Expected behavior. • receive: Accept the VPC endpoint. • reject: Reject the VPC endpoint.
endpoint_id_list	Yes	Array of strings	Lists VPC endpoint IDs.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Update an endpoint connection.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/vpcep-service/connections
```

```
{
  "action": "receive",
  "endpoint_id_list": [ "f132bb14-e1d5-4f25-9f7c-a29e4c8effd4" ]
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.

Status Code	Description
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.10.5 Modifying the VPCEP Service Whitelist of a Cluster

Function

This API is used to modify the endpoint service whitelist.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/vpcepservice/permissions

Table 5-217 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster to be modified. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-218 Request body parameters

Parameter	Mandatory	Type	Description
vpc_permissions	Yes	Array of strings	Whitelist (user account ID).

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Modify the endpoint service whitelist.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/vpcepservice/permissions
{
  "vpc_permissions" : [ "dffef70c1db243cbb373e45791165281" ]
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.11 Parameter Configuration

5.11.1 Modifying Cluster Parameters

Function

The core configuration information of an Elasticsearch cluster is stored in the **elasticsearch.yml** file, while that of an OpenSearch cluster is stored in the **opensearch.yml** file. You can modify a cluster's configuration file for purposes like accelerating queries, modifying cross-domain access configuration, adjusting the internal cache size, and managing the task queue size.

Before calling this API, call the [Obtaining the Parameter Configuration List](#) API to verify the cluster's parameter settings, and modify them as required.

Constraints

Set the values of any custom parameters to those supported by Elasticsearch. Otherwise, the cluster may fail to restart. Exercise caution when performing this operation.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/ymls/update

Table 5-219 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster whose parameter settings you want to be modify. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-220 Request body parameters

Parameter	Mandatory	Type	Description
edit	Yes	UpdateYmlsReqEdit object	Definition: Configuration file information. Constraints: N/A

Parameter	Mandatory	Type	Description
inst_type	No	String	Node type. KooSearch clusters contain different types of nodes. kos: search control node kos-doc: document parsing node

Table 5-221 UpdateYmlsReqEdit

Parameter	Mandatory	Type	Description
modify	No	UpdateYmlsReqEditModify object	Definition: Modify parameter settings. Constraints: • modify, delete, and reset: at least one of them will take effect, and each operation can only be used once each time. • Cannot perform multiple operations on the same parameter at the same time.
delete	No	UpdateYmlsReqEditModify object	Definition: Delete custom parameter settings. Constraints: • Irrelevant to the input parameter value. • modify, delete, and reset: at least one of them will take effect, and each operation can be used only once. • Cannot perform multiple operations on the same parameter at the same time.

Parameter	Mandatory	Type	Description
reset	No	UpdateYmlsReqEditModify object	Definition: Reset parameter settings. Constraints: - Irrelevant to the input parameter value. - Custom parameters cannot be reset. - modify, delete, and reset: at least one of them will take effect, and each operation can be used only once. - Cannot perform multiple operations on the same parameter at the same time.

Table 5-222 UpdateYmlsReqEditModify

Parameter	Mandatory	Type	Description
elasticsearch.yml	Yes	Object	Definition: Parameter configuration list. The value is the JSON data to be modified. Constraints: This parameter is also used for OpenSearch clusters.

Response Parameters

Status code: 200

Table 5-223 Response body parameters

Parameter	Type	Description
acknowledged	Boolean	Definition: Parameter settings modified successfully or not. Value range: true: The modification is successful. false: The modification fails.

Parameter	Type	Description
externalMessage	String	Definition: Error description. Value range: If acknowledged was set to true , null is returned for this field.
httpErrorResponse	String	Definition: HTTP error information. The default value is null. Value range: N/A

Example Requests

Modify parameter settings.

POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ymls/update

```
{
  "edit" : {
    "modify" : {
      "elasticsearch.yml" : {
        "thread_pool.force_merge.size" : 1
      }
    }
  }
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "acknowledged" : true
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.

Status Code	Description
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.11.2 Obtaining the Task List of Parameter Configurations

Function

This API is used to obtain the parameter configuration task list of a cluster. It can obtain the parameter configuration list ID, cluster ID, and parameter modification records. After a cluster is created, you can call this API to check the cluster parameter settings.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/ymls/joblists

Table 5-224 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster you want to query. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-225 Query Parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Definition: The start value of the query. The default value is 1, indicating that the query starts from the first task. Constraints: N/A Value range: 1-1000 Default value: 1
limit	No	Integer	Parameter description: Number of tasks to be queried. The default value is 10 , indicating that 10 tasks are queried at a time. Constraints: N/A Options: 1-1000 Default value: 10

Request Parameters

None

Response Parameters

Status code: 200

Table 5-226 Response body parameters

Parameter	Type	Description
configList	Array of configListRsp objects	Definition: List of historical configuration changes. Value range: N/A
totalSize	Integer	Definition: Number of configuration tasks. Value range: N/A

Table 5-227 configListRsp

Parameter	Type	Description
id	String	Definition: Configuration list ID. Value range: N/A
clusterId	String	Definition: Cluster ID. Value range: N/A
createAt	String	Definition: Configuration list creation time Value range: Format: Unix timestamp.

Parameter	Type	Description
status	String	Definition: Cluster task execution status. Value range: • success: The task is successful. • failed: The task failed. • running: The task is being executed.
finishedAt	String	Definition: End time. If the creation has not been completed, the end time is null. Value range: Format: Unix timestamp.
modifyDeleteReset	String	Definition: History of parameter setting modifications. Value range: N/A
failedMsg	String	Definition: Returned error message. If the status is success, the value of this parameter is null. Value range: N/A

Example Requests

Obtaining the Parameter Configuration Task List of a Cluster

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/ymls/joblists
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "configList" : [ {
    "id" : "7ed7258a-60a8-46fe-8814-52819d491b80",
    "clusterId" : "4213d908-f5dc-4633-8401-cfd7175fca0c",
    "createAt" : "1755745344000",
    "status" : "success",
    "finishedAt" : "1755745355000",
    "modifyDeleteReset" : "{\"modify\":{\"elasticsearch.yml\":{\"thread_pool.force_merge.size\":\"1\"}}}",
    "failedMsg" : ""
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.11.3 Obtaining the Parameter Configuration List

Function

This API is used to obtain the parameter settings list of a cluster in the form of an AML file. The core configuration information of an Elasticsearch cluster is stored in the **elasticsearch.yml** file, while that of an OpenSearch cluster is stored in the **opensearch.yml** file.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/ymls/template

Table 5-228 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-229 Response body parameters

Parameter	Type	Description
configurations	Object	Definition: Cluster parameter configuration list. Value range: The key value in the object is subject to the actual situation. The value has the following attributes: • id: parameter ID. • key: parameter name. • value: parameter value. • defaultValue: parameter default value. • regex: parameter constraint. • desc: parameter description in Chinese. • type: parameter type description. • moduleDesc: parameter function description in Chinese. • modifyEnable: whether the parameter can be modified. The value can be true or false. • enableValue: parameter value that can be changed. • fileName: name of the file that contains the parameters. The default value is elasticsearch.yml, which is also used by OpenSearch clusters. • version: version information. • descENG: parameter description in English. • moduleDescENG: parameter function description in English. • instType: node type.

Example Requests

Obtain the parameter settings of a cluster in a YAML file.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/5c77b71c-5b35-4f50-8984-76387e42451a/ymls/template
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "configurations" : {
    "http.cors.allow-credentials" : {
      "id" : "b462d13c-294b-4e0f-91d3-58be2ad02b99",
      "key" : "http.cors.allow-credentials",
      "value" : "false",
      "defaultValue" : "false",
      "regex" : "^(true|false)$",
      "desc" : "Indicates whether to return Access-Control-Allow-Credentials in the header during cross-domain access. The value is of the Boolean type and can be true or false",
      "type" : "Boolean",
      "moduleDesc" : "Cross-domain access",
      "modifyEnable" : "true",
      "enableValue" : "true,false",
      "fileName" : "elasticsearch.yml",
      "version" : null,
      "descENG" : "Whether to return the Access-Control-Allow-Credentials of the header during cross-domain access. The value is a Boolean value and the options are true and false.",
      "moduleDescENG" : "Cross-domain Access"
    }
  }
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request before retry.
409	The request could not be completed due to a conflict with the current state of the resource. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
412	The server did not meet one of the preconditions contained in the request.

Error Codes

See [Error Codes](#).

5.12 Intelligent O&M

5.12.1 Obtaining the Intelligent O&M Task List and Details

Function

CSS provides intelligent O&M to check potential risks in clusters. After a scan task is completed, check the risks identified by the task, confirm and handle these risks

in a timely manner based on risk handling suggestions. This API is used to obtain the intelligent O&M task list and details.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/ai-ops

Table 5-230 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-231 Query Parameters

Parameter	Mandatory	Type	Description
limit	No	Integer	Definition: Pagination parameter, indicating the maximum number of records on a page. The default value is 10, indicating that 10 tasks are queried at a time. Constraints: N/A Value range: 1-1000 Default value: 10
offset	No	Integer	Definition: Offset. All VPC endpoint services after this offset will be queried. The offset must be an integer greater than 0 but less than the number of VPC endpoint services. Constraints: N/A Value range: 0-1000 Default value: 0
report	No	String	Definition: Obtain the latest report or a historical report. Constraints: N/A Value range: • current: Obtain only the latest check report. • history: Obtain only historical reports. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-232 Response body parameters

Parameter	Type	Description
aiops_list	Array of aiops_list objects	Definition: Cluster risk detection task details. Value range: N/A
check_type	String	Definition: Check scope Value range: - full_detection: full check - unavailability_detection: cluster unavailability check - partial_detection: partial check
trigger_type	String	Definition: Trigger type Value range: - manual - auto

Table 5-233 aiops_list

Parameter	Type	Description
id	String	Definition: ID of a cluster risk detection task. Value range: N/A
name	String	Definition: Name of a cluster risk detection task. Value range: N/A

Parameter	Type	Description
status	Integer	Definition: Task execution status. Value range: • 150: The function is disabled. • 200: The function is enabled. • 300: A message has been sent.
summary	summary object	Definition: Cluster risk overview. Value range: N/A
create_time	String	Definition: Time when a cluster risk detection task was created. Value range: Format: Unix timestamp.
task_risks	Array of AiOpsRiskObject objects	Definition: Cluster risk details. Value range: N/A

Table 5-234 summary

Parameter	Type	Description
highRisk	Integer	Definition: Number of high-risk items. Value range: 0-100
mediumRisk	Integer	Definition: Number of medium-risk items. Value range: 0-100
normal	String	Definition: Number of medium-risk items. Value range: 0-100

Parameter	Type	Description
suggestion	Integer	Definition: Number of items that are normal. Value range: 0-100

Table 5-235 AiOpsRiskObject

Parameter	Type	Description
id	String	Definition: Check item ID Value range: N/A Default value: N/A
name	String	Definition: Check item name Value range: N/A Default value: N/A
desc	String	Definition: Check item description Value range: N/A Default value: N/A
level	String	Definition: Risk level of the check result Value range: N/A Default value: N/A

Parameter	Type	Description
result	String	Definition: Check result Value range: N/A Default value: N/A
suggestion	String	Definition: Handling suggestion Value range: N/A Default value: N/A

Example Requests

This API is used to obtain the intelligent O&M task list and details.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ai-ops
```

