

Cloud Search Service

User Guide

Issue 04
Date 2024-09-06



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

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

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks 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 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.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Getting Started.....	1
1.1 Getting Started with Elasticsearch.....	1
2 Creating and Accessing a Cluster.....	7
2.1 Deploying a Cross-AZ Cluster.....	7
2.2 Clusters in Security Mode.....	9
2.3 Creating an Elasticsearch Cluster in Security Mode.....	14
2.4 Creating an Elasticsearch Cluster in Non-Security Mode.....	22
2.5 Accessing an Elasticsearch Cluster.....	28
2.6 Access Cluster.....	29
2.7 Viewing Cluster Information.....	40
3 Scaling In/Out a Cluster.....	44
3.1 Overview.....	44
3.2 Scaling Out a Cluster.....	45
3.3 Changing Specifications.....	47
3.4 Scaling in a Cluster.....	48
3.5 Removing Specified Nodes.....	50
3.6 Replacing a Specified Node.....	51
3.7 Adding Master/Client Nodes.....	52
3.8 Changing the Security Mode.....	53
4 Upgrading Versions.....	58
5 Importing Data to Elasticsearch.....	64
5.1 Using Logstash to Import Data to Elasticsearch.....	64
5.2 Using Kibana or APIs to Import Data to Elasticsearch.....	70
6 Managing Elasticsearch Clusters.....	74
6.1 Cluster and Storage Capacity Statuses.....	74
6.2 Introduction to the Cluster List.....	75
6.3 Index Backup and Restoration.....	76
6.4 Binding an Enterprise Project.....	82
6.5 Restarting a Cluster.....	82
6.6 Migrating Cluster Data.....	84
6.7 Deleting a Cluster.....	85

6.8 Managing Tags.....	85
6.9 Public Network Access.....	87
6.10 Managing Logs.....	89
6.11 Managing Plugins.....	91
6.12 Hot and Cold Data Storage.....	92
6.13 Configuring Parameters.....	93
6.14 VPC Endpoint Service.....	94
6.15 Kibana Public Access.....	97
7 Vector Retrieval.....	100
7.1 Description.....	100
7.2 Cluster Planning for Vector Retrieval.....	101
7.3 Creating a Vector Index.....	102
7.4 Querying Vectors.....	108
7.5 Optimizing the Performance of Vector Retrieval.....	112
7.6 (Optional) Pre-Building and Registering a Center Point Vector.....	113
7.7 Managing the Vector Index Cache.....	114
7.8 Sample Code for Vector Search on a Client.....	115
8 Working with Kibana.....	119
8.1 Logging In to Kibana.....	119
8.2 Creating a User and Granting Permissions by Using Kibana.....	119
8.3 Managing Index Statuses.....	125
8.3.1 Creating and Managing Indexes.....	125
8.3.2 Changing Policies.....	127
8.4 How Do I Connect a User-built Kibana with Elasticsearch on Huawei Cloud?.....	128
8.5 Kibana Usage Restrictions.....	129
9 Elasticsearch SQL.....	130
10 Connecting a Cluster to a Dedicated Load Balancer.....	135
10.1 Scenario.....	135
10.2 Connecting to a Dedicated Load Balancer.....	137
10.3 Sample Code for Two-Way Authentication During the Access to a Cluster.....	145
11 Enhanced Features.....	149
11.1 Storage-Compute Decoupling.....	149
11.1.1 Context.....	149
11.1.2 Freezing an Index.....	149
11.1.3 Configuring Cache.....	157
11.1.4 Enhanced Cold Data Query Performance.....	159
11.2 Flow Control.....	162
11.2.1 Flow Control.....	162
11.2.1.1 Context.....	163
11.2.1.2 HTTP/HTTPS Flow Control.....	164

11.2.1.3 Memory Flow Control.....	166
11.2.1.4 Global Path Whitelist for Flow Control.....	169
11.2.1.5 Request Sampling.....	170
11.2.1.6 Flow Control.....	171
11.2.1.7 Access Logs.....	174
11.2.1.8 CPU Flow Control.....	176
11.2.1.9 One-click Traffic Blocking.....	178
11.3 Large Query Isolation.....	179
11.3.1 Context.....	179
11.3.2 Procedure.....	179
11.4 Index Monitoring.....	184
11.4.1 Context.....	184
11.4.2 Enabling Index Monitoring.....	184
11.4.3 Checking the Index Read and Write Traffic.....	185
11.4.4 Checking Index Monitoring Information.....	187
11.4.5 kibana-monitor.....	191
11.5 Enhanced Monitoring.....	195
11.5.1 P99 Latency Monitoring.....	195
11.5.2 HTTP Status Code Monitoring.....	196
12 Monitoring.....	198
12.1 Supported Metrics.....	198
12.2 Configuring Cluster Monitoring.....	207
13 Auditing.....	210
13.1 Key Operations Recorded by CTS.....	210
13.2 Viewing Audit Logs.....	211
A Change History.....	212

1 Getting Started

1.1 Getting Started with Elasticsearch

This section describes how to use Elasticsearch for product search. You can use the Elasticsearch search engine of CSS to search for data based on the scenario example. The basic operation process is as follows:

- [Step 1: Create a Cluster](#)
- [Step 2: Import Data](#)
- [Step 3: Search for Data](#)
- [Step 4: Delete the Cluster](#)

Scenario Description

A women's clothing brand builds an e-commerce website. It uses traditional databases to provide a product search function for users. However, due to an increase in the number of users and business growth, the traditional databases have slow response and low accuracy. To improve user experience and user retention, the e-commerce website plans to use Elasticsearch to provide the product search function for users.

This section describes how to use Elasticsearch to provide the search function for users.

Assume that the e-commerce website provides the following data:

```
{
  "products":[
    {"productName":"Latest art shirts for women in autumn 2017","size":"L"},
    {"productName":"Latest art shirts for women in autumn 2017","size":"M"},
    {"productName":"Latest art shirts for women in autumn 2017","size":"S"},
    {"productName":"Latest jeans for women in spring 2018","size":"M"},
    {"productName":"Latest jeans for women in spring 2018","size":"S"},
    {"productName":"Latest jeans for women in spring 2017","size":"L"},
    {"productName":"Latest casual pants for women in spring 2017","size":"S"}
  ]
}
```

Step 1: Create a Cluster

Create a cluster using Elasticsearch as the search engine. In this example, suppose that you create a cluster named **Sample-ESCluster**. This cluster is used only for getting started with Elasticsearch. For this cluster, you are advised to select **ess.spec-4u8g** for **Node Specifications**, **High I/O** for **Node Storage Type**, and **40 GB** for **Node Storage Capacity**. For details, see [Creating an Elasticsearch Cluster in Non-Security Mode](#).

Create a cluster using Elasticsearch as the search engine. In this example, suppose that you create a cluster named **Sample-ESCluster**. This cluster is used only for getting started with Elasticsearch. For this cluster, you are advised to select **ess.spec-4u8g** for **Node Specifications**, **High I/O** for **Node Storage Type**, and **40 GB** for **Node Storage Capacity**. For details, see [Creating an Elasticsearch Cluster in Security Mode](#) or [Creating an Elasticsearch Cluster in Non-Security Mode](#).

After you create the cluster, switch to the cluster list to view the created cluster. If the **Status** of the cluster is **Available**, the cluster is created successfully.


Step 2: Import Data

CSS supports importing data to Elasticsearch using Logstash, Kibana, or APIs. Kibana lets you visualize your Elasticsearch data. The following procedure illustrates how to import data to Elasticsearch using Kibana.

1. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column to go to the Kibana login page.
 - Non-security cluster: The Kibana console is displayed.
 - Security cluster: Enter the username and password on the login page and click **Log In** to go to the Kibana console. The default username is **admin** and the password is the one specified during cluster creation.
2. In the navigation pane of Kibana on the left, choose **Dev Tools**.

The text box on the left is the input box. The triangle icon in the upper right corner of the input box is the command execution button. The text box on the right area is the result output box.

Figure 1-1 Console page



NOTE

The Kibana UI varies depending on the Kibana version.

3. On the **Console** page, run the following command to create index named **my_store**:

(Versions later than 7.x)

```
PUT /my_store
{
  "settings": {
    "number_of_shards": 1
  },
  "mappings": {
    "properties": {
      "productName": {
        "type": "text",
        "analyzer": "ik_smart"
      },
      "size": {
        "type": "keyword"
      }
    }
  }
}
```

The command output is similar to the following:

```
{
  "acknowledged" : true,
  "shards_acknowledged" : true,
  "index" : "my_store"
}
```

4. On the **Console** page, run the following command to import data to index named **my_store**:

(Versions later than 7.x)

```
POST /my_store/_doc/_bulk
{"index":{}}
{"productName":"Latest art shirts for women in autumn 2017","size":"L"}
{"index":{}}
{"productName":"Latest art shirts for women in autumn 2017","size":"M"}
{"index":{}}
{"productName":"Latest art shirts for women in autumn 2017","size":"S"}
{"index":{}}
{"productName":"Latest jeans for women in spring 2018","size":"M"}
{"index":{}}
{"productName":"Latest jeans for women in spring 2018","size":"S"}
{"index":{}}
{"productName":"Latest casual pants for women in spring 2017","size":"L"}
{"index":{}}
{"productName":"Latest casual pants for women in spring 2017","size":"S"}
```

If the value of the **errors** field in the command output is **false**, the data is imported successfully.

Step 3: Search for Data

- **Full-text search**

If you access the e-commerce website and want to search for commodities whose names include "spring jeans", enter "spring jeans" to begin your search. The following example shows the command to be executed on Kibana and the command output.

Command to be executed on Kibana:

(Versions later than 7.x)

```
GET /my_store/_search
{
  "query": {"match": {
```

```
"productName": "spring jeans"
  }}
}
```

The command output is similar to the following:

```
{
  "took" : 3,
  "timed_out" : false,
  "_shards" : {
    "total" : 1,
    "successful" : 1,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : {
      "value" : 4,
      "relation" : "eq"
    },
    "max_score" : 1.7965372,
    "hits" : [
      {
        "_index" : "my_store",
        "_type" : "_doc",
        "_id" : "9xf6VHIBfCl6SDJw7H5",
        "_score" : 1.7965372,
        "_source" : {
          "productName": "Latest jeans for women in spring 2018",
          "size" : "M"
        }
      },
      {
        "_index" : "my_store",
        "_type" : "_doc",
        "_id" : "-Bf6VHIBfCl6SDJw7H5",
        "_score" : 1.7965372,
        "_source" : {
          "productName": "Latest jeans for women in spring 2018",
          "size" : "S"
        }
      },
      {
        "_index" : "my_store",
        "_type" : "_doc",
        "_id" : "-Rf6VHIBfCl6SDJw7H5",
        "_score" : 0.5945667,
        "_source" : {
          "productName": "Latest casual pants for women in spring 2017",
          "size" : "L"
        }
      },
      {
        "_index" : "my_store",
        "_type" : "_doc",
        "_id" : "-hf6VHIBfCl6SDJw7H5",
        "_score" : 0.5945667,
        "_source" : {
          "productName": "Latest casual pants for women in spring 2017",
          "size" : "S"
        }
      }
    ]
  }
}
```

- Elasticsearch supports word segmentation. The preceding command segments "spring jeans" into "spring" and "jeans".

- Elasticsearch supports full-text search. The preceding command searches for the information about all commodities whose names include "spring" or "jeans".
 - Unlike traditional databases, Elasticsearch can return results in milliseconds by using inverted indexes.
 - Elasticsearch supports sorting by score. In the command output, information about the first two commodities contains both "spring" and "jeans", while that about the last two products contain only "spring". Therefore, the first two commodities rank prior to the last two due to high keyword match.
- **Aggregation result display**

The e-commerce website provides the function of displaying aggregation results. For example, it classifies commodities corresponding to "spring" based on the size so that you can collect the number of products of different sizes. The following example shows the command to be executed on Kibana and the command output.

Command to be executed on Kibana:

(Versions later than 7.x)

```
GET /my_store/_search
{
  "query": {
    "match": { "productName": "spring" }
  },
  "size": 0,
  "aggs": {
    "sizes": {
      "terms": { "field": "size" }
    }
  }
}
```

The command output is similar to the following:

(Versions later than 7.x)

```
{
  "took" : 3,
  "timed_out" : false,
  "_shards" : {
    "total" : 1,
    "successful" : 1,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : {
      "value" : 4,
      "relation" : "eq"
    },
    "max_score" : null,
    "hits" : [ ]
  },
  "aggregations" : {
    "sizes" : {
      "doc_count_error_upper_bound" : 0,
      "sum_other_doc_count" : 0,
      "buckets" : [
        {
          "key" : "S",
          "doc_count" : 2
        },
        {
          "key" : "L",
```

```
    "doc_count" : 1
  },
  {
    "key" : "M",
    "doc_count" : 1
  }
]
}
```

Step 4: Delete the Cluster

Once you understand the process and method of using Elasticsearch, you can perform the following steps to delete the cluster you created for the example and its data to avoid resource wastage.

NOTE

After you delete a cluster, its data cannot be restored. Exercise caution when deleting a cluster.

1. Log in to the CSS management console. In the navigation pane on the left, choose **Clusters > Elasticsearch**.
2. Locate the row that contains cluster **Sample-ESCluster** and click **More > Delete** in the **Operation** column.
3. In the displayed dialog box, enter the name of the cluster to be deleted and click **OK**.

2 Creating and Accessing a Cluster

2.1 Deploying a Cross-AZ Cluster

To prevent data loss and minimize the cluster downtime in case of service interruption, CSS supports cross-AZ cluster deployment. When creating a cluster, you can select two or three AZs in the same region. The system will automatically allocate nodes to these AZs.

Allocating Nodes

If you select two or three AZs when creating a cluster, CSS automatically enables the cross-AZ HA function and properly allocates nodes to different AZs. [Table 2-1](#) describes how the nodes are allocated.

 **NOTE**

- When creating a cluster, ensure that the number of selected nodes is no less than the number of AZs. Otherwise, cross-AZ deployment is not supported.
- If you enable master nodes when deploying a cross-AZ cluster, the master nodes will also be distributed to different AZs.
- The node quantity difference between any two AZs is no more than one.

Table 2-1 Number of nodes and AZ distribution

Nodes	Single AZ	Two AZs		Three AZs		
	AZ1	AZ1	AZ2	AZ1	AZ2	AZ3
1	1	Not supported		Not supported		
2	2	1	1	Not supported		
3	3	2	1	1	1	1
4	4	2	2	2	1	1
...

Setting Replicas

To maximize HA performance, properly determine the number of replicas.

- In two-AZ deployment, if one AZ becomes unavailable, the other AZ continues to provide services. In this case, at least one replica is required. Elasticsearch has one replica by default. You can retain the default value if you do not require higher read performance.
- In three-AZ deployment, if one AZ becomes unavailable, the other AZs continue to provide services. In this case, at least one replica is required. Elasticsearch has one replica by default. If you need more replicas to improve the cluster's ability to handle queries, modify `[<code>/topic/body/section/ul/li/p/idp:point/strong {<code>"}</code>] settings (strong) to change the number of replicas.`

You can run the following command to modify the number of index replicas:

```
curl -XPUT http://ip:9200/{index_name}/_settings -d '{"number_of_replicas":2}'
```

Alternatively, run the following command to specify the number of replicas in the template:

```
curl -XPUT http://ip:9200/_template/templatename -d '{"template": "*", "settings": {"number_of_replicas": 2}}'
```

NOTE

- **ip**: private network address
- **index_name**: index name
- **number_of_replicas**: number of replicas after modification. The value in the preceding command indicates that two replicas are required.

Possible Service Interruptions

The following table describes the possible service interruptions when an AZ of a two- or three-AZ cluster is faulty.

Table 2-2 Possible service interruptions

AZs	Master Nodes	Service Interruption Analysis
2	0	<ul style="list-style-type: none"> • When the number of nodes is an even number: <ul style="list-style-type: none"> – If half of data nodes are faulty, replace one node in the faulty AZ before you select the master node. • When the number of nodes is an odd number: <ul style="list-style-type: none"> – If the faulty AZ contains one more node than the normal AZ, you need to replace one node in the faulty AZ before you select the master node. For details about how to replace nodes, contact technical support. – If the faulty AZ contains one less node than the normal AZ, services will not be interrupted and you can select the master node.

AZs	Master Nodes	Service Interruption Analysis
2	3	<p>There is a 50% possibility for service interruption. When two dedicated master nodes are allocated to one AZ and another master node is allocated to the other AZ:</p> <ul style="list-style-type: none"> • If service interruption happens in the AZ with one master node, you can select a master node from the AZ that has two dedicated master nodes. • If service interruption happens in the AZ with two dedicated master nodes, you have no choice in the remaining AZ, because it has only one dedicated master node. In this case, services will be interrupted and you need to contact technical support.
3	0	<p>If you configure four nodes in three AZs, each AZ will have at least one node. If the AZ with two nodes is faulty, the services will be interrupted. You are not advised to configure four nodes when selecting three AZs.</p> <p>Generally, service interruption will not occur.</p>
3	3	Service interruption does not occur.

2.2 Clusters in Security Mode

When creating an Elasticsearch cluster, you can enable the security mode for it. Identity authentication is required when users access a security cluster. You can also authorize and encrypt security clusters.

Context

You can create clusters in multiple security modes. For details about the differences between security modes, see [Table 2-3](#).

Table 2-3 Cluster security modes

Security Mode	Scenario	Advantage	Disadvantage
Non-Security Mode	Intranet services and test scenarios	Simple. Easy to access.	Poor security. Anyone can access such clusters.

Security Mode	Scenario	Advantage	Disadvantage
Security Mode + HTTP Protocol	User permissions can be isolated, which is applicable to scenarios sensitive to cluster performance.	Security authentication is required for accessing such clusters, which improves cluster security. Accessing a cluster through HTTP protocol can retain the high performance of the cluster.	Cannot be accessed from the public network.
Security Mode + HTTPS Protocol	Scenarios that require high security and public network access.	Security authentication is required for accessing such clusters, which improves cluster security. HTTPS protocol allows public network to access such clusters.	The performance of clusters using HTTPS is 20% lower than that of using HTTP.

Identity Verification

To access a security cluster, you need to enter the username and password. The identity verification is required for the following two types of users:

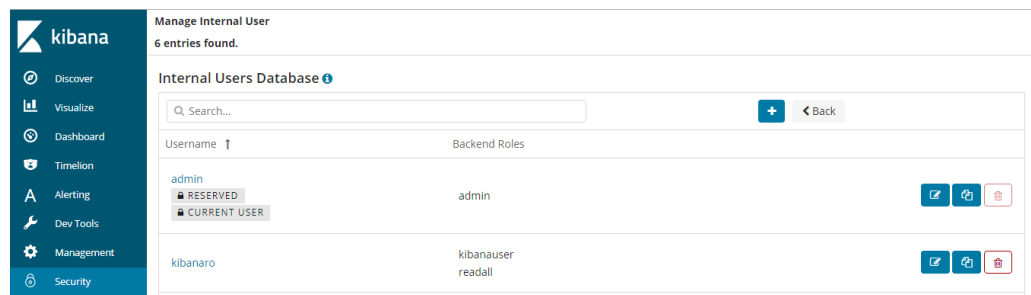
- Administrator: The default administrator username is **admin**, and the password is the one specified during cluster creation.
- Users: Enter the username and password created through Kibana.

Authorization

On the **Kibana** console, click **Security** to control user permissions in Elasticsearch clusters. You can configure hierarchical user permissions by cluster, index, document, and field. For details, see [Creating a User and Granting Permissions by Using Kibana](#).

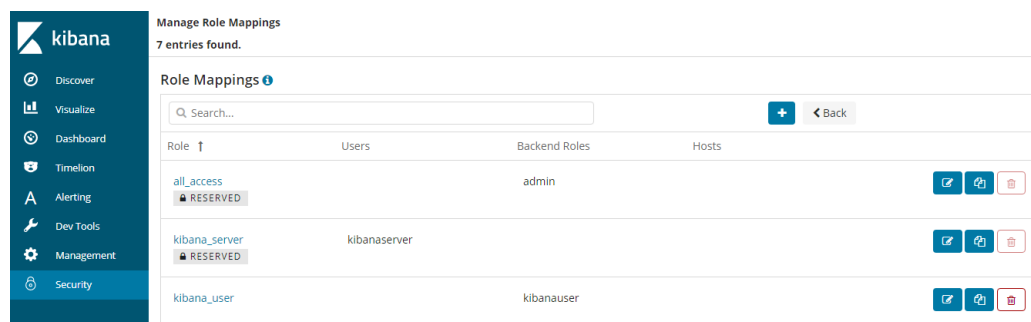
You can add or delete users, and map users to different roles for permissions control.

Figure 2-1 Configuring users



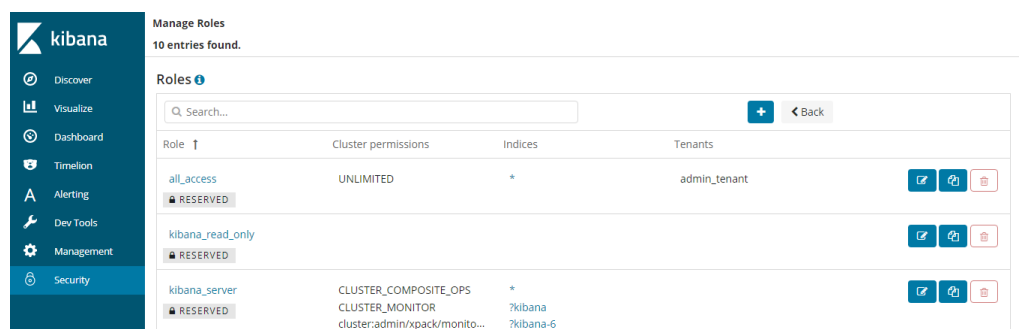
You can use role mapping to configure roles and map a user, backend role, and host name to a role.

Figure 2-2 Role mapping



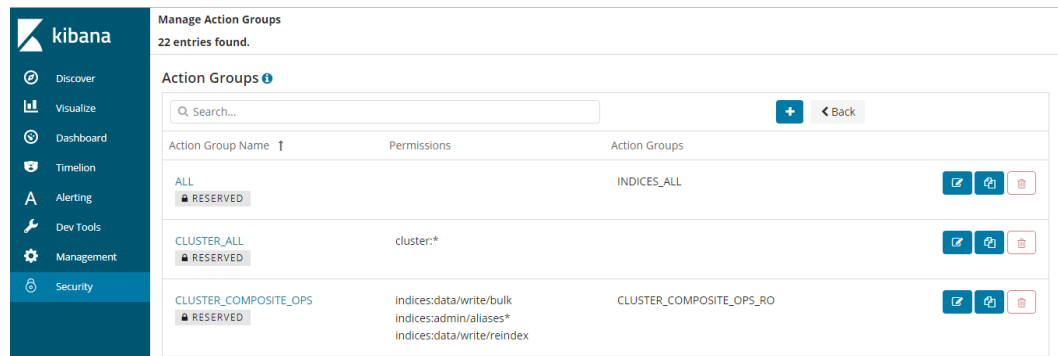
You can set permissions for each role to access clusters, indexes and documents and assign Kibana tenants different roles.