Example Responses

Status code: 200

Request succeeded.

```
[ {
  "id" : "0f896cea-c56a-4235-8cf2-69018229c903",
  "name" : "aiops-1763365380724",
  "status" : 200,
  "summary" : {
    "highRisk" : 0,
    "mediumRisk" : 0,
    "normal" : 1,
    "suggestion" : 0
  },
  "trigger_type" : "manual",
  "check_type" : "full_detection",
  "create_time" : 1763365380723,
  "task_risks" : [ {
    "id" : "nodeDiskCheck",
    "name" : "name",
    "desc" : "desc",
    "level" : "normal",
    "result" : "result",
    "suggestion" : "suggestion"
  } ]
} ]
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.12.2 Creating a Cluster Detection Task

Function

This API is used to create a cluster detection task.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/ai-ops

Table 5-236 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the target cluster. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-237 Request header parameters

Parameter	Mandatory	Type	Description
X-Language	No	String	Language

Table 5-238 Request body parameters

Parameter	Mandatory	Type	Description
check_type	No	String	Definition: Check type Value range: - full_detection (all check items) - unavailability_detection (cluster unavailability check items) - partial_detection (partial check items; select by setting check_items.) Default value: N/A
check_items	No	Array of strings	Definition: Select a subset of all check items, and enter the check item list. Value range: N/A Default value: N/A
name	No	String	Definition: A custom name of the check report. Value range: N/A Default value: N/A
description	No	String	Definition: A custom description of the check report. Value range: 0 to 255 characters Default value: N/A

Parameter	Mandatory	Type	Description
alarm	No	AlarmInfo object	Definition: Check report sending. Currently, it is not supported. Value range: N/A Default value: N/A

Table 5-239 AlarmInfo

Parameter	Mandatory	Type	Description
level	No	String	Report level. This parameter is not supported currently.
smn_topic	No	String	SMN topic. This parameter is not supported currently.

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Create a cluster detection task.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ai-ops
{
  "check_type": "full_detection",
  "check_items": [ ]
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.

Status Code	Description
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.12.3 Deleting a Detection Task

Function

This API is used to delete a detection task.

Calling Method

For details, see [Calling APIs](#).

URI

DELETE /v1.0/{project_id}/clusters/{cluster_id}/ai-ops/{aiops_id}

Table 5-240 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster to be deleted. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
aiops_id	Yes	String	ID of a detection task

Request Parameters

None

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Delete a detection task.

```
DELETE https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ai-ops/e19c9298-465e-42ad-a0ae-b6b552222925
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.

Status Code	Description
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.12.4 Obtaining SMN Topics Available for Intelligent O&M Alarms

Function

This API is used to obtain SMN topics available for intelligent O&M alarms.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/domains/{domain_id}/ai-ops/smn-topics

Table 5-241 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
domain_id	Yes	String	Domain account ID

Request Parameters

None

Response Parameters

Status code: 200

Table 5-242 Response body parameters

Parameter	Type	Description
topicsName	Array of strings	Subject name

Example Requests

Obtain the intelligent O&M task list and details.

```
GET /v1.0/6204a5bd270343b5885144cf9c8c158d/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ai-ops
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "topicsName": [ "aiops-test" ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.12.5 Viewing Intelligent O&M Check Items

Function

CSS provides intelligent O&M to check potential risks in clusters. This API is used to obtain the check items for intelligent O&M.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/ai-ops/detector

Table 5-243 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-244 Response body parameters

Parameter	Type	Description
full_detection	Array of AiOpsDetector objects	All check items
unavailability_detection	Array of AiOpsDetector objects	Cluster unavailability check items

Table 5-245 AiOpsDetector

Parameter	Type	Description
id	String	Definition: Check item ID. Value range: N/A
name	String	Definition: Check item name. Value range: N/A
desc	String	Definition: Check item description. Value range: N/A

Example Requests

Viewing Intelligent O&M Check Items

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ai-ops/detector
```

Example Responses

Status code: 200

Intelligent O&M check items

```
{  
  "full_detection": [{  
    "id": "nodeDiskCheck",
```

```
"name" : "name",
"desc" : "desc"
}],
"unavailability_detection" : [ {
  "id" : "nodeDiskCheck",
  "name" : "name",
  "desc" : "desc"
} ]
}
```

Status Codes

Status Code	Description
200	Intelligent O&M check items
400	Invalid request. Modify the request directly and do not attempt to retry it.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request failed to be processed because of the update of the conflict request.
412	The server does not meet one of the preconditions that the client puts on the request.

Error Codes

See [Error Codes](#).

5.12.6 Viewing Scheduled Check Settings for Intelligent O&M

Function

CSS provides intelligent O&M to check potential risks in clusters on a daily basis. This API is used to obtain the scheduled check settings for intelligent O&M.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/ai-ops/setting

Table 5-246 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-247 Response body parameters

Parameter	Type	Description
check_type	String	Definition: Check type Value range: - full_detection (all check items) - unavailability_detection (cluster unavailability check items) - partial_detection (partial check items; select by setting check_items.) Default value: N/A
check_items	Array of strings	Definition: Select a subset of all check items, and enter the check item list. Value range: N/A Default value: N/A
period	String	Definition: Intelligent O&M automatic check time, in the form of time followed by time zone, for example, 00:00 GMT+08:00. Value range: N/A Default value: N/A

Example Requests

Viewing Scheduled Check Settings for Intelligent O&M

GET https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ai-ops/setting

Example Responses

Status code: 200

```
{
  "check_type": "full_detection",
  "check_items": [ "nodeDiskCheck", "dataNodeDiskBalanceCheck" ],
  "period": "23:00 GMT+08:00"
}
```

Status Codes

Status Code	Description
200	
400	Invalid request. Modify the request directly and do not attempt to retry it.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request failed to be processed because of the update of the conflict request.
412	The server does not meet one of the preconditions that the client puts on the request.

Error Codes

See [Error Codes](#).

5.12.7 Updating Scheduled Check Settings for Intelligent O&M

Function

CSS provides intelligent O&M to check potential risks in clusters on a daily basis. This API is used to configure scheduled checks for intelligent O&M.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/ai-ops/setting

Table 5-248 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-249 Request body parameters

Parameter	Mandatory	Type	Description
check_type	Yes	String	Definition: Check type Value range: - full_detection (all check items) - unavailability_detection (cluster unavailability check items) - partial_detection (partial check items; select by setting check_items.) Default value: N/A
check_items	No	Array of strings	Definition: Select a subset of all check items, and enter the check item list. Value range: N/A Default value: N/A
period	Yes	String	Definition: Intelligent O&M automatic check time, in the form of time followed by time zone, for example, 00:00 GMT +08:00. Value range: N/A Default value: N/A

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

None

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request directly and do not attempt to retry it.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request failed to be processed because of the update of the conflict request.
412	The server does not meet one of the preconditions that the client puts on the request.

Error Codes

See [Error Codes](#).

5.12.8 Disabling Scheduled Checks for Intelligent O&M

Function

CSS provides intelligent O&M to check potential risks in clusters on a daily basis. This API is used to disable scheduled checks for intelligent O&M.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/ai-ops/close

Table 5-250 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200
Request succeeded.
None

Example Requests

Disabling Scheduled Checks for Intelligent O&M
PUT https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/ai-ops/close

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request directly and do not attempt to retry it.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request failed to be processed because of the update of the conflict request.
412	The server does not meet one of the preconditions that the client puts on the request.