Figure 2-3 Configuring role permissions



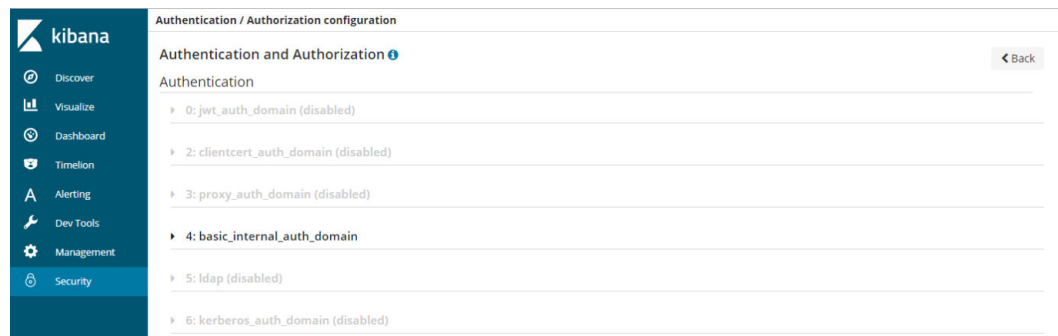
You can set action groups, assign the groups to roles, and configure the roles' permission for accessing indexes and documents.

Figure 2-4 Configuring action groups



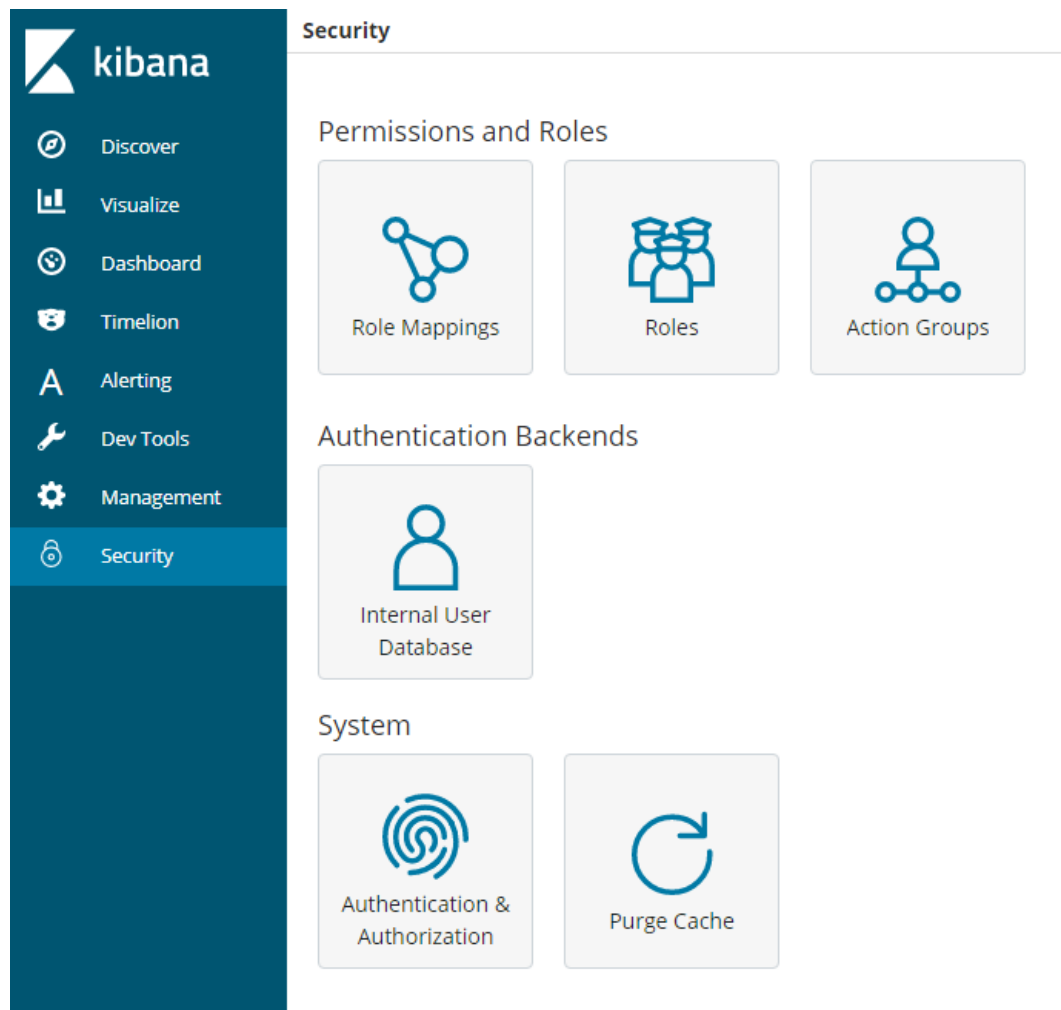
You can view the parameters of authentication and authorization for the current cluster. You can also run the **securityadmin** command to modify the configuration.

Figure 2-5 Viewing cluster parameters



You can also clear the security cache.

Figure 2-6 Clearing the security cache



Encryption

When key data is transferred between nodes or through the HTTP protocol, SSL/TLS encryption is used to ensure data security.

You can perform the preceding functions on Kibana, using **.yaml** files (not recommended), or by calling RESTful APIs. For more information about the security mode, see [Security](#).

Resetting the Administrator Password

If you want to change the administrator password of a security cluster or you have forgotten the password, reset the password.

1. On the **Clusters** page, locate the target cluster whose password you want to reset and click the cluster name. The **Cluster Information** page is displayed.
2. In the **Configuration** area, click **Reset** next to **Reset Password**.

 **NOTE**

- The password can contain 8 to 32 characters.
- The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The following special characters are supported: ~!@#\$\$%^&*()-_+=\|{}];;<.>/?
- Do not use the administrator name, or the administrator name spelled backwards.
- You are advised to change the password periodically.

2.3 Creating an Elasticsearch Cluster in Security Mode

This section describes how to create an Elasticsearch cluster in security mode.

 **NOTE**

- Public IP address access and Kibana public access can be used only after security mode is enabled.

Context

- When creating a cluster, the number of nodes that can be added varies according to the node type. For details, see [Table 2-4](#).

Table 2-4 Number of nodes in different types

Node Type	Number
ess	ess: 1-32
ess, ess-master	ess: 1-200 ess-master: an odd number ranging from 3 to 9
ess, ess-client	ess: 1-32 ess-client: 1-32
ess, ess-cold	ess: 1-32 ess-cold: 1-32
ess, ess-master, ess-client	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32
ess, ess-master, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-cold: 1-32
ess, ess-client, ess-cold	ess: 1-32 ess-client: 1-32 ess-cold: 1-32

Node Type	Number
ess, ess-master, ess-client, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32 ess-cold: 1-32
Details about the four node types: <ul style="list-style-type: none"> ● ess: the default node type that is mandatory for cluster creation. The other three node types are optional. ● ess-master: master node ● ess-client: client node ● ess-cold: cold data node 	

Procedure

1. Log in to the CSS management console.
2. Click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.
3. Specify **Region** and **AZ**.




Table 2-5 Parameter description for Region and AZ

Parameter	Description
Region	Select a region for the cluster from the drop-down list on the right.
AZ	Select AZs associated with the cluster region. You can select a maximum of three AZs. For details, see Deploying a Cross-AZ Cluster .

4. Configure basic cluster information.

Table 2-6 Description of basic parameters

Parameter	Description
Version	Select a cluster version from the drop-down list box.

Parameter	Description
Name	<p>Cluster name, which contains 4 to 32 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed and the value must start with a letter.</p> <p>NOTE After a cluster is created, you can modify the cluster name as required. Click the name of a cluster to be modified. On the displayed Basic Information page, click  next to the cluster name. After the modification is completed, click  to save the modification. If you want to cancel the modification, click .</p>

5. Configure cluster specifications.

Table 2-7 Parameter description

Parameter	Description
Nodes	<p>Number of nodes in a cluster.</p> <ul style="list-style-type: none"> • If neither a master node nor client node is enabled, the nodes specified by this parameter are used to serve as both the master node and client node. Nodes provide the cluster management, data storage, cluster access, and data analysis functions. To ensure data stability in the cluster, it is recommended that you set this parameter to a value no less than 3. • If only the master node function is enabled, nodes specified by this parameter are used to store data and provide functions of client nodes. • If both the master and client node functions are enabled, the nodes specified by this parameter are only used for storing data. • If only the client node function is enabled, nodes specified by this parameter are used to store data and provide functions of the master node.
CPU Architecture	Currently, x86 and Kunpeng are supported. The supported type is determined by the actual regional environment.
Node Specifications	Specifications of nodes in a cluster. You can select a specified specification based on your needs. Each cluster supports only one specification.
Node Storage Type	In the current version, the following options are available: Common I/O , High I/O , and Ultra-high I/O .
Node Storage Capacity	Storage space. Its value varies with node specifications. The node storage capacity must be a multiple of 20.

Parameter	Description
Master node	<p>The master node manages all nodes in the cluster. If more than 20 nodes are required to store and analyze a large amount of data, you are advised to enable the master node to ensure cluster stability. Otherwise, you are advised to set only the Nodes parameter and use the nodes as both master and client nodes.</p> <p>After enabling the master node, specify Node Specifications, Nodes, and Node Storage Type. The value of Nodes must be an odd number equal to or greater than 3. Up to nine nodes are supported. The value of Node Storage Capacity is fixed. You can select a storage type based on your needs.</p>
Client node	<p>The client node allows clients to access clusters and analyze data. If more than 20 nodes are required to store and analyze a large amount of data, you are advised to enable the client node to ensure cluster stability. Otherwise, you are advised to set only the Nodes parameter and use the nodes as both master and client nodes.</p> <p>After enabling the client node, specify Node Specifications, Nodes and Node Storage Type. The value of Nodes ranges from 1 to 32. The value of Node Storage Capacity is fixed. You can select a storage type based on your needs.</p>
Cold data node	<p>The cold data node is used to store historical data, for which query responses can be returned in minutes. If you do not require a quick query response, store historical data on cold data nodes to reduce costs.</p> <p>After enabling cold data node, configure Node Specifications, Nodes, Node Storage Type, and Node Storage Capacity. The value of Nodes ranges from 1 to 32. Select Node Storage Type and Node Storage Capacity as requirement.</p> <p>After the cold data node is enabled, CSS automatically adds cold and hot tags to related nodes.</p>

6. Set the enterprise project.
When creating a CSS cluster, you can bind an enterprise project to the cluster if you have enabled the enterprise project function. You can select an enterprise project created by the current user from the drop-down list on the right or click **View Project Management** to go to the **Enterprise Project Management** console and create a new project or view existing projects.
7. Click **Next: Configure Network**. Configure the cluster network.

Table 2-8 Network configuration parameters

Parameter	Description
VPC	<p>A VPC is a secure, isolated, and logical network environment.</p> <p>Select the target VPC. Click View VPC to enter the VPC management console and view the created VPC names and IDs. If no VPCs are available, create one.</p> <p>NOTE The VPC must contain CIDRs. Otherwise, cluster creation will fail. By default, a VPC will contain CIDRs.</p>
Subnet	<p>A subnet provides dedicated network resources that are isolated from other networks, improving network security.</p> <p>Select the target subnet. You can access the VPC management console to view the names and IDs of the existing subnets in the VPC.</p>
Security Group	<p>A security group is a collection of access control rules for ECSs that have the same security protection requirements and are mutually trusted in a VPC. To view more details about the security group, click View Security Group.</p> <p>NOTE</p> <ul style="list-style-type: none"> • For cluster access purposes, ensure that the security group contains port 9200. • If your cluster version is 7.6.2 or later, ensure that all the ports used for communication between nodes in the same security group are allowed. If such settings cannot be configured, ensure at least the access to port 9300 is allowed.
Security Mode	<p>After the security mode is enabled, communication will be encrypted and authentication required for the cluster.</p> <ul style="list-style-type: none"> • The default Administrator Username is admin. • Set and confirm the Administrator Password. This password will be required when you access this cluster.

Parameter	Description
HTTPS Access	<p>HTTPS access can be enabled only after the security mode of the cluster is enabled. After HTTPS access is enabled, communication is encrypted when you access the cluster.</p> <p>NOTE Security clusters use HTTPS for communication, which is much slower than non-security clusters that use HTTP for communication. If you want fast read performance and the permission provided by the security mode to isolate resources (such as indexes, documents, and fields), you can disable the HTTPS Access function. After HTTPS Access is disabled, HTTP protocol is used for cluster communication. In this case, data security cannot be ensured and public IP address cannot be used.</p>
Public IP Address	<p>If HTTPS Access is enabled, you can configure Public Network Access and obtain an IP address for public network access. This IP address can be used to access this security cluster through the public network. For details, see Public Network Access.</p>

8. Click **Next: Configure Advanced Settings**. Configure the automatic snapshot creation and other functions.
 - a. Configure **Cluster Snapshot**. Set basic configuration and snapshot configuration.

The cluster snapshot function is enabled by default. You can also disable this function as required. To store automatic snapshots in OBS, an agency will be created to access OBS. Additional cost will be incurred if snapshots are stored in standard storage.

Table 2-9 Cluster snapshot parameter

Parameter	Description
OBS bucket	<p>Select an OBS bucket for storing snapshots from the drop-down list box. You can also click Create Bucket on the right to create an OBS bucket.</p> <p>The created or existing OBS bucket must meet the following requirements:</p> <ul style="list-style-type: none"> • Storage Class is Standard.

Parameter	Description
Backup Path	<p>Storage path of the snapshot in the OBS bucket.</p> <p>The backup path configuration rules are as follows:</p> <ul style="list-style-type: none"> The backup path cannot contain the following characters: \:*?"<> The backup path cannot start with a slash (/). The backup path cannot start or end with a period (. The backup path cannot contain more than 1,023 characters.
IAM Agency	<p>IAM agency authorized by the current account to CSS access or maintain data stored in the OBS bucket. You can also click Create IAM Agency on the right to create an IAM agency.</p> <p>The created or existing IAM agency must meet the following requirements:</p> <ul style="list-style-type: none"> Agency Type must be Cloud service. Set Cloud Service to Elasticsearch or CSS. The agency must have the Tenant Administrator permission for the OBS project in .

Table 2-10 Automatic snapshot creation parameter

Parameter	Description
Snapshot Name Prefix	<p>The snapshot name prefix contains 1 to 32 characters and must start with a lowercase letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed. A snapshot name consists of a snapshot name prefix and a timestamp, for example, snapshot-1566921603720.</p>
Time Zone	<p>Time zone for the backup time, which cannot be changed. Specify Backup Started Time based on the time zone.</p>
Backup Start Time	<p>The time when the backup starts automatically every day. You can specify this parameter only in full hours, for example, 00:00 or 01:00. The value ranges from 00:00 to 23:00. Select a time from the drop-down list.</p>
Retention Period (days)	<p>The number of days that snapshots are retained in the OBS bucket. The value ranges from 1 to 90. You can specify this parameter as required. The system automatically deletes expired snapshots every hour at half past the hour.</p>

Figure 2-7 Setting parameters for automatic snapshot creation

Snapshot Name Prefix ✕ ?

Time Zone GMT+08:00

Backup Start Time ▼ ?

Retention Period (days) ?

- b. Configure advanced settings for the cluster.
 - **Default:** The **VPC Endpoint Service**, **Kibana Public Access**, and **Tag** functions are disabled by default. You can manually enable these functions after the cluster is created.
 - **Custom:** You can enable the **VPC Endpoint Service**, **Kibana Public Access**, and **Tag** functions as required.

Table 2-11 Parameters for advanced settings

Parameter	Description
VPC Endpoint Service	After enabling this function, you can obtain a private domain name for accessing the cluster in the same VPC. For details, see VPC Endpoint Service .
Kibana Public Access	You can configure this parameter only when security mode is enabled for a cluster. After enabling this function, you can obtain a public IP address for accessing Kibana. For details, see Kibana Public Access .
Tag	Adding tags to clusters can help you identify and manage your cluster resources. You can customize tags or use tags predefined by Tag Management Service (TMS). For details, see Managing Tags .

9. Click **Next: Confirm**. Check the configuration and click **Next** to create a cluster.
10. Click **Back to Cluster List** to switch to the **Clusters** page. The cluster you created is listed on the displayed page and its status is **Creating**. If the cluster is successfully created, its status will change to **Available**.

If the cluster creation fails, create the cluster again.

2.4 Creating an Elasticsearch Cluster in Non-Security Mode

This section describes how to create an Elasticsearch cluster in non-security mode.

Context

- When creating a cluster, the number of nodes that can be added varies according to the node type. For details, see [Table 2-12](#).

Table 2-12 Number of nodes in different types

Node Type	Number
ess	ess: 1-32
ess, ess-master	ess: 1-200 ess-master: an odd number ranging from 3 to 9
ess, ess-client	ess: 1-32 ess-client: 1-32
ess, ess-cold	ess: 1-32 ess-cold: 1-32
ess, ess-master, ess-client	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32
ess, ess-master, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-cold: 1-32
ess, ess-client, ess-cold	ess: 1-32 ess-client: 1-32 ess-cold: 1-32
ess, ess-master, ess-client, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32 ess-cold: 1-32

Node Type	Number
Details about the four node types: <ul style="list-style-type: none"> ● ess: the default node type that is mandatory for cluster creation. The other three node types are optional. ● ess-master: master node ● ess-client: client node ● ess-cold: cold data node 	

Procedure




1. Log in to the CSS management console.
2. Click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.
3. Specify **Region** and **AZ**.

Table 2-13 Parameter description for Region and AZ

Parameter	Description
Region	Select a region for the cluster from the drop-down list on the right.
AZ	Select AZs associated with the cluster region. You can select a maximum of three AZs. For details, see Deploying a Cross-AZ Cluster .

4. Configure basic cluster information.

Table 2-14 Description of basic parameters

Parameter	Description
Version	Select a cluster version from the drop-down list box.
Name	Cluster name, which contains 4 to 32 characters. Only letters, numbers, hyphens (-), and underscores (_) are allowed and the value must start with a letter. NOTE After a cluster is created, you can modify the cluster name as required. Click the name of a cluster to be modified. On the displayed Basic Information page, click  next to the cluster name. After the modification is completed, click  to save the modification. If you want to cancel the modification, click  .

5. Configure cluster specifications.

Table 2-15 Parameter description

Parameter	Description
Nodes	<p>Number of nodes in a cluster.</p> <ul style="list-style-type: none"> • If neither a master node nor client node is enabled, the nodes specified by this parameter are used to serve as both the master node and client node. Nodes provide the cluster management, data storage, cluster access, and data analysis functions. To ensure data stability in the cluster, it is recommended that you set this parameter to a value no less than 3. • If only the master node function is enabled, nodes specified by this parameter are used to store data and provide functions of client nodes. • If both the master and client node functions are enabled, the nodes specified by this parameter are only used for storing data. • If only the client node function is enabled, nodes specified by this parameter are used to store data and provide functions of the master node.
CPU Architecture	Currently, x86 and Kunpeng are supported. The supported type is determined by the actual regional environment.
Node Specifications	Specifications of nodes in a cluster. You can select a specified specification based on your needs. Each cluster supports only one specification.
Node Storage Type	In the current version, the following options are available: Common I/O , High I/O , and Ultra-high I/O .
Node Storage Capacity	Storage space. Its value varies with node specifications. The node storage capacity must be a multiple of 20.
Master node	<p>The master node manages all nodes in the cluster. If more than 20 nodes are required to store and analyze a large amount of data, you are advised to enable the master node to ensure cluster stability. Otherwise, you are advised to set only the Nodes parameter and use the nodes as both master and client nodes.</p> <p>After enabling the master node, specify Node Specifications, Nodes, and Node Storage Type. The value of Nodes must be an odd number equal to or greater than 3. Up to nine nodes are supported. The value of Node Storage Capacity is fixed. You can select a storage type based on your needs.</p>

Parameter	Description
Client node	<p>The client node allows clients to access clusters and analyze data. If more than 20 nodes are required to store and analyze a large amount of data, you are advised to enable the client node to ensure cluster stability. Otherwise, you are advised to set only the Nodes parameter and use the nodes as both master and client nodes.</p> <p>After enabling the client node, specify Node Specifications, Nodes and Node Storage Type. The value of Nodes ranges from 1 to 32. The value of Node Storage Capacity is fixed. You can select a storage type based on your needs.</p>
Cold data node	<p>The cold data node is used to store historical data, for which query responses can be returned in minutes. If you do not require a quick query response, store historical data on cold data nodes to reduce costs.</p> <p>After enabling cold data node, configure Node Specifications, Nodes, Node Storage Type, and Node Storage Capacity. The value of Nodes ranges from 1 to 32. Select Node Storage Type and Node Storage Capacity as requirement.</p> <p>After the cold data node is enabled, CSS automatically adds cold and hot tags to related nodes.</p>

6. Set the enterprise project.

When creating a CSS cluster, you can bind an enterprise project to the cluster if you have enabled the enterprise project function. You can select an enterprise project created by the current user from the drop-down list on the right or click **View Project Management** to go to the **Enterprise Project Management** console and create a new project or view existing projects.

7. Set network specifications of the cluster.

Table 2-16 Parameter description

Parameter	Description
VPC	<p>A VPC is a secure, isolated, and logical network environment.</p> <p>Select the target VPC. Click View VPC to enter the VPC management console and view the created VPC names and IDs. If no VPCs are available, create one.</p> <p>NOTE The VPC must contain CIDRs. Otherwise, cluster creation will fail. By default, a VPC will contain CIDRs.</p>

Parameter	Description
Subnet	<p>A subnet provides dedicated network resources that are isolated from other networks, improving network security.</p> <p>Select the target subnet. You can access the VPC management console to view the names and IDs of the existing subnets in the VPC.</p>
Security Group	<p>A security group is a collection of access control rules for ECSs that have the same security protection requirements and are mutually trusted in a VPC. To view more details about the security group, click View Security Group.</p> <p>NOTE</p> <ul style="list-style-type: none"> For cluster access purposes, ensure that the security group contains port 9200. If your cluster version is 7.6.2 or later, ensure that all the ports used for communication between nodes in the same security group are allowed. If such settings cannot be configured, ensure at least the access to port 9300 is allowed.
Security Mode	Security mode is disabled.

8. Click **Next: Configure Advanced Settings**. Configure the automatic snapshot creation and other functions.
 - a. Configure **Cluster Snapshot**. Set basic configuration and snapshot configuration.

The cluster snapshot function is enabled by default. You can also disable this function as required. To store automatic snapshots in OBS, an agency will be created to access OBS. Additional cost will be incurred if snapshots are stored in standard storage.

Table 2-17 Cluster snapshot parameter

Parameter	Description
OBS bucket	<p>Select an OBS bucket for storing snapshots from the drop-down list box. You can also click Create Bucket on the right to create an OBS bucket.</p> <p>The created or existing OBS bucket must meet the following requirements:</p> <ul style="list-style-type: none"> Storage Class is Standard.