Error Codes

See [Error Codes](#).

5.13 Cluster management

5.13.1 Updating Cluster Routes

Function

This interface is used to update a cluster route.

NOTE

Modifying the cluster route may cause service network exceptions and affect the network connection between the client and the cluster. Exercise caution when performing this operation.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/route

Table 5-251 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the target cluster. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-252 Request body parameters

Parameter	Mandatory	Type	Description
configtype	Yes	String	Definition: Operation type. Constraints: N/A Value range: • add_ip: Add a cluster route. • del_ip: Delete a cluster route. Default value: N/A

Parameter	Mandatory	Type	Description
configkey	Yes	String	Definition: Route IP address, that is, the IP address of the server where the public network data source is located. Constraints: The value of this parameter cannot start with 0 . Value range: N/A Default value: N/A
configvalue	Yes	String	Definition: Subnet mask of a route. If the IP address contains 16 bits, set the subnet mask to 255.255.0.0 . If the IP address contains 24 bits, set the subnet mask to 255.255.255.0 . Constraints: The value of this parameter cannot start with 0 . Value range: N/A Default value: N/A

Response Parameters

Status code: 200

Request succeeded.

None

Example Requests

Update cluster routes.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/route
{
  "configtype" : "add_ip",
  "configkey" : "10.5.2.1",
  "configvalue" : "255.255.255.255"
}
```

Example Responses

None

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.13.2 Obtaining a Cluster Route

Function

By default, a CSS Elasticsearch cluster cannot access a target service—such as a client program or a Lightweight Directory Access Protocol (LDAP) service—that is located in a different VPC. To enable this access, you need to configure cluster routes. This API is used to obtain cluster route information, including the route IP address and the total number of routes.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/route

Table 5-253 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the target cluster. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-254 Query Parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Definition: The start value of the query. The default value is 1, indicating that the query starts from the first route. Constraints: N/A Value range: 1-1000 Default value: 1

Parameter	Mandatory	Type	Description
limit	No	Integer	Parameter description: Number of routes to be queried. The default value is 10 , indicating that 10 routes are queried at a time. Constraints: N/A Options: 1-1000 Default value: 10

Request Parameters

None

Response Parameters

Status code: 200

Table 5-255 Response body parameters

Parameter	Type	Description
routeResps	Array of RouteRespsResource objects	Definition: Route IP address. Value range: N/A
totalSize	Integer	Definition: Total number of routes. Value range: N/A

Table 5-256 RouteRespsResource

Parameter	Type	Description
ipAddress	String	Definition: IP address. Value range: N/A

Parameter	Type	Description
ipNetMask	String	Definition: Subnet mask. Value range: N/A
updateAt	String	Definition: Route IP address update time. Value range: N/A

Example Requests

Obtain cluster routes.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/route
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "routeResps" : [ {
    "ipAddress" : "1.2.3.65",
    "ipNetMask" : "255.255.255.255",
    "updateAt" : "2023-07-17T08:09:20"
  }, {
    "ipAddress" : "1.1.1.1",
    "ipNetMask" : "255.255.255.255",
    "updateAt" : "2023-07-17T08:08:50"
  } ],
  "totalSize" : 2
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.

Status Code	Description
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.14 Load Balancing

5.14.1 Querying Dedicated Load Balancers Supported by a Cluster

Function

This API is used to query the dedicated load balancers supported by an Elasticsearch or OpenSearch cluster.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/loadbalancers

Table 5-257 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-258 Response body parameters

Parameter	Type	Description
loadbalancers	Array of LoadbalancersResource objects	Definition: List of dedicated load balancers.

Table 5-259 LoadbalancersResource

Parameter	Type	Description
id	String	Definition: Load balancer ID. Value range: N/A
name	String	Definition: Load balancer name. Value range: N/A

Parameter	Type	Description
l7_flavor_id	String	Definition: Layer-7 protocol ID. Value range: N/A
ip_target_enable	Boolean	Definition: Whether to enable the cross-VPC backend. Value range: • true: Enable the cross-VPC backend. • false: Disable the cross-VPC backend.

Example Requests

Query dedicated load balancers supported by a cluster.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/loadbalancers
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "loadbalancers" : [ {
    "id" : "5d45faad-6cb3-479b-96b8-3e2de0cc6268",
    "name" : "elb-css",
    "l7_flavor_id" : "9c8c2425-e061-4bf8-ac65-cd1db92b18e1",
    "ip_target_enable" : true
  }, {
    "id" : "5d45faad-6cb3-479b-96b8-3e2de0cc6269",
    "name" : "elb-b832",
    "l7_flavor_id" : "9c8c2425-e061-4bf8-ac65-cd1db92b18e1",
    "ip_target_enable" : true
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.

Status Code	Description
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.14.2 Enabling or Disabling Load Balancers for a Cluster

Function

This API is used to enable or disable load balancers for a cluster.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/loadbalancers/es-switch

Table 5-260 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster to be modified. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-261 Request body parameters

Parameter	Mandatory	Type	Description
enable	Yes	Boolean	Definition: Enable or disable the ES load balancer. Constraints: N/A Value range: • true: Yes • false: This option will be disabled. Default value: N/A
agency	No	String	Definition: Name of the agency used to configure load balancing. This parameter is mandatory when enable is set to true . Constraints: N/A Value range: N/A Default value: N/A

Parameter	Mandatory	Type	Description
elb_id	No	String	Definition: Load balancer ID. This parameter is mandatory when enable is set to true . Constraints: N/A Value range: N/A Default value: N/A
type	No	String	Definition: Load balancer type. This parameter is not required for an observability cluster. Constraints: N/A Value range: searchTool: Enable or disable an Elasticsearch/OpenSearch load balancer. Default value: searchTool

Response Parameters

Status code: 200

Table 5-262 Response body parameters

Parameter	Type	Description
elb_id	String	Load balancer ID.

Example Requests

Enable the load balancer.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/58ee0f27-70b3-47e0-ac72-9e3df6cd15cd/loadbalancers/es-switch
```

```
{
  "enable" : true,
  "elb_id" : "5d45faad-6cb3-479b-96b8-3e2de0cc6268",
  "agency" : "css_elb_agency"
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "elb_id" : "5d45faad-6cb3-479b-96b8-3e2de0cc6268"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.14.3 Obtain information about the load balancers of a cluster.

Function

CSS integrates shared load balancers and allows you to bind public network access and enable the VPC Endpoint service. Dedicated load balancers provide more functions and higher performance than shared load balancers.

This API is used to obtain the ELB load balancer information of a cluster, including the server certificate name, server certificate ID, ELB switch information, load balancers, listeners, and health check results.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/es-listeners

Table 5-263 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Table 5-264 Query Parameters

Parameter	Mandatory	Type	Description
type	No	String	Definition: Search type. Constraints: N/A Value range: searchTool: Query details about the load balancers of an Elasticsearch/OpenSearch cluster. Default value: searchTool

Request Parameters

None

Response Parameters

Status code: 200

Table 5-265 Response body parameters

Parameter	Type	Description
serverCertName	String	Definition: Server certificate name. Value range: N/A
serverCertId	String	Definition: Server certificate ID. Value range: N/A
cacertName	String	Definition: The name of the CA certificate. Value range: N/A
cacertId	String	Definition: ID of the CA certificate. Value range: N/A
elb_enable	Boolean	Definition: Whether ELB is enabled. Value range: • true: ELB is enabled. • false: ELB is disabled.
authentication_type	String	Definition: Authentication mode. Value range: • single: one-way authentication. • double: two-way authentication.

Parameter	Type	Description
loadBalancer	LoadBalancer object	Definition: Load balancer information. Value range: N/A
listener	Elbv3Listener object	Definition: Listener object information. Value range: N/A
healthmonitors	Array of Member objects	Definition: Health check result set. Value range: N/A

Table 5-266 LoadBalancer

Parameter	Type	Description
id	String	Definition: Load balancer ID. Value range: N/A
name	String	Definition: Load balancer name. Value range: N/A
guaranteed	String	Definition: Whether a dedicated load balancer is used. Value range: • false: The load balancer is a shared load balancer. • true: The load balancer is a dedicated load balancer.

Parameter	Type	Description
billing_info	String	Definition: Resource billing information. Value range: - If the value is left blank, the resource is billed on a pay-per-use basis. - If the value is not left blank, the resource is billed on a yearly/ monthly basis.
description	String	Definition: Description. Value range: N/A
vpc_id	String	Definition: ID of the VPC to which the ELB instance belongs. Value range: N/A
provisioning_status	String	Definition: Provisioning status of the load balancer Value range: ACTIVE: in use PENDING_DELETE: deleting
listeners	Array of IdListWrapper objects	Definition: Associated listener list. Value range: N/A
vip_address	String	Definition: IPv4 virtual IP address bound to the load balancer. Value range: N/A
vip_port_id	String	Definition: Port ID bound to the private IPv4 IP address of the load balancer. Value range: N/A

Parameter	Type	Description
ipv6_vip_address	String	Definition: IPv6 address of the load balancer. Value range: N/A
publicips	Array of PublicIpInfo objects	Definition: The EIP bound to the load balancer Value range: N/A

Table 5-267 IdListWrapper

Parameter	Type	Description
id	String	Definition: Listener ID. Value range: N/A

Table 5-268 PublicIpInfo

Parameter	Type	Description
publicip_id	String	Definition: EIP configuration ID. Value range: N/A
publicip_address	String	Definition: Specifies the EIP. Value range: N/A
ip_version	Integer	Definition: IP address version. Value range: • 4: IPv4. • 6: IPv6.

Table 5-269 Elbv3Listener

Parameter	Type	Description
id	String	Definition: Listener ID. Value range: N/A
name	String	Definition: Listener name. Value range: N/A
protocol	String	Definition: Protocol used by the listener. Value range: N/A
protocol_port	Integer	Definition: Port used by the listener. Value range: N/A
ipgroup	ListenerIpGroup object	Definition: ipgroup information in the listener object. Value range: N/A

Table 5-270 ListenerIpGroup

Parameter	Type	Description
ipgroup_id	String	Definition: ID of the access control group associated with the listener. This parameter is mandatory during creation and is optional during update. Value range: N/A

Parameter	Type	Description
enable_ipgroup	Boolean	Definition: Status of an access control group. Value range: • True: Access control is enabled. • False: Access control is disabled.

Table 5-271 Member

Parameter	Type	Description
name	String	Definition: Specifies the backend server name. Value range: N/A
address	String	Definition: Specifies the IP address bound to the backend server. Value range: N/A
protocol_port	Integer	Definition: Specifies the port used by the backend server. Value range: N/A
operating_status	String	Definition: Specifies the operating status of the backend server. Value range: • ONLINE: The backend server is running normally. • NO_MONITOR: No health check is configured for the backend server group to which the backend server belongs. • OFFLINE: The cloud server used as the backend server is stopped or does not exist.

Parameter	Type	Description
instance_id	String	Definition: ID of the instance used as the backend server. Value range: An empty value indicates that the instance associated with the member is not a real device.

Example Requests

Obtain information about the load balancers of a cluster.