Parameter	Description
Backup Path	<p>Storage path of the snapshot in the OBS bucket.</p> <p>The backup path configuration rules are as follows:</p> <ul style="list-style-type: none"> • The backup path cannot contain the following characters: \:*?"<> • The backup path cannot start with a slash (/). • The backup path cannot start or end with a period (.). • The backup path cannot contain more than 1,023 characters.
IAM Agency	<p>IAM agency authorized by the current account to CSS access or maintain data stored in the OBS bucket. You can also click Create IAM Agency on the right to create an IAM agency.</p> <p>The created or existing IAM agency must meet the following requirements:</p> <ul style="list-style-type: none"> • Agency Type must be Cloud service. • Set Cloud Service to Elasticsearch or CSS. • The agency must have the Tenant Administrator permission for the OBS project in Global service.

Table 2-18 Automatic snapshot creation parameter

Parameter	Description
Snapshot Name Prefix	<p>The snapshot name prefix contains 1 to 32 characters and must start with a lowercase letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed. A snapshot name consists of a snapshot name prefix and a timestamp, for example, snapshot-1566921603720.</p>
Time Zone	<p>Time zone for the backup time, which cannot be changed. Specify Backup Started Time based on the time zone.</p>
Backup Start Time	<p>The time when the backup starts automatically every day. You can specify this parameter only in full hours, for example, 00:00 or 01:00. The value ranges from 00:00 to 23:00. Select a time from the drop-down list.</p>
Retention Period (days)	<p>The number of days that snapshots are retained in the OBS bucket. The value ranges from 1 to 90. You can specify this parameter as required. The system automatically deletes expired snapshots every hour at half past the hour.</p>

Figure 2-8 Setting parameters for automatic snapshot creation

Snapshot Name Prefix	snapshot	x	?
Time Zone	GMT+08:00		
Backup Start Time	00:00	v	?
Retention Period (days)	-	35	+ ?

- b. Configure advanced settings for the cluster.
 - **Default:** The **Kibana Public Access**, and **Tag** functions are disabled by default. You can manually enable these functions after the cluster is created.
 - **Custom:** You can enable the **Tag** functions as required.

Table 2-19 Parameters for advanced settings

Parameter	Description
VPC Endpoint Service	After enabling this function, you can obtain a private domain name for accessing the cluster in the same VPC. For details, see VPC Endpoint Service .
Kibana Public Access	Clusters in non-security mode cannot access Kibana through the Internet.
Tag	Adding tags to clusters can help you identify and manage your cluster resources. You can customize tags or use tags predefined by Tag Management Service (TMS). For details, see Managing Tags .

9. Click **Next: Confirm**. Check the configuration and click **Next** to create a cluster.
10. Click **Back to Cluster List** to switch to the **Clusters** page. The cluster you created is listed on the displayed page and its status is **Creating**. If the cluster is successfully created, its status will change to **Available**.

If the cluster creation fails, create the cluster again.

2.5 Accessing an Elasticsearch Cluster

Elasticsearch clusters have built-in Kibana and Cerebro components. You can quickly access an Elasticsearch cluster through Kibana and Cerebro.

Access a Cluster Through Kibana

1. Log in to the CSS management console.
2. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column to go to the Kibana login page.
 - Non-security cluster: The Kibana console is displayed.
 - Security cluster: Enter the username and password on the login page and click **Log In** to go to the Kibana console. The default username is **admin** and the password is the one specified during cluster creation.
3. After the login is successful, you can access the Elasticsearch cluster through Kibana.

Accessing a Cluster Through Cerebro

1. Log in to the CSS management console.
2. On the **Clusters** page, locate the target cluster and click **More > Cerebro** in the **Operation** column to go to the Cerebro login page.
 - Non-security cluster: Click the cluster name on the Cerebro login page to go to the Cerebro console.
 - Security cluster: Click the cluster name on the Cerebro login page, enter the username and password, and click **Authenticate** to go to the Cerebro console. The default username is **admin** and the password is the one specified during cluster creation.
3. After the login is successful, you can access the Elasticsearch cluster through Cerebro.

2.6 Access Cluster

After a cluster is created, you can access the cluster to use Elasticsearch and perform operations, such as, defining index data, importing data, and searching for data. For more information about Elasticsearch, see the [Elasticsearch Reference](#). You can use any of the following methods to access a cluster:

- [Accessing a Cluster Using Kibana on the Management Console](#)
- [Accessing a Cluster by Calling Elasticsearch APIs on the ECS That Is Located in the Same VPC as the Cluster](#)
- [Accessing a Cluster Using Java API in Non-security Mode](#)
- [Accessing a Cluster Using the Java API in Security Mode with Elasticsearch](#)

Accessing a Cluster Using Kibana on the Management Console

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters**.
3. On the displayed **Clusters** page, locate the row containing the target cluster and click **Access Kibana** in the **Operation** column.

 NOTE

Normally when you click **Access Kibana**, a new window will be displayed. If no new window is displayed when you click **Access Kibana**, the pop-up has been blocked. In this case, manage pop-up blocking to allow access to the pop-up (the Kibana access address).

4. On the Kibana page that is displayed, you can create indexes, query indexes and documents, and analyze document fields. For details about how to import data to Elasticsearch, see the following sections:
 - [Using Logstash to Import Data to Elasticsearch](#)
 - [Using Kibana or APIs to Import Data to Elasticsearch](#)

Accessing a Cluster by Calling Elasticsearch APIs on the ECS That Is Located in the Same VPC as the Cluster

The ECS that you use to access the cluster by calling Elasticsearch APIs, must meet the following requirements. For details about how to create and log in to an ECS, see Logging In to a Linux ECS in the *Elastic Cloud Server User Guide* or Logging In to a Windows ECS.

- Sufficient disk space is allocated for the ECS.
- The ECS and the cluster must be in the same VPC.
- The security group of the ECS must be the same as that of the cluster.

If this requirement is not met, modify the ECS security group or configure the inbound and outbound rules of the ECS security group to allow the ECS security group to be accessed by all security groups of the cluster. For details, see the *Virtual Private Cloud User Guide*.

- Configure security group rule settings of the target CSS cluster. Set **Protocol** to **TCP** and **Port Range** to **9200** or a port range including port **9200** for both the outbound and inbound directions.

To access a cluster by calling Elasticsearch APIs on the ECS that is located in the same VPC as the cluster, perform the following steps:

1. Create and then log in to an ECS that meets the requirements.
2. To access a cluster, use the private network address and port number of one node in the cluster. You can obtain the private network addresses of nodes from the **Private Network Address** column in the cluster list. If there is only one node in the cluster, the private network address and port number of this node are displayed. If there are multiple nodes in the cluster, private network addresses and port numbers of all nodes are displayed.

Assume that there are two nodes in a cluster. Value **10.62.179.32:9200 10.62.179.33:9200** indicates that the private network addresses of the two nodes are **10.62.179.32** and **10.62.179.33** respectively, and port **9200** is used to access both nodes.

3. Run the cURL command to execute the API or call the API by using a program before accessing the cluster. For details about Elasticsearch operations and APIs, see the [Elasticsearch Reference](#).

For example, run the following cURL command to view the index information in the cluster. In this example, the private access IP address of one node in the cluster is **10.62.179.32** and the port number is **9200**.

- If the cluster you access does not have the security mode enabled, run the following command:

```
curl 'http://10.62.179.32:9200/_cat/indices'
```
- If the cluster you access has the security mode enabled, access the cluster using HTTPS and add the username, password and **-u** to the cURL command.

```
curl -u username:password -k 'https://10.62.179.32:9200/_cat/indices'
```

NOTE

In the preceding command, the private network address and port number of only one node in the cluster are used. If the node fails, the command will fail to be executed. If the cluster contains multiple nodes, you can replace the private network address and port number of the faulty node with those of any available node in the cluster. If the cluster contains only one node, restore the node and execute the command again.

If encryption has not been enabled for the communication with the cluster, the command output is similar to that shown in the following figure.

Figure 2-9 Command output

```
SZX1000355659:/home/elasticsearch-5.5.2/bin # curl 'http://10.62.179.32:9200/_cat/indices'
green open new_twitter BSVY8wt0SIWSXGzZ5U5mzw 5 1 3 0 21.3kb 10.6kb
green open .kibana ks71z4ggTUCy2UDWkXqEgw 1 1 2 0 22.8kb 11.4kb
green open tweets_1 FXOn8ykvQrmvBISyKbFRHA 5 1 0 0 1.5kb 810b
green open my_store AWybpSpLQPK_2T4cWDN-TQ 5 1 20 0 41.2kb 20.6kb
green open my_index QF5ARy2VQ6G0t8BzZEI86g 5 1 1 0 9kb 4.5kb
green open tweets_2 uLdSGZ0BS7uamlQfxDlEsQ 5 1 0 0 1.5kb 810b
green open twitter lzPIdrmRQpeBglI76SAYGA 5 2 3 0 40.5kb 13.5kb
green open my_index2 oLjbtIBPSNqeVIXgHOXQHg 5 1 0 0 1.8kb 955b
```

Accessing a Cluster Using Java API in Non-security Mode

For clusters in the non-security mode, you can use the `TransportClient` or `RestHighLevelClient` method to create a client. `RestHighLevelClient` is recommended.

- Create a client using the default method of the `TransportClient` class.

```
Settings settings = ImmutableSettings.settingsBuilder().put("client.transport.sniff",false).build();
TransportClient client = new TransportClient(settings) .addTransportAddress(new
InetSocketAddress("host1", 9300));
```
- Create a client using the default method of the `RestHighLevelClient` class.

```
RestHighLevelClient client = new RestHighLevelClient(
RestClient.builder(
new HttpHost("localhost", 9200, "http"));
```

Accessing a Cluster Using the Java API in Security Mode with Elasticsearch

After enabling the security mode function for a cluster, you need to access the cluster via HTTPS and enter your username and password for authentication.

Create a client using the `RestHighLevelClient` class.

Configure `CredentialsProvider` and `SSLIOSessionStrategy` as the member variables of the user-defined `SecuredHttpClientConfigCallback` class.

`SecuredHttpClientConfigCallback`: encapsulates all user-defined connection parameters.

`CredentialsProvider`: Elasticsearch API, which is used to encapsulate the username and password using the method provided by Elasticsearch.

SSLIOSessionStrategy: configures SSL parameters, including the SSL domain name authentication mode and certificate processing mode. The SSLContext class is used to encapsulate related settings.

You can access a cluster through either of the following modes: ignore certificates and use certificates.

- Ignore all certificates and skip certificate authentication.

Use the default X509TrustManager to construct the TrustManager. You do not need to override any method.

Use the TrustManager created in the last step as the parameter to construct SSLContext using the default method.

```
static TrustManager[] trustAllCerts = new TrustManager[] { new X509TrustManager() {
    @Override
    public void checkClientTrusted(X509Certificate[] chain, String authType) throws
CertificateException {

    }
    @Override
    public void checkServerTrusted(X509Certificate[] chain, String authType) throws
CertificateException {

    }
    @Override
    public X509Certificate[] getAcceptedIssuers() {
        return null;
    }
}
};
final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
credentialsProvider.setCredentials(AuthScope.ANY,
    new UsernamePasswordCredentials(userName, password));
SSLContext sc = null;
try{
    sc = SSLContext.getInstance("SSL");
    sc.init(null, trustAllCerts, new SecureRandom());
}catch(KeyManagementException e){
    e.printStackTrace();
}catch(NoSuchAlgorithmException e){
    e.printStackTrace();
}
SSLIOSessionStrategy sessionStrategy = new SSLIOSessionStrategy(sc, new
NullHostNameVerifier());

SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,credentialsProvider);

RestClientBuilder builder = RestClient.builder(new HttpHost(clusterAddress, 9200,
"https")).setHttpClientConfigCallback(httpClientConfigCallback);

RestHighLevelClient client = new RestHighLevelClient(builder);
```

- Upload the downloaded certificate (**CloudSearchService.cer**) to access a cluster.