```
GET https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/es-listeners
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "loadBalancer": {
    "id": "5d45faad-6cb3-479b-96b8-3e2de0cc6268",
    "name": "elb-css",
    "guaranteed": "true",
    "description": "",
    "listeners": [ {
      "id": "011d14fa-908b-4cd9-b0d6-0768ddc6cb71"
    }, {
      "id": "04b86029-c281-4490-a3bd-5ea1266658ba"
    } ],
    "publicips": [ {
      "publicip_id": "f678e23c-96a5-49e8-8ea2-bd8b47d41e78",
      "publicip_address": "100.1.1.1",
      "ip_version": 4
    } ],
    "vpc_id": "4f3deec3-efa8-4598-bf91-560aad1377a3",
    "provisioning_status": "ACTIVE",
    "vip_address": "10.0.0.1",
    "ipv6_vip_address": "2409:27ff:2003:3e:1fd::f3"
  },
  "listener": {
    "id": "41ff041d-c7b9-4142-9167-fa93d54f97da",
    "name": "css-searchServer",
    "protocol": "HTTPS",
    "protocol_port": 9265
  },
  "healthmonitors": [ {
    "instance_id": "bac86342-2222-43e6-817f-57f040a174a6",
    "name": "",
    "address": "10.0.0.87",
    "protocol_port": 9200,
    "operating_status": "ONLINE"
  }, {
    "instance_id": "d935b82c-f94b-4ae0-9997-ddc90885d8c6",
    "name": "",
    "address": "10.0.0.61",
    "protocol_port": 9200,
    "operating_status": "ONLINE"
  } ]
}
```

```
  },  
  "serverCertName" : "server1",  
  "serverCertId" : "82375af01c0d40f6a44c15962c570625",  
  "elb_enable" : true,  
  "authentication_type" : "single"  
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.14.4 Configuring Load Balancing Listeners for a Cluster

Function

This API is used to configure load balancing listeners for a cluster.

Calling Method

For details, see [Calling APIs](#).

URI

POST /v1.0/{project_id}/clusters/{cluster_id}/es-listeners

Table 5-272 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A
cluster_id	Yes	String	Definition: ID of the cluster that you want to rename. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

Table 5-273 Request body parameters

Parameter	Mandatory	Type	Description
protocol	Yes	String	Protocol type. HTTP and HTTPS are supported.
protocol_port	Yes	Integer	Port
server_cert_id	No	String	Server certificate ID. This parameter is mandatory when protocol is set to HTTPS .
ca_cert_id	No	String	ID of the CA certificate. This parameter is mandatory when protocol is set to HTTPS and bidirectional authentication is used.

Parameter	Mandatory	Type	Description
type	No	String	Type: searchTool indicates configuring listeners for Elasticsearch/OpenSearch. The default value is searchTool .

Response Parameters

Status code: 200

Table 5-274 Response body parameters

Parameter	Type	Description
elb_id	String	Load balancer ID.

Example Requests

Configure ELB listeners for a cluster.

```
POST https://{Endpoint}/v1.0/{project_id}/clusters/4f3deec3-efa8-4598-bf91-560aad1377a3/es-listeners
{
  "protocol": "HTTPS",
  "protocol_port": 9200,
  "server_cert_id": "1b74ae4afad741e6bc9497429d3d0838",
  "ca_cert_id": "2a7b050f598e4ea8a61cffbe2bdb312f"
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "elb_id": "56a83547-a55a-407a-8245-7f9b3a0606c0"
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.

Status Code	Description
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.14.5 Updating Load Balancing Listeners for a Cluster

Function

This API is used to update load balancing listeners for a cluster.

Calling Method

For details, see [Calling APIs](#).

URI

PUT /v1.0/{project_id}/clusters/{cluster_id}/es-listeners/{listener_id}

Table 5-275 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the target cluster. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A
listener_id	Yes	String	ID of the target listener.

Request Parameters

Table 5-276 Request body parameters

Parameter	Mandatory	Type	Description
listener	Yes	EsListenerRequest object	Listener object
type	No	String	Type: searchTool indicates configuring load balancers for Elasticsearch/OpenSearch. The default value is searchTool .

Table 5-277 EsListenerRequest

Parameter	Mandatory	Type	Description
default_tls_container_ref	Yes	String	ID of the server certificate used by the listener
client_ca_tls_container_ref	No	String	ID of the CA certificate used by the listener. This parameter is mandatory when bidirectional authentication is to be updated.

Response Parameters

Status code: 200

Table 5-278 Response body parameters

Parameter	Type	Description
listener	EsListenerResponse object	Listener object

Table 5-279 EsListenerResponse

Parameter	Type	Description
protocol	String	Protocol used by the listener
id	String	Listener ID
name	String	Listener name
protocol_port	String	Port used by the listener
ipgroup	EslpgroupResource object	Access control information of the listener object

Table 5-280 EslpgroupResource

Parameter	Type	Description
ipgroup_id	String	Definition: ID of the access control group associated with the listener. Value range: N/A
enable_ipgroup	Boolean	Definition: Status of an access control group. Value range: • true: Enable the access control listener (which can be deleted directly). • false: Disable the access control listener.

Parameter	Type	Description
type	String	Definition: Type of an access control group. Value range: • white: Configure a whitelist. Only IP addresses on the whitelist can have access. • black: Configure a blacklist. IP addresses on the blacklist are blocked access.

Example Requests

Update ELB listeners for a cluster.

PUT https://{Endpoint}/v1.0/{project_id}/clusters/58ee0f27-70b3-47e0-ac72-9e3df6cd15cd/es-listeners/4425eb63-78ce-4c63-b60e-492befdac0e7

```
{
  "listener" : {
    "default_tls_container_ref" : "9c82320a8dd049eba49b28e70a2c0c3e"
  }
}
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "listener" : {
    "id" : "9c82320a8dd049eba49b28e70a2c0c3e",
    "name" : "name",
    "protocol" : "tcp",
    "protocol_port" : 8080,
    "ipgroup" : {
      "ipgroup_id" : "4f3deec3-efa8-4598-bf91-560aad1377a3",
      "enable_ipgroup" : true,
      "type" : "white"
    }
  }
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.

Status Code	Description
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.14.6 Querying the Certificate List

Function

This API is used to query the certificate list.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/clusters/{cluster_id}/elb/certificates

Table 5-281 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Parameter	Mandatory	Type	Description
cluster_id	Yes	String	Definition: ID of the cluster to be queried. For details about how to obtain the cluster ID, see Obtaining the Cluster ID . Constraints: N/A Value range: Cluster ID. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-282 Response body parameters

Parameter	Type	Description
certificates	Array of CertificatesResource objects	Definition: Certificate list information. Value range: N/A

Table 5-283 CertificatesResource

Parameter	Type	Description
id	String	Definition: Certificate ID. Value range: N/A
name	String	Definition: Certificate name. Value range: N/A

Parameter	Type	Description
type	String	Definition: Type of the SL certificate. Value range: • server: server certificates • client: CA certificates

Example Requests

This API is used to query the certificate list.

GET https://{Endpoint}/v1.0/{project_id}/clusters/58ee0f27-70b3-47e0-ac72-9e3df6cd15cd/elb/certificates

Example Responses

Status code: 200

Request succeeded.

```
{
  "certificates" : [ {
    "id" : "8c415d2e2f4c4fdbbfc9c21c98d7832b",
    "name" : "server1",
    "type" : "server"
  }, {
    "id" : "8c415d2e2f4c4fdbbfc9c21c98d7832b",
    "name" : "ca1",
    "type" : "client"
  }, {
    "id" : "8c415d2e2f4c4fdbbfc9c21c98d7832b",
    "name" : "server-css",
    "type" : "server"
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	Invalid request. Modify the request instead of retrying.
409	The request cannot be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the requested update failed due to a conflict.
412	The server does not meet one of the requirements that the requester puts on the request.

Error Codes

See [Error Codes](#).

5.15 Node Specifications Query

5.15.1 Listing Disk Types

Function

Obtain the disk types supported by each AZ.

Calling Method

For details, see [Calling APIs](#).

URI

GET /v1.0/{project_id}/disktypes

Table 5-284 Path Parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Definition: Project ID. For details about how to obtain the project ID and name, see Obtaining the Project ID and Name . Constraints: N/A Value range: Project ID of the account. Default value: N/A

Request Parameters

None

Response Parameters

Status code: 200

Table 5-285 Response body parameters

Parameter	Type	Description
diskTypes	Array of DiskType objects	Definition: Disk type list. Value range: N/A

Table 5-286 DiskType

Parameter	Type	Description
availabilityZone	String	Definition: Availability zone (AZ). Value range: N/A
volumeNames	Array of strings	Definition: Storage type. Value range: • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O • ESSD: extreme SSD

Status code: 403**Table 5-287** Response body parameters

Parameter	Type	Description
diskTypes	Array of DiskType objects	Definition: Disk type list. Value range: N/A

Table 5-288 DiskType

Parameter	Type	Description
availabilityZone	String	Definition: Availability zone (AZ). Value range: N/A
volumeNames	Array of strings	Definition: Storage type. Value range: • SATA: common I/O • SAS: high I/O • SSD: ultra-high I/O • ESSD: extreme SSD

Example Requests

An example request for obtaining the disk types supported by each AZ.

```
GET https://{Endpoint}/v1.0/{project_id}/disktypes
```

Example Responses

Status code: 200

Request succeeded.

```
{
  "diskTypes": [ {
    "availabilityZone": "cn-north-1a",
    "volumeNames": [ "SATA" ]
  }, {
    "availabilityZone": "cn-north-1b",
    "volumeNames": [ "SATA", "SAS" ]
  }, {
    "availabilityZone": "cn-north-1c",
    "volumeNames": [ "SATA", "SAS", "SSD" ]
  } ]
}
```

Status Codes

Status Code	Description
200	Request succeeded.
400	The request is invalid. The client should modify the request instead of re-initiating it.

Status Code	Description
403	The request is rejected. The server has received the request and understood it, but refuses to respond to it. The client should not repeat the request without modifications.

Error Codes

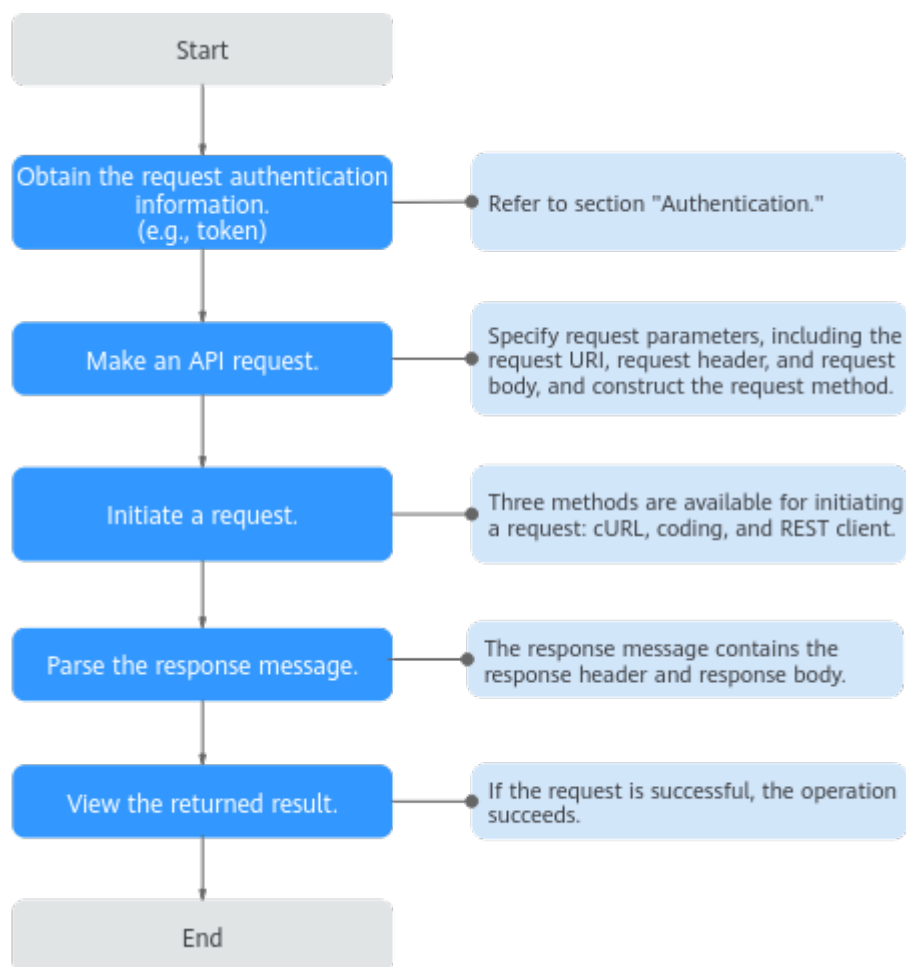
See [Error Codes](#).

6 Examples

6.1 Creating a Pay-per-Use Elasticsearch Cluster

This section describes how to create a CSS cluster by using APIs. [Figure 6-1](#) illustrates the API calling process.

Figure 6-1 API calling process



 NOTE

The token obtained on IAM is valid for only 24 hours. If you want to use one token for authentication, you can cache it to avoid frequent calling.

Involved APIs

If you use a token for authentication, you must obtain the token and add **X-Auth-Token** to the request header of the API when making an API call.

- API for obtaining tokens from IAM
- API for creating CSS clusters

Procedure

1. Obtain the token. Send **POST https://IAM endpoint/v3/auth/tokens**.
Obtain the token by following instructions in [Authentication](#).
The value of **X-Subject-Token** in the response header is the user token.
2. Add **Content-type** and **X-Auth-Token** to the request header.
 - **Content-Type**: The request body type or format. Its default value is **application/json**.
 - **X-Auth-Token**: Enter the user token obtained in 1.
3. Send a cluster creation request and specify the following parameters in the request body:

```
POST /v1.0/6204a5bd270343b5885144cf9c8c158d/clusters
```

```
{
  "cluster" : {
    "instance" : {
      "flavorRef" : "ess.spec-4u16g",
      "volume" : {
        "volume_type" : "COMMON",
        "size" : 120
      },
    },
    "nics" : {
      "vpclId" : "{VPC ID}",
      "netId" : "{NET ID}",
      "securityGroupId" : "{Security group ID}"
    },
    "availability_zone" : "{AZ CODE}"
  },
  "datastore" : {
    "version" : "{cluster-version}",
    "type" : "elasticsearch"
  },
  "name" : "cluster-name",
  "instanceNum" : 3,
  "backupStrategy" : {
    "period" : "16:00 GMT+08:00",
    "prefix" : "snapshot",
    "keepday" : 7,
    "frequency" : "DAY",
    "bucket" : "css-obs-backup",
    "basePath" : "css_repository/obs-path",
    "agency" : "css_obs_agency"
  },
  "httpsEnable" : true,
  "authorityEnable" : true,
  "adminPwd" : "{password}",
  "enterprise_project_id" : "0",
}
```

```
"tags" : [ {  
  "key" : "k1",  
  "value" : "v1"  
}, {  
  "key" : "k2",  
  "value" : "v2"  
}  
]  
}
```

Check the response message. The following is an example response:

```
{  
  "cluster": {  
    "id": "ef683016-871e-48bc-bf93-74a29d60d214",  
    "name": "ES-Test"  
  }  
}
```

If the request is successful, 200 OK is returned.

If the request fails, an error code and error information are returned. For details, see section [Status Codes](#).

7 Appendixes

7.1 Status Codes

Table 7-1 describes the status codes.

Table 7-1 Status code description

Status Code	Code	Description
100	Continue	The client continues sending the 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	Switching protocols. The target protocol must be later than the source protocol. For example, the current HTTPS protocol is switched to a later version.
200	OK	The request is processed successfully.
201	Created	The request for creating a resource has been fulfilled.
202	Accepted	The request has been accepted, but the processing has not been completed.
203	Non-Authoritative Information	The server successfully processed the request, but is returning information that may be from another source.
204	NoContent	The server has successfully processed the request, but has not returned any content. The status code is returned in response to an HTTPS OPTIONS request.

Status Code	Code	Description
205	Reset Content	The server has fulfilled the request, but the requester is required to reset the content.
206	Partial Content	The server has processed certain GET requests.
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 the user or user agent (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	Retrieve a location. The response to the request can be found under a different URI and should be retrieved 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 must be accessed through a proxy.
306	Unused	The HTTPS status code is no longer used.
400	BadRequest	Invalid request. The client should not repeat the request without modifications.
401	Unauthorized	The 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 server received and understood the request but refused to fulfill it, because the request is set to deny access. The client should not repeat the request without modifications.

Status Code	Code	Description
404	NotFound	The requested resource cannot be found. The client should not repeat the request without modifications.
405	MethodNotAllowed	The method specified in the request 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 request timed out. The client may repeat the request without modifications at any later time.
409	Conflict	The request could not be processed due to a conflict. This status code indicates that the resource that the client attempts to create already exists, or the request fails to be processed because of the update of the conflict request.
410	Gone	The requested resource is no longer available. The status code indicates that the requested resource has been deleted.
411	Length Required	The server refuses to process the request without a defined Content-Length.
412	Precondition Failed	The server does not meet one of the preconditions that the requester puts on the request.
413	Request Entity Too Large	The request is larger than that a server is able to process. The server may close the connection to prevent the client from continuing the request. If the server cannot process the request temporarily, the response will contain a Retry-After header field.
414	Request-URI Too Large	The URI provided was too long for the server to process.
415	Unsupported Media Type	The server is unable to process the media format in the request.

Status Code	Code	Description
416	Requested range not satisfiable	The requested range is invalid.
417	Expectation Failed	The server fails 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 more requests than its rate limit is allowed within a given amount of time, or the server has received more requests than it is able to process within a given amount of time. In this case, it is advisable for the client to re-initiate requests after the time specified in the Retry-After header of the response expires.
500	InternalServerError	The server is able to receive the request but it could not understand the request.
501	Not Implemented	The server does not support the requested function.
502	Bad Gateway	The server is acting as a gateway or proxy and receives an invalid request from a 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 when the Timeout parameter is specified in the request.
505	HTTP Version not supported	The server does not support the HTTPS protocol version used in the request.

7.2 Error Codes

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.0001	Incorrect parameters.	Incorrect parameters.	Change the parameters based on the displayed information.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.0006	The request is null. Enter the request parameters.	The request is null. Enter the request parameters.	Enter the request parameters based on the error information.
400	CSS.0017	Invalid parameter letter case.	Invalid parameter letter case.	Set the parameter based on the naming rules.
400	CSS.0021	The account is frozen.	The account is frozen.	Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.0026	The log size is too large.	The log size is too large.	Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.0036	Internal error.	Internal error.	Try again later or Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.0047	The cluster does not exist or has been deleted.	The cluster does not exist or has been deleted.	The operation cannot be performed because the cluster does not exist.
400	CSS.0048	The OBS bucket does not exist.	The OBS bucket does not exist.	Check whether the OBS bucket exists or whether the user has the permission to access the bucket.
400	CSS.0050	Invalid cluster administrator password.	Invalid cluster administrator password.	Change the password based on the API reference.
400	CSS.0051	The current policy does not allow the #actionName # operation.	The current policy does not allow the #actionName # operation.	Grant the required permissions to the account.
400	CSS.0052	Insufficient CPU quota.	Insufficient CPU quota.	Release unused clusters, reduce the number of new CPUs, or increase the quota.
400	CSS.0053	Insufficient resource quota on the node.	Insufficient resource quota on the node.	Release unused clusters, reduce the number of new nodes, or increase the quota.
400	CSS.0054	Insufficient memory quota.	Insufficient memory quota.	Release unused clusters, reduce the memory size, or increase the quota.
400	CSS.0055	Insufficient disk quota.	Insufficient disk quota.	Release unused clusters, reduce the disk size, or increase the quota.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.0060	Insufficient disk capacity quota.	Insufficient disk capacity quota.	Release unused resources, reduce the disk capacity, or increase the quota.
400	CSS.0061	Insufficient EIP quota.	Insufficient EIP quota.	Release unused resources or increase the quota.
400	CSS.0062	Insufficient ELB quota.	Insufficient ELB quota.	Release unused resources or increase the quota.
400	CSS.0063	Insufficient VPC endpoint quota.	Insufficient VPC endpoint quota.	Release unused resources or increase the quota.
400	CSS.0064	Insufficient subnet IP address quota.	Insufficient subnet IP address quota.	Release unused resources or increase the quota.
400	CSS.0065	The disk has been sold out.	The disk has been sold out.	Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.0066	Invalid password.	Invalid password.	Change the password based on the API reference.
400	CSS.0068	The flavor has been sold out.	The flavor has been sold out.	Release unused resources.
400	CSS.0069	The disk has been sold out.	The disk has been sold out.	Release resources or increase the quota.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.0070	Interconnection with the enterprise project is not supported.	Interconnection with the enterprise project is not supported.	This site does not support enterprise projects.
400	CSS.0071	Invalid bandwidth size.	Invalid bandwidth size.	Change the bandwidth size based on the error information.
400	CSS.0072	Failed to create the SMN agency. Check whether the account has the required permission.	Failed to create the SMN agency. Check whether the account has the required permission.	Grant the required permissions to the account.
400	CSS.0074	The number of data nodes (including cold nodes) must be greater than or equal to three.	The number of data nodes (including cold nodes) must be greater than or equal to three.	Scale out the cluster or cancel the copy verification.
400	CSS.0075	The target flavor is not supported.	The target flavor is not supported.	The target flavor must be smaller than the original one.
400	CSS.0076	Invalid node role type.	Invalid node role type.	Change the role type based on the API parameter description.
400	CSS.0077	The cluster is being scaled out or scaled in.	The cluster is being scaled out or scaled in.	Perform the operation after the cluster status becomes normal.
400	CSS.0078	The cluster is in a status that does not support upgrade.	The cluster is in a status that does not support upgrade.	Perform the operation after the cluster status becomes normal.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.0080	The default agency already exists. Check whether the agency has the required permissions before using it.	The default agency already exists. Check whether the agency has the required permissions before using it.	Check whether the account has the required permissions.
400	CSS.0081	Failed to create the agency. Check whether the account has the permission to create an agency.	Failed to create the agency. Check whether the account has the permission to create an agency.	Grant the required permissions to the account.
400	CSS.0082	The record does not exist.	The record does not exist.	Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.0083	Some indexes have not been backed up. Perform full index backup in the snapshot module of the cluster.	Some indexes have not been backed up. Perform full index backup in the snapshot module of the cluster.	Perform full index backup in the snapshot module of the cluster.
400	CSS.1111	Invalid scale-out parameters.	Invalid scale-out parameters.	Check whether the scale-out parameters are correct.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.1112	The number of instances has reached the quota.	The number of instances has reached the quota.	Delete some clusters or increase the quota.
400	CSS.1119	The Kibana public access already exists.	The Kibana public access already exists.	Do not repeatedly bind an IP address for Kibana public access.
400	CSS.1120	The Kibana public access does not exist.	The Kibana public access does not exist.	Bind an IP address for Kibana public access.
400	CSS.1122	To change the security mode, disable the public access and try again.	To change the security mode, disable the public access and try again.	Disable the public access based on the error information.
400	CSS.1123	To change the security mode, disable the Kibana public access and try again.	To change the security mode, disable the Kibana public access and try again.	Disable the Kibana public access based on the error information.
400	CSS.2026	This is a weak password. Enter a strong password.	This is a weak password. Enter a strong password.	Set a strong password based on the error information.
400	CSS.3011	Invalid retention period.	Invalid retention period.	Check whether the value is an integer in the range 1 to 90.
400	CSS.3016	The backup file is being deleted. Try again later.	The backup file is being deleted. Try again later.	Perform the operation after the backup file is deleted.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.3026	The backup file does not exist.	The backup file does not exist.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.3030	The new instance volume cannot be smaller than the old instance volume.	The new instance volume cannot be smaller than the old instance volume.	Change the parameter values based on the displayed information.
400	CSS.5002	Invalid database type.	Invalid database type.	Currently, only Elasticsearch is supported.
400	CSS.5003	Invalid database version.	Invalid database version.	Elasticsearch and Kibana support the versions 7.6.2 and 7.10.2.
400	CSS.5004	The datastore parameter is not specified.	The datastore parameter is not specified.	Specify the datastore parameter.
400	CSS.5005	The database type or version is not supported.	The database type or version is not supported.	Elasticsearch and Kibana support the versions 7.6.2 and 7.10.2.
400	CSS.5006	Flavor is illegal.	Flavor is illegal.	Change the flavors based on the API parameter description.
400	CSS.5007	The selected flavors do not exist.	The selected flavors do not exist.	Change the flavors based on the API parameter description.
400	CSS.5008	Invalid volume type.	Invalid volume type.	Select a correct volume type.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5009	Invalid volume size.	Invalid volume size.	Enter a valid size.
400	CSS.5011	Invalid AZ.	Invalid AZ.	Select a correct AZ.
400	CSS.5017	Invalid retention period of automated backups.	Invalid retention period of automated backups.	Check whether the retention period of automated backups is set to 1 to 90 days.
400	CSS.5018	Invalid automated backup period.	Invalid automated backup period.	Only the full hour backup is supported. Add the time zone at the end of the time. The format is HH:mm z.
400	CSS.5020	Invalid automated backup start time.	Invalid automated backup start time.	Set a valid backup time based on the API parameter description.
400	CSS.5021	The VPC does not exist or does not belong to the user.	The VPC does not exist or does not belong to the user.	Check whether the VPC is correct.
400	CSS.5023	The security group does not exist or does not belong to the VPC.	The security group does not exist or does not belong to the VPC.	Check whether the security group is correct.
400	CSS.5033	The volume size is not an integer multiple of 10.	The volume size is not an integer multiple of 10.	Check whether the volume size is a multiple of 4 and 10.
400	CSS.5036	Datastore does not exist.	Datastore does not exist.	Currently, only Elasticsearch is supported.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5037	The parameter does not exist.	The parameter does not exist.	Check the parameter value range based on the API reference.
400	CSS.5038	The value of the parameter is out of range.	The value of the parameter is out of range.	Change the parameter values based on the displayed information.
400	CSS.5039	Duplicate parameters.	Duplicate parameters.	Delete duplicate parameters.
400	CSS.5044	Invalid backup prefix.	Invalid backup prefix.	Change the backup prefix based on the API parameter description.
400	CSS.5046	NICS is left blank.	NICS is left blank.	Specify NICS (VPC and subnet information).
400	CSS.5047	Invalid instance quantity.	Invalid instance quantity.	Check the parameter value range based on the API parameter description.
400	CSS.5048	Invalid extended parameters.	Invalid extended parameters.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5049	Invalid cluster name.	Invalid cluster name.	Change the cluster name.
400	CSS.5050	The cluster name already exists.	The cluster name already exists.	Change the cluster name.
400	CSS.5051	Invalid IP address.	Invalid IP address.	Enter a correct IP address.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5052	Available zone is illegal.	Available zone is illegal.	Select a correct AZ.
400	CSS.5055	Invalid flavor ID.	Invalid flavor ID.	Change the flavor ID based on the API parameter description.
400	CSS.5056	Invalid string length.	Invalid string length.	Check the supported string length based on the displayed information.
400	CSS.5057	The cluster name contains invalid characters or the name length is not allowed.	The cluster name contains invalid characters or the name length is not allowed.	Change the cluster name.
400	CSS.5059	Invalid database username.	Invalid database username.	Currently, only Elasticsearch is supported.
400	CSS.5061	Invalid volume type.	Invalid volume type.	Check whether the volume type is supported based on the API reference.
400	CSS.5063	Invalid value of the Specific parameter.	Invalid value of the Specific parameter.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5068	Invalid hard disk parameters.	Invalid hard disk parameters.	Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5069	Invalid node quantity.	Invalid node quantity.	Change the number of nodes based on the API parameter description.
400	CSS.5073	Instances of this type cannot be restarted separately.	Instances of this type cannot be restarted separately.	Do not select a type that is not contained in the resource.
400	CSS.5074	Subnet does not belong to VPC	Subnet does not belong to VPC	Check and enter correct subnet parameters.
400	CSS.5077	Invalid cluster name.	Invalid cluster name.	Change the cluster name.
400	CSS.5078	Invalid disk size.	Invalid disk size.	Change the disk size based on the API parameter description.
400	CSS.5087	Inconsistent disk specifications are used for the same type of instances.	Inconsistent disk specifications are used for the same type of instances.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5088	The SSL is not globally consistent.	The SSL is not globally consistent.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5090	In this status, the current operation is not allowed.	In this status, the current operation is not allowed.	Rectify the fault based on the error information and then perform the operation again.
400	CSS.5091	The number of nodes has reached the upper limit. Contact the administrator.	The number of nodes has reached the upper limit. Contact the administrator.	Change the number of nodes.
400	CSS.5092	Invalid storage information.	Invalid storage information.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5093	Invalid node specifications.	Invalid node specifications.	Change the node specifications based on the description about the API used to obtain the instance specifications list.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5094	Invalid node scale-out type.	Invalid node scale-out type.	Check and enter a valid node type.
400	CSS.5095	The current instance type does not support this operation.	The current instance type does not support this operation.	The local disk does not support this operation.
400	CSS.5098	Invalid EIP binding type.	Invalid EIP binding type.	Change the EIP binding type based on the API parameter description.
400	CSS.5106	Invalid instance type.	Invalid instance type.	Enter a correct instance type.
400	CSS.5107	Invalid cluster mode.	Invalid cluster mode.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5110	Inconsistent parameter groups are used for the same type of instances.	Inconsistent parameter groups are used for the same type of instances.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5111	The backup file ID does not exist.	The backup file ID does not exist.	Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5130	Invalid agency name.	Invalid agency name.	Perform operations based on the error information.
400	CSS.5134	The total number of master nodes must be an odd number.	The total number of master nodes must be an odd number.	Change the parameters based on the error information.
400	CSS.5135	The disk size must be a multiple of 20.	The disk size must be a multiple of 20.	Change the parameters based on the error information.
400	CSS.5137	Invalid tag key.	Invalid tag key.	Change the parameter based on the error information.
400	CSS.5138	Invalid tag value.	Invalid tag value.	Change the parameter based on the error information.
400	CSS.5140	The cluster does not support this operation.	The cluster does not support this operation.	Perform this operation after the cluster status becomes normal.
400	CSS.5141	Tag not found.	Tag not found.	Check whether the tag in the request body belongs to the cluster.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5158	Only bound public IP addresses can be unbound from clusters.	Only bound public IP addresses can be unbound from clusters.	Check the public access binding status.
400	CSS.5159	Only unbound public IP addresses can be bound to clusters.	Only unbound public IP addresses can be bound to clusters.	Check the public access binding status.
400	CSS.5160	Failed to bind the public IP address. Please try again.	Failed to bind the public IP address. Please try again.	Try again later or Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5164	Failed to update the public network access whitelist. Please try again.	Failed to update the public network access whitelist. Please try again.	Try again later or Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5167	No data nodes are selected.	No data nodes are selected.	Select at least one data node.
400	CSS.5169	The selected version does not support the security mode.	The selected version does not support the security mode.	For details about how to create a security cluster, see the Cloud Search Service User Guide.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5170	The volume parameter is not required for local disk specifications.	The volume parameter is not required for local disk specifications.	Delete the specified volume based on the error information.
400	CSS.5171	If httpsEnable is set to true, authorityEnable cannot be set to false.	If httpsEnable is set to true, authorityEnable cannot be set to false.	If httpsEnable is set to true, authorityEnable cannot be set to false.
400	CSS.5172	The number of instances must be greater than or equal to the number of AZs.	The number of instances must be greater than or equal to the number of AZs.	Increase the number of instances.
400	CSS.5174	Failed to obtain VPC information. Try again later.	Failed to obtain VPC information. Try again later.	Check whether the VPC exists and the network is available.
400	CSS.5177	The https enable parameter is not specified.	The https enable parameter is not specified.	Specify the httpsEnable parameter (mandatory).
400	CSS.5178	Invalid time zone.	Invalid time zone.	Change the time zone format based on the API parameter description.
400	CSS.5180	Kibana public access is not enabled.	Kibana public access is not enabled.	Enable Kibana public access.
400	CSS.5181	The whitelist does not exist.	The whitelist does not exist.	The current whitelist is empty. Add IP addresses to the whitelist and then perform related operations.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5182	The VPC endpoint service is not enabled.	The VPC endpoint service is not enabled.	Enable the VPC endpoint service and then perform related operations.
400	CSS.5301	This operation is not allowed because the cluster has been frozen.	This operation is not allowed because the cluster has been frozen.	If the account is frozen due to arrears, renew the account. Otherwise, Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.5302	Invalid subscription period type.	Invalid subscription period type.	Modify the subscription period type based on the API parameter description.
400	CSS.5303	Invalid subscription period.	Invalid subscription period.	Modify the subscription period based on the API parameter description.
400	CSS.5304	Invalid auto-renewal parameter.	Invalid auto-renewal parameter.	Modify the renewal parameters based on the API parameter description.
400	CSS.5305	Invalid automated payment parameters.	Invalid automated payment parameters.	Modify the automated payment parameters based on the API parameter description.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.5306	An unpaid order exists.	An unpaid order exists.	Pay for or cancel the order.
400	CSS.5307	An order is being processed.	An order is being processed.	Wait until the order is processed.
400	CSS.5308	The bandwidth exceeds the limit.	The bandwidth exceeds the limit.	Reduce the bandwidth.
400	CSS.5310	If httpsEnable is set to false, public network access cannot be enabled. If authorityEnable is set to false, Kibana public network access cannot be enabled.	If httpsEnable is set to false, public network access cannot be enabled. If authorityEnable is set to false, Kibana public network access cannot be enabled.	Enable the related functions based on the error information.
400	CSS.5313	Insufficient permissions on the VPC endpoint. Check whether the user has the VPC endpoint administrator role.	Insufficient permissions on the VPC endpoint. Check whether the user has the VPC endpoint administrator role.	Check whether the user has the required permissions based on the error information.
400	CSS.6000	Failed to create the cluster. Contact technical support or try again later.	Failed to create the cluster. Contact technical support or try again later.	Try again later. If the fault persists, Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.6001	Failed to scale out the cluster. Contact technical support or try again later.	Failed to scale out the cluster. Contact technical support or try again later.	Try again later. If the fault persists, Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.6003	Failed to restore the cluster. Contact technical support or try again later.	Failed to restore the cluster. Contact technical support or try again later.	Try again later. If the fault persists, Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.6004	Failed to create the node due to an ECS exception. Contact technical support or try again later. Error code: #failedReason #.	Failed to create the node due to an ECS exception. Contact technical support or try again later. Error code: #failedReason #.	Try again later. If the fault persists, Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.

Status Code	Error Codes	Error Message	Description	Solution
400	CSS.6009	Failed to create the snapshot. Contact technical support or try again later.	Failed to create the snapshot. Contact technical support or try again later.	Try again later. If the fault persists, Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
400	CSS.6012	Failed to restore the snapshot because the OBS object does not exist or has been deleted.	Failed to restore the snapshot because the OBS object does not exist or has been deleted.	Check whether the OBS bucket exists.
400	CSS.6013	Parameters cannot be modified because another operation is being performed on the cluster.	Parameters cannot be modified because another operation is being performed on the cluster.	Ensure that the cluster is available before modifying parameters.
400	CSS.6014	Unchangeable parameters are contained.	Unchangeable parameters are contained.	Delete the unchangeable parameters if any and try again.
400	CSS.6015	Parameters are not in the YAML format.	Parameters are not in the YAML format.	Modify the parameter format and try again.
403	CSS.0010	The instance ID or user ID may be null, or the operation is not authorized.	The instance ID or user ID may be null, or the operation is not authorized.	Check the account permission.

Status Code	Error Codes	Error Message	Description	Solution
403	CSS.0015	No resources are found or the access is denied.	No resources are found or the access is denied.	Insufficient permission. Add the required permission and try again.
403	CSS.5045	The current tenant is not allowed to perform this operation.	The current tenant is not allowed to perform this operation.	Check whether the tenant is frozen or has required permissions.
404	CSS.0002	The instance does not exist.	The instance does not exist.	Check whether the permission and resource exist.
404	CSS.0007	The VM information of the DB instance cannot be found.	The VM information of the DB instance cannot be found.	Try again later or Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
404	CSS.0016	Resource not found.	Resource not found.	Check whether the resource exists based on the displayed information.
404	CSS.0022	The instance does not exist or has been deleted.	The instance does not exist or has been deleted.	Check whether the resource exists.
404	CSS.0028	The object does not exist.	The object does not exist.	Check whether the current resource exists based on the displayed information.
404	CSS.1103	The instance does not exist!	The instance does not exist!	Check whether the cluster exists.

Status Code	Error Codes	Error Message	Description	Solution
404	CSS.3027	The cluster does not exist!	The cluster does not exist!	The cluster does not exist. No operations are supported.
404	CSS.5012	The region or AZ does not exist.	The region or AZ does not exist.	Select a correct region or AZ.
404	CSS.5022	The subnet does not exist or does not belong to the VPC.	The subnet does not exist or does not belong to the VPC.	Check and enter a subnet that belongs to the VPC.
404	CSS.5042	The source instance does not exist.	The source instance does not exist.	The request body contains an instance that does not exist in the cluster.
409	CSS.0011	This operation cannot be performed because another operation is being performed on the instance or the instance is faulty. Please try again later.	This operation cannot be performed because another operation is being performed on the instance or the instance is faulty. Please try again later.	Try again later.
409	CSS.0012	Operation is not allowed because the cluster has a standby instance.	Operation is not allowed because the cluster has a standby instance.	Perform the operation after the cluster backup is complete.
409	CSS.0029	The object is being created. Please wait.	The object is being created. Please wait.	The cluster or snapshot is being created. Do not perform other operations.
409	CSS.0031	The object already exists.	The object already exists.	Change the object name.

Status Code	Error Codes	Error Message	Description	Solution
409	CSS.0043	The snapshot name already exists.	The snapshot name already exists.	Change the snapshot name.
409	CSS.1102	The instance name already exists.	The instance name already exists.	Change the instance name.
409	CSS.3004	Backup is in progress.	Backup is in progress.	Perform the operation after the backup succeeded.
413	CSS.0030	The number of snapshots reaches the quota.	The number of snapshots reaches the quota.	Release the resources that do not need to be backed up.
415	CSS.0004	Invalid operation.	Invalid operation.	Change the parameter based on the error information.
415	CSS.5108	The operation is not supported.	The operation is not supported.	The operation is not supported.
419	CSS.0023	Insufficient resource.	Insufficient resource.	Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
422	CSS.0040	The parameter cannot be processed.	The parameter cannot be processed.	Try again later or Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.

Status Code	Error Codes	Error Message	Description	Solution
500	CSS.0005	Server error.	Server error.	Internal error. Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
500	CSS.1004	Database service error.	Database service error.	Try again later or Contact technical support. You can log in to the Huawei Cloud official website and click Contact Us in the upper right corner. Technical support will respond as soon as possible.
500	CSS.9999	Failed to process the request.	Failed to process the request.	Check whether the request path and request body are correct based on the displayed information.

7.3 Obtaining a Project ID and Name

A project ID or project name is required in some API requests. You need to obtain the project ID and name before calling an API.

Obtaining a Project ID and Name from the Console

1. Log in to the [CSS management console](#).
2. In the upper-right corner of the page, hover your mouse over the username and choose **My Credentials** from the drop-down list. The **API Credentials** page is displayed.
3. In the project list, view the **Project ID** and **Project Name**.

Obtaining a Project ID by Calling an API

The API for obtaining a project ID is **GET https://{iam-endpoint}/v3/projects**. *{iam-endpoint}* indicates the endpoint of IAM, which can be obtained from [Regions and Endpoints](#).

The following is an example response. For example, if CSS is deployed in region **xxx**, the value of **name** in the response body is **xxx**. The value of **id** in **projects** is the project ID.