You need to download the certificate (**CloudSearchService.cer**) from the cluster management page.

Upload the certificate to the client and use the keytool to convert the certificate into a format that can be read by Java. (The default password of the keytool is **changeit**).

```
keytool -import -alias Alias -keystore Certificate_output_path_and_alias -file Certificate_upload_path
```

Customize the TrustManager class, which is inherited from the X509TrustManager. Read the certificate generated in the last step, add it to the trust certificate, and rewrite the three methods of the X509TrustManager interface.

Use the TrustManager created in the last step as the parameter to construct SSLContext using the default method.

```
public static class MyX509TrustManager implements X509TrustManager {

    X509TrustManager sunJSSEX509TrustManager;
    MyX509TrustManager() throws Exception {
        File file = new File("certification file path");
        if (file.isFile() == false) {
            throw new Exception("Wrong Certification Path");
        }
        System.out.println("Loading KeyStore " + file + "...");
        InputStream in = new FileInputStream(file);
        KeyStore ks = KeyStore.getInstance("JKS");
        ks.load(in, "changeit".toCharArray());
        TrustManagerFactory tmf =
            TrustManagerFactory.getInstance("SunX509", "SunJSSE");
        tmf.init(ks);
        TrustManager tms [] = tmf.getTrustManagers();
        for (int i = 0; i < tms.length; i++) {
            if (tms[i] instanceof X509TrustManager) {
                sunJSSEX509TrustManager = (X509TrustManager) tms[i];
                return;
            }
        }
        throw new Exception("Couldn't initialize");
    }
}

final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
credentialsProvider.setCredentials(AuthScope.ANY,
    new UsernamePasswordCredentials(userName, password));

SSLContext sc = null;
try{
    TrustManager[] tm = {new MyX509TrustManager()};
    sc = SSLContext.getInstance("SSL", "SunJSSE");
    sc.init(null, tm, new SecureRandom());
}catch (Exception e) {
    e.printStackTrace();
}

SSLIOSessionStrategy sessionStrategy = new SSLIOSessionStrategy(sc, new
NullHostNameVerifier());

SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,credentialsProvider);
RestClientBuilder builder = RestClient.builder(new HttpHost(clusterAddress, 9200, "https"))
    .setHttpClientConfigCallback(httpClientConfigCallback);
RestHighLevelClient client = new RestHighLevelClient(builder);
```

- **Sample code**

When the code is running, transfer three parameters: access address, cluster login username, and password. The request is **GET /_search{"query": {"match_all": {}}}**.

 **NOTE**

The client version in the POM file must be the same as the cluster version.

ESSecuredClient class (Ignore certificates)

```
package securitymode;

import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.HttpHost;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.elasticsearch.action.search.SearchRequest;
```

```
import org.elasticsearch.action.search.SearchResponse;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.index.query.QueryBuilders;
import org.elasticsearch.search.SearchHit;
import org.elasticsearch.search.SearchHits;
import org.elasticsearch.search.builder.SearchSourceBuilder;

import java.io.IOException;
import java.security.KeyManagementException;
import java.security.NoSuchAlgorithmException;
import java.security.SecureRandom;
import java.security.cert.CertificateException;
import java.security.cert.X509Certificate;

import javax.net.ssl.HostnameVerifier;
import javax.net.ssl.SSLContext;
import javax.net.ssl.SSLSession;
import javax.net.ssl.TrustManager;
import javax.net.ssl.X509TrustManager;

public class ESsecuredClientIgnoreCerDemo {

    public static void main(String[] args) {
        String clusterAddress = args[0];
        String userName = args[1];
        String password = args[2];
        // Create a client.
        RestHighLevelClient client = initESClient(clusterAddress, userName, password);
        try {
            // Search match_all, which is equivalent to {"query":{"match_all":{"}}}.
            SearchRequest searchRequest = new SearchRequest();
            SearchSourceBuilder searchSourceBuilder = new SearchSourceBuilder();
            searchSourceBuilder.query(QueryBuilders.matchAllQuery());
            searchRequest.source(searchSourceBuilder);

            SearchResponse searchResponse = client.search(searchRequest, RequestOptions.DEFAULT);
            System.out.println("query result: " + searchResponse.toString());
            SearchHits hits = searchResponse.getHits();
            for (SearchHit hit : hits) {
                System.out.println(hit.getSourceAsString());
            }
            System.out.println("query success");
            Thread.sleep(2000L);
        } catch (InterruptedException | IOException e) {
            e.printStackTrace();
        } finally {
            try {
                client.close();
                System.out.println("close client");
            } catch (IOException e) {
                e.printStackTrace();
            }
        }
    }

    private static RestHighLevelClient initESClient(String clusterAddress, String userName, String password) {
        final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
        credentialsProvider.setCredentials(AuthScope.ANY, new
        UsernamePasswordCredentials(userName, password));
        SSLContext sc = null;
        try {
            sc = SSLContext.getInstance("SSL");
            sc.init(null, trustAllCerts, new SecureRandom());
        } catch (KeyManagementException | NoSuchAlgorithmException e) {
            e.printStackTrace();
        }
    }
}
```



```
    }
    SSLIOSessionStrategy sessionStrategy = new SSLIOSessionStrategy(sc, new
NullHostNameVerifier());
    SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
    credentialsProvider);
    RestClientBuilder builder = RestClient.builder(new HttpHost(clusterAddress, 9200, "https"))
    .setHttpClientConfigCallback(httpClientConfigCallback);
    RestHighLevelClient client = new RestHighLevelClient(builder);
    return client;
}

static TrustManager[] trustAllCerts = new TrustManager[] {
    new X509TrustManager() {
        @Override
        public void checkClientTrusted(X509Certificate[] chain, String authType) throws
CertificateException {
        }

        @Override
        public void checkServerTrusted(X509Certificate[] chain, String authType) throws
CertificateException {
        }

        @Override
        public X509Certificate[] getAcceptedIssuers() {
            return null;
        }
    }
};

public static class NullHostNameVerifier implements HostnameVerifier {
    @Override
    public boolean verify(String arg0, SSLSession arg1) {
        return true;
    }
}
}
```

ESSecuredClient class (Uses certificates)

```
package securitymode;

import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.HttpHost;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.elasticsearch.action.search.SearchRequest;
import org.elasticsearch.action.search.SearchResponse;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.index.query.QueryBuilders;
import org.elasticsearch.search.SearchHit;
import org.elasticsearch.search.SearchHits;
import org.elasticsearch.search.builder.SearchSourceBuilder;

import java.io.File;
import java.io.FileInputStream;
import java.io.IOException;
import java.io.InputStream;
import java.security.KeyStore;
import java.security.SecureRandom;
import java.security.cert.CertificateException;
import java.security.cert.X509Certificate;

import javax.net.ssl.HostnameVerifier;
```

```
import javax.net.ssl.SSLContext;
import javax.net.ssl.SSLSession;
import javax.net.ssl.TrustManager;
import javax.net.ssl.TrustManagerFactory;
import javax.net.ssl.X509TrustManager;

public class ESSecuredClientWithCerDemo {

    public static void main(String[] args) {
        String clusterAddress = args[0];
        String userName = args[1];
        String password = args[2];
        String cerFilePath = args[3];
        String cerPassword = args[4];
        // Create a client.
        RestHighLevelClient client = initESClient(clusterAddress, userName, password, cerFilePath,
cerPassword);
        try {
            // Search match_all, which is equivalent to {"query":{"match_all":{"}}}
            SearchRequest searchRequest = new SearchRequest();
            SearchSourceBuilder searchSourceBuilder = new SearchSourceBuilder();
            searchSourceBuilder.query(QueryBuilders.matchAllQuery());
            searchRequest.source(searchSourceBuilder);

            // query
            SearchResponse searchResponse = client.search(searchRequest, RequestOptions.DEFAULT);
            System.out.println("query result: " + searchResponse.toString());
            SearchHits hits = searchResponse.getHits();
            for (SearchHit hit : hits) {
                System.out.println(hit.getSourceAsString());
            }
            System.out.println("query success");
            Thread.sleep(2000L);
        } catch (InterruptedException | IOException e) {
            e.printStackTrace();
        } finally {
            try {
                client.close();
                System.out.println("close client");
            } catch (IOException e) {
                e.printStackTrace();
            }
        }
    }

    private static RestHighLevelClient initESClient(String clusterAddress, String userName, String
password,
String cerFilePath, String cerPassword) {
        final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
        credentialsProvider.setCredentials(AuthScope.ANY, new
UsernamePasswordCredentials(userName, password));
        SSLContext sc = null;
        try {
            TrustManager[] tm = {new MyX509TrustManager(cerFilePath, cerPassword)};
            sc = SSLContext.getInstance("SSL", "SunJSSE");
            //You can also use SSLContext sslContext = SSLContext.getInstance("TLSv1.2");
            sc.init(null, tm, new SecureRandom());
        } catch (Exception e) {
            e.printStackTrace();
        }

        SSLIOSessionStrategy sessionStrategy = new SSLIOSessionStrategy(sc, new
NullHostNameVerifier());
        SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
credentialsProvider);
        RestClient.Builder builder = RestClient.builder(new HttpHost(clusterAddress, 9200, "https"))
.setHttpClientConfigCallback(httpClientConfigCallback);
        RestHighLevelClient client = new RestHighLevelClient(builder);
    }
}
```

```
    return client;
}

public static class MyX509TrustManager implements X509TrustManager {
    X509TrustManager sunJSSEX509TrustManager;

    MyX509TrustManager(String cerFilePath, String cerPassword) throws Exception {
        File file = new File(cerFilePath);
        if (!file.isFile()) {
            throw new Exception("Wrong Certification Path");
        }
        System.out.println("Loading KeyStore " + file + "...");
        InputStream in = new FileInputStream(file);
        KeyStore ks = KeyStore.getInstance("JKS");
        ks.load(in, cerPassword.toCharArray());
        TrustManagerFactory tmf = TrustManagerFactory.getInstance("SunX509", "SunJSSE");
        tmf.init(ks);
        TrustManager[] tms = tmf.getTrustManagers();
        for (TrustManager tm : tms) {
            if (tm instanceof X509TrustManager) {
                sunJSSEX509TrustManager = (X509TrustManager) tm;
                return;
            }
        }
        throw new Exception("Couldn't initialize");
    }

    @Override
    public void checkClientTrusted(X509Certificate[] chain, String authType) throws
CertificateException {

    }

    @Override
    public void checkServerTrusted(X509Certificate[] chain, String authType) throws
CertificateException {

    }

    @Override
    public X509Certificate[] getAcceptedIssuers() {
        return new X509Certificate[0];
    }
}

public static class NullHostNameVerifier implements HostnameVerifier {
    @Override
    public boolean verify(String arg0, SSLSession arg1) {
        return true;
    }
}
}
```