```
{
  "projects": [
    {
      "domain_id": "65382450e8f64ac0870cd180d14exxxx",
      "is_domain": false,
      "parent_id": "65382450e8f64ac0870cd180d14exxxx",
      "name": "xxx", //Project name, the name of the deployment zone.
      "description": "",
      "links": {
        "next": null,
        "previous": null,
        "self": "https://www.example.com/v3/projects/a4a5d4098fb4474fa22cd05f897dxxxx"
      },
      "id": "a4a5d4098fb4474fa22cd05f897dxxxx", //Project ID
      "enabled": true
    }
  ],
  "links": {
    "next": null,
    "previous": null,
    "self": "https://www.example.com/v3/projects"
  }
}
```

7.4 Obtaining the Cluster ID

A cluster ID (**cluster_id**) is required for some URLs when an API is called. To obtain the cluster ID, perform the following steps:

Obtaining the Cluster ID by Calling an API

You can obtain the cluster ID by calling the API.

The API for obtaining the cluster ID is **GET https://{Endpoint}/v1.0/{project_id}/clusters**, where *{Endpoint}* indicates the IAM endpoint, which can be obtained from . For the project ID, see [Obtaining a Project ID and Name](#). For API authentication, see [Authentication](#).

The following is an example response. The value of **id** under **clusters** is the cluster ID.