SecuredHttpClientConfigCallback class

```
import org.apache.http.client.CredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.common.Nullable;
import java.util.Objects;
class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
    @Nullable
    private final CredentialsProvider credentialsProvider;
    /**
     * The {@link SSLIOSessionStrategy} for all requests to enable SSL / TLS encryption.
     */
    private final SSLIOSessionStrategy sslStrategy;
    /**
     * Create a new {@link SecuredHttpClientConfigCallback}.
     */
}
```

```
*
 * @param credentialsProvider The credential provider, if a username/password have been supplied
 * @param sslStrategy The SSL strategy, if SSL / TLS have been supplied
 * @throws NullPointerException if {@code sslStrategy} is {@code null}
 */
SecuredHttpClientConfigCallback(final SSLIOSessionStrategy sslStrategy,
                               @Nullable final CredentialsProvider credentialsProvider) {
    this.sslStrategy = Objects.requireNonNull(sslStrategy);
    this.credentialsProvider = credentialsProvider;
}
/**
 * Get the {@link CredentialsProvider} that will be added to the HTTP client.
 *
 * @return Can be {@code null}.
 */
@Nullable
CredentialsProvider getCredentialsProvider() {
    return credentialsProvider;
}
/**
 * Get the {@link SSLIOSessionStrategy} that will be added to the HTTP client.
 *
 * @return Never {@code null}.
 */
SSLIOSessionStrategy getSSLStrategy() {
    return sslStrategy;
}
/**
 * Sets the {@linkplain HttpAsyncClientBuilder#setDefaultCredentialsProvider(CredentialsProvider)
credential provider},
 *
 * @param httpClientBuilder The client to configure.
 * @return Always {@code httpClientBuilder}.
 */
@Override
public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
    // enable SSL / TLS
    httpClientBuilder.setSSLStrategy(sslStrategy);
    // enable user authentication
    if (credentialsProvider != null) {
        httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
    }
    return httpClientBuilder;
}
}
```

pom.xml file

```
<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/
XMLSchema-instance"
        xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/
maven-4.0.0.xsd">
    <modelVersion>4.0.0</modelVersion>
    <groupId>1</groupId>
    <artifactId>ESClient</artifactId>
    <version>1.0-SNAPSHOT</version>
    <name>ESClient</name>
    <!-- FIXME change it to the project's website -->
    <url>http://www.example.com</url>
    <properties>
        <project.build.sourceEncoding>UTF-8</project.build.sourceEncoding>
        <maven.compiler.source>1.7</maven.compiler.source>
        <maven.compiler.target>1.7</maven.compiler.target>
    </properties>
    <dependencies>
        <dependency>
            <groupId>junit</groupId>
            <artifactId>junit</artifactId>
            <version>4.11</version>
            <scope>test</scope>
```

```
</dependency>

<dependency>
  <groupId>org.elasticsearch</groupId>
  <artifactId>elasticsearch</artifactId>
  <version>7.6.2</version>
</dependency>
<dependency>
  <groupId>org.elasticsearch.client</groupId>
  <artifactId>elasticsearch-rest-high-level-client</artifactId>
  <version>7.6.2</version>
</dependency>
<dependency>
  <groupId>org.apache.logging.log4j</groupId>
  <artifactId>log4j-api</artifactId>
  <version>2.7</version>
</dependency>
<dependency>
  <groupId>org.apache.logging.log4j</groupId>
  <artifactId>log4j-core</artifactId>
  <version>2.7</version>
</dependency>
</dependencies>
<build>
  <pluginManagement><!-- lock down plugins versions to avoid using Maven defaults (may be
moved to parent pom) -->
    <plugins>
      <!-- clean lifecycle, see https://maven.apache.org/ref/current/maven-core/
lifecycles.html#clean_Lifecycle -->
      <plugin>
        <artifactId>maven-clean-plugin</artifactId>
        <version>3.1.0</version>
      </plugin>
      <!-- default lifecycle, jar packaging: see https://maven.apache.org/ref/current/maven-core/
default-bindings.html#Plugin_bindings_for_jar_packaging -->
      <plugin>
        <artifactId>maven-resources-plugin</artifactId>
        <version>3.0.2</version>
      </plugin>
      <plugin>
        <artifactId>maven-compiler-plugin</artifactId>
        <version>3.8.0</version>
      </plugin>
      <plugin>
        <artifactId>maven-surefire-plugin</artifactId>
        <version>2.22.1</version>
      </plugin>
      <plugin>
        <artifactId>maven-jar-plugin</artifactId>
        <version>3.0.2</version>
      </plugin>
      <plugin>
        <artifactId>maven-install-plugin</artifactId>
        <version>2.5.2</version>
      </plugin>
      <plugin>
        <artifactId>maven-deploy-plugin</artifactId>
        <version>2.8.2</version>
      </plugin>
      <!-- site lifecycle, see https://maven.apache.org/ref/current/maven-core/
lifecycles.html#site_Lifecycle -->
      <plugin>
        <artifactId>maven-site-plugin</artifactId>
        <version>3.7.1</version>
      </plugin>
    </plugins>
  </pluginManagement>
</build>
</dependency>
```

```

    <plugin>
      <artifactId>maven-project-info-reports-plugin</artifactId>
      <version>3.0.0</version>
    </plugin>
  </plugins>
</pluginManagement>
</build>
</project>


```

2.7 Viewing Cluster Information

On the **Cluster Information** page, you can view the information about a cluster, including the private network address, public IP address, version, and node.

1. Log in to the CSS management console.
2. Choose **Clusters > Elasticsearch**. The cluster list page is displayed.
3. Click a cluster name to go to the **Cluster Information** page and view the basic information about the cluster.

Table 2-20 Parameters for configuring basic information

Type	Parameter	Description
Cluster Information	Name	User-defined cluster name. You can click  on the right to change the cluster name.
	ID	Unique ID of a cluster, which is automatically generated by the system. Each cluster in the same region has a unique ID.
	Version	Cluster version information.
	Cluster Status	Current status of a cluster
	Task Status	Current task status of a cluster. If no task is in progress, -- is displayed.
	Created	Time when a cluster was created
	Cluster Storage Capacity (GB)	Storage capacity of a cluster
	Used Cluster Storage (GB)	Used storage capacity of a cluster
Configuration	Region	Region where a cluster is located
	AZ	AZ where a cluster is located
	VPC	VPC to which the cluster belongs
	Subnet	Subnet to which the cluster belongs

Type	Parameter	Description
	Security Group	<p>Security group to which a cluster belongs.</p> <p>To change the security group of a cluster, click Change Security Group on the right.</p> <p>NOTICE Before changing the security group, ensure that the port 9200 required for service access has been enabled. Incorrect security group configuration may cause service access failures. Exercise caution when performing this operation.</p>
	Security Mode	<p>Security mode of a cluster.</p> <ul style="list-style-type: none"> • Enabled: The current cluster is a security cluster. • Disabled: The current cluster is a non-security cluster. <p>For details about how to change the security mode of a cluster, see Changing the Security Mode.</p>
	Reset Password	<p>This parameter is displayed only for security clusters.</p> <p>Click Reset to change the password of the administrator account admin of the security cluster.</p> <p>NOTE Requirements for administrator passwords:</p> <ul style="list-style-type: none"> • The password can contain 8 to 32 characters. • The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The following special characters are supported: ~!@#%&*()-_+=\ [{}];:;<.>/? • Do not use the administrator name, or the administrator name spelled backwards. • You are advised to change the password periodically.
	Enterprise Project	<p>Enterprise project to which a cluster belongs.</p> <p>You can click the project name to view the basic information about the enterprise project.</p>

Type	Parameter	Description
	Access Control	<p>Whether to set access control for a cluster. This parameter is displayed only for clusters with public network access enabled.</p> <ul style="list-style-type: none"> • Enabled: Only IP addresses in the whitelist can access the cluster through the public network. • Disabled: Any IP address can access the cluster through the public network. <p>Click Set to configure the access control and the whitelist.</p>
	Bandwidth	<p>The bandwidth for public network access. This parameter is displayed only for clusters with public network access enabled.</p> <p>Click Edit to change the bandwidth size.</p>
	HTTPS Access	<p>Whether to enable the HTTPS access protocol for a cluster.</p> <ul style="list-style-type: none"> • Disabled: The HTTP protocol is used for cluster access. • Enabled: The HTTPS protocol is used for cluster access. Only security clusters can enable this function. If HTTPS Access is enabled, you can click Download Certificate to obtain the CER security certificate for accessing the security cluster. <p>NOTE For details about how to change the access mode of a cluster in security mode, see Switching the Protocol of Security Clusters.</p>
	Private Network Address	<p>Private IP address and port number of a cluster, which can be used to access the cluster. If the cluster has only one node, the IP address and port number of only one node are displayed, for example, 10.62.179.32:9200. If the cluster has multiple nodes, the IP addresses and port numbers of all nodes are displayed, for example, 10.62.179.32:9200,10.62.179.33:9200.</p>

Type	Parameter	Description
Node	Node Specifications	Specifications of nodes in a cluster
	Node Storage Type	Storage capacity and storage type of nodes in a cluster
	Nodes	Number of nodes in a cluster

3 Scaling In/Out a Cluster

3.1 Overview

You can scale in or out a cluster and change cluster specifications. In this way, you can improve cluster efficiency and reduce O&M costs.

Scaling Out a Cluster

- If a data node (ess) processes many data writing and querying requests and responds slowly, you can expand its storage capacity to improve its efficiency. If some nodes turn unavailable due to the excessive data volume or misoperations, you can add new nodes to ensure the cluster availability.
- **Cold data nodes** (ess-cold) are used to share the workload of data nodes. To prevent cold data loss, you can expand the storage capacity of the cold data node or add new ones.

Changing Specifications

- If the allocation of new indexes or shards takes too long or the node coordination and scheduling are inefficient, you can change the master node (ess-master) specifications.
- If too many tasks need to be distributed or too many results have been aggregated, you can change the client node (ess-client) specifications.
- If the writing and query of a node suddenly become slow, you can change the data node (ess) specifications.
- If cold data query becomes slow, you can change the cold node (ess-cold) specifications.

Scaling in a Cluster

- If a cluster can process existing data without fully using its resources, you can scale in the cluster to reduce costs.

Removing Specified Nodes

- If a cluster can process existing data without fully using its nodes, you can remove one or more specified nodes from the cluster to reduce costs.

Replacing a Specified Node

- If a node in the cluster is faulty, you can create a new node with the same specifications to replace it.

Adding Master/Client Nodes

- If the workloads on the data plane of a cluster increase, you can dynamically scale the cluster by adding master/client nodes.

Changing the Security Mode

After a cluster is created, its security mode can be changed using the following methods:

- Change a non-security cluster to a security cluster that uses HTTP or HTTPS protocol.
- Change a security cluster that uses HTTP or HTTPS protocol to a non-security cluster.
- Change the protocol of a security cluster.

3.2 Scaling Out a Cluster

If the workloads on the data plane of a cluster change, you can scale out the cluster by increasing the number or capacity of its nodes. Services are not interrupted during cluster scale-out.

Prerequisites

- The target cluster is available and has no tasks in progress.
- The target cluster has sufficient quotas available.

Constraints

- Node specifications cannot be modified during cluster scale-out.
- If you change the number and storage capacity of a specified type of node, nodes in other types will not be changed.
- The quota of nodes in different types varies. For details, see [Table 3-1](#).

Table 3-1 Number of nodes in different types

Node Type	Number
ess	ess: 1-32
ess, ess-master	ess: 1-200 ess-master: an odd number ranging from 3 to 9
ess, ess-client	ess: 1-32 ess-client: 1-32
ess, ess-cold	ess: 1-32 ess-cold: 1-32

Node Type	Number
ess, ess-master, ess-client	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32
ess, ess-master, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-cold: 1-32
ess, ess-client, ess-cold	ess: 1-32 ess-client: 1-32 ess-cold: 1-32
ess, ess-master, ess-client, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32 ess-cold: 1-32
<p>Details about the four node types:</p> <ul style="list-style-type: none"> ● ess: the default node type that is mandatory for cluster creation. The other three node types are optional. ● ess-master: master node ● ess-client: client node ● ess-cold: cold data node 	

Procedure

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. On the **Modify Configuration** page, choose the **Scale Cluster** tab and click **Scale out** to set parameters.
 - **Action**: Select **Scale out**.
 - **New Nodes**: The number of default nodes. For details about the quota, see [Table 3-1](#).
 - **Node Storage Capacity**: The storage capacity of the default nodes. The quota is determined by the **New Node Specifications**. The value must be a multiple of 20.