```
{
  "totalSize": 1,
  "clusters": [ {
    "datastore": {
      "type": "elasticsearch",
      "version": "7.10.2",
      "subVersion": "7.10.2_24.3.0_0827",
      "isEosCluster": false,
      "supportSecuritymode": false
    },
    "instances": [ {
```

```
{
  "status" : "200",
  "type" : "ess",
  "id" : "{INSTANCE_ID}",
  "name" : "css-8bc5-ess-esn-1-1",
  "specCode" : "ess.spec-4u8g",
  "azCode" : "{AZ_CODE}",
  "volume" : {
    "type" : "ULTRAHIGH",
    "size" : 40,
    "resourceIds" : [ "{RESOURCE_ID}" ]
  },
  "ip" : "192.168.0.122",
  "resourceId" : "{RESOURCE_ID}"
},
"publicKibanaResp" : {
  "eipSize" : 10,
  "publicKibanaIp" : "100.95.152.28:9200",
  "elbWhiteListResp" : null,
  "bandwidthResourceId" : "18bec13f-5cc1-4631-867f-33505d15be12"
},
"elbWhiteList" : {
  "whiteList" : "",
  "enableWhiteList" : false
},
"updated" : "2023-10-09T02:07:13",
"name" : "css-8bc5",
"publicIp" : "100.85.222.202",
"created" : "2023-10-09T02:07:13",
"id" : "{CLUSTER_ID}",
"status" : "200",
"endpoint" : "192.168.0.122:9200",
"vpId" : "{VPC_ID}",
"subnetId" : "{SUBNET_ID}",
"securityGroupId" : "{SECURITY_GROUP_ID}",
"bandwidthResourceId" : "{BANDWIDTH_RESOURCE_ID}",
"bandwidthSize" : 3,
"httpsEnable" : true,
"authorityEnable" : true,
"diskEncrypted" : false,
"backupAvailable" : false,
"actionProgress" : { },
"actions" : [ ],
"enterpriseProjectId" : "0",
"tags" : [ ],
"period" : true
}
}
```

Obtaining the Cluster ID from the GaussDB(DWS) Console

1. Log in to the [CSS management console](#).
2. In the navigation pane on the left, click **Clusters**.
3. In the cluster list, find the cluster and copy the cluster ID in the **Name/ID** column.

Figure 7-1 Copying the cluster ID



- 4. Alternatively, click the cluster name to go to its details page.
- 5. Check the cluster ID under **Basic Information**.

Figure 7-2 Checking the cluster ID