NOTE

If a cluster has master, client, or cold data nodes, you can change the number of these nodes and expand the storage capacity of cold data nodes. For details about the quotas of the master, client, and cold data node, see [Table 3-1](#).

4. Click **Next: Scale Now**.
5. Confirm the information and click **Submit**.
6. Click **Back to Cluster List** to switch to the **Clusters** page. The **Task Status** is **Scaling out**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled out.

3.3 Changing Specifications

If the workloads on the data plane of a cluster change, you can change its node specifications as needed.

Prerequisites

- The target cluster is available and has no tasks in progress.
- The target cluster has sufficient quotas available.
- When changing the node specifications, ensure that all service data has copies so the services will not be interrupted.

Run the **GET _cat/indices?v** command in Kibana. If the returned **rep** value is greater than **0**, the data has copies. If the returned **rep** value is **0**, the data has no copies. In this case, create snapshot for the cluster by referring to [Manually Creating a Snapshot](#).

- If the data volume is large, it may take long to modify the node specifications. You are advised to modify specifications during off-peak hours.

Constraints

- The node number and storage capacity cannot be modified when you change the node specifications.
- After decreasing cluster specifications, the cluster performance will deteriorate and service capabilities will be affected. Exercise caution when performing this operation.
- If a cluster has multiple node types, you can change the specifications of only one type at a time. After the change, nodes in other types still maintain their original specifications.
- Kibana is unavailable during specification change.
- During the specification modification, the nodes are stopped and restarted in sequence. It is a rolling process.

Procedure

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. On the **Modify Configuration** page, choose the **Scale Cluster** tab and click **Change Specifications** to set parameters.
 - **Action**: select **Change Specifications**.
 - **New Node Specifications** The specifications of the default data nodes. Select the specifications from the drop-down list as required.

- If a cluster has master nodes, client nodes, or cold data nodes, you can change their specifications.
- 4. Click **Next: Scale Now**.
- 5. Confirm the information and click **Submit**.
- 6. In the displayed **Verify Index Copy** dialog box, select **Verify index copies** if you need. Click **OK**.
 - If you selected **Verify index copies** and the cluster has no master node, indexes must have at least one copy and the cluster must have at least three nodes.
 - If you selected **Verify index copies** and the cluster has no master node, indexes must have at least one copy.
- 7. Click **Back to Cluster List** to switch to the **Clusters** page. The **Cluster Status** is **Configuration modified**. When **Cluster Status** changes to **Available**, the cluster specifications have been successfully modified.

3.4 Scaling in a Cluster

If a cluster can process existing data without fully using its resources, you can scale in the cluster to reduce costs. Services are not interrupted during cluster scale-in.

Prerequisites

The target cluster is available and has no tasks in progress.

Constraints

- The node specifications and storage capacity cannot be modified during scale-in.
- If you change the number and storage capacity of a specified type of node, nodes in other types will not be changed.
- Ensure that the disk usage after scale-in is less than 80% and each AZ of each node type has at least one node.
- When scaling in a cluster, the data in the node to be deleted is migrated to other nodes. The timeout threshold for data migration is five hours. If data migration is not complete within 5 hours, the cluster scale-in fails. You are advised to perform scale-in for multiple times when the cluster has huge amounts of data.
- For a cluster without master nodes, the number of remaining data nodes (including cold data nodes and other types of nodes) after scale-in must be greater than half of the original node number, and greater than the maximum number of index replicas.
- The quota of nodes in different types varies. For details, see [Table 3-2](#).

Table 3-2 Number of nodes in different types

Node Type	Number
ess	ess: 1-32

Node Type	Number
ess, ess-master	ess: 1-200 ess-master: an odd number ranging from 3 to 9
ess, ess-client	ess: 1-32 ess-client: 1-32
ess, ess-cold	ess: 1-32 ess-cold: 1-32
ess, ess-master, ess-client	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32
ess, ess-master, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-cold: 1-32
ess, ess-client, ess-cold	ess: 1-32 ess-client: 1-32 ess-cold: 1-32
ess, ess-master, ess-client, ess-cold	ess: 1-200 ess-master: an odd number ranging from 3 to 9 ess-client: 1-32 ess-cold: 1-32
<p>Details about the four node types:</p> <ul style="list-style-type: none"> ● ess: the default node type that is mandatory for cluster creation. The other three node types are optional. ● ess-master: master node ● ess-client: client node ● ess-cold: cold data node 	

Procedure

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. On the **Modify Configuration** page, choose the **Scale Cluster** tab and click **Scale in** to set parameters.

- **Action:** Select **Scale in**.
 - **New Nodes:** The number of default nodes. For details about the quota, see [Table 3-2](#).
 - If a cluster has master nodes, client nodes, or cold data nodes, you can change the node number. For details about the quotas of the master, client, and cold data node, see [Table 3-2](#).
4. Click **Next: Scale Now**.
 5. Confirm the information and click **Submit**.
 6. Click **Back to Cluster List** to switch to the **Clusters** page. The **Task Status** is **Scaling in**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled in.

3.5 Removing Specified Nodes

If a cluster can process existing data without fully using its nodes, you can remove one or more specified nodes from the cluster to reduce costs. Services will not be interrupted during the removal of specified nodes.

Prerequisites

The target cluster is available and has no tasks in progress.

Constraints

- Ensure that the disk usage after scale-in is less than 80% and each AZ of each node type has at least one node.
- In a cross-AZ cluster, the difference between the numbers of the same type nodes in different AZs cannot exceed 1.
- For a cluster without master nodes, the number of removed data nodes and cold data nodes in a scale-in must be fewer than half of the original number of data nodes and cold data nodes, and the number of remaining data nodes and cold 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 original master node number. After scale-in, there has to be an odd number of master nodes, and there has to be at least three of them.

Procedure

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. On the **Modify Configuration** page, click the **Scale In** tab.
4. Select the target nodes.
5. Click **Next: Scale Now**.
6. Confirm the information and click **Submit**.

7. Click **Back to Cluster List** to switch to the **Clusters** page. The **Task Status** is **Scaling in**. When **Cluster Status** changes to **Available**, the cluster has been successfully scaled in.

3.6 Replacing a Specified Node

If a node in the cluster is faulty, you can create a new node with the same specifications to replace it.

Prerequisites

The target cluster is available and has no tasks in progress.

Constraints

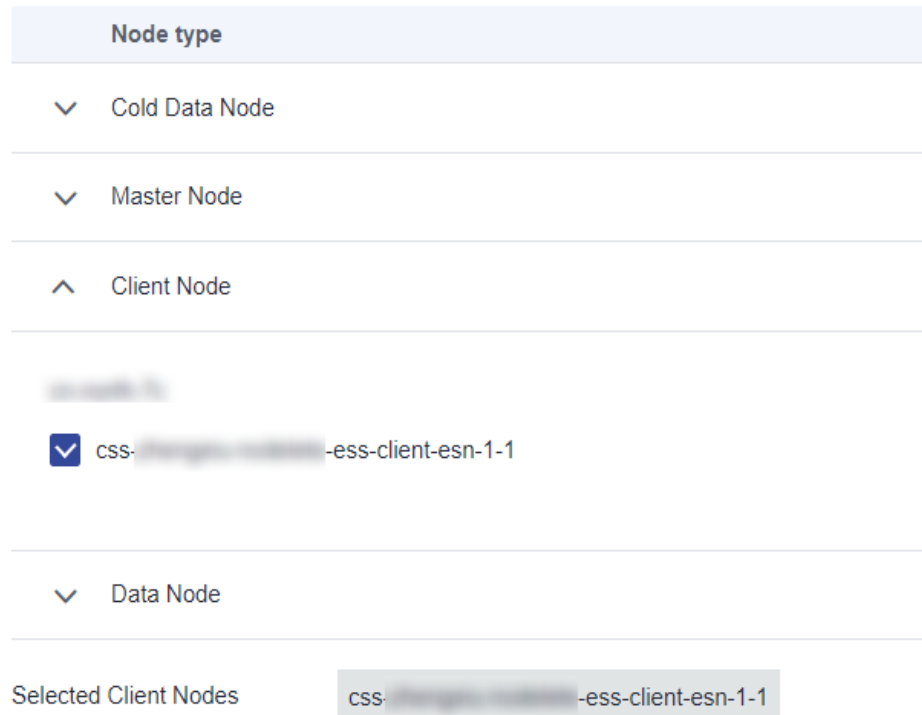
- Only one node can be replaced at a time.
- The ID, IP address, specifications, and AZ of the new node will be the same as those of the original one.
- The configurations you modified manually will not be retained after node replacement. For example, if you have manually added a return route to the original node, you need to add it to the new node again after the node replacement is complete.
- If the node you want to replace is a data node (ess) or cold data node (ess-cold), pay attention to the following precautions:
 - a. Before a data node or cold data node is replaced, its data needs to be migrated to other nodes. To properly store the data, ensure the maximum sum of replicas and primary shards of an index is smaller than the total number of data nodes (ess and ess-cold nodes) in the cluster. The node replacement duration depends heavily on the migration speed.
 - b. Clusters whose version is earlier than 7.6.2 cannot have closed indexes. Otherwise, data nodes or cold data nodes cannot be replaced.
 - c. The AZ of the node to be replaced must have two or more data nodes (including ess and ess-cold).
 - d. If the cluster of the node to be replaced does not have a master node (ess-master), the number of available data nodes (including ess and ess-cold) in the cluster must be greater than or equal to 3.
 - e. The preceding precautions do not apply if you are replacing a master node (ess-master) or client node (ess-client).
 - f. The precautions 1 to 4 do not apply if you are replacing a faulty node, regardless of its type. Faulty nodes are not included in **_cat/nodes**.

Procedure

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. On the **Modify Configuration** page, click the **Replace Node** tab.

4. Select the target node.

Figure 3-1 Replacing a specified node



5. Click **Submit**.
6. Click **Back to Cluster List** to switch to the **Clusters** page. The **Task Status** is **Upgrading**. When **Cluster Status** changes to **Available**, the node has been successfully replaced.

3.7 Adding Master/Client Nodes

If workloads on the data plane of a cluster increase, you can add master or client nodes as needed. Services are not interrupted while they are added.

Prerequisites

The target cluster is available and has no tasks in progress.

Constraints

- If a cluster already has master and client nodes, the **Add Master/Client Node** tab is not displayed on the **Modify Configuration** page. In this case, you need to add the master or client nodes by referring to [Scaling Out a Cluster](#).
- When you add master or client nodes, the number of nodes that can be configured varies depending on the node type. For details, see [Table 3-3](#).

Table 3-3 Number of nodes in different types

Node Flavor	Number
Master node	An odd number ranging from 3 to 9
Client node	1 to 32

Procedure

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. On the **Modify Configuration** page, choose the **Add Master/Client Node** tab.
4. Select the target node type and set the node specifications, quantity, and storage.
 - Master and client nodes cannot be added at the same time.
 - If a cluster already has a master or client node, you can only add nodes of the other type.

Figure 3-2 Adding a master or client node

5. Click **Next: Scale Now**.
6. Confirm the information and click **Submit**.
Return to the cluster list page. The **Task Status** of the cluster is **Scaling out**.
 - If you added a master node and **Cluster Status** changed to **Available**, the master node has been successfully added.

NOTICE

If the cluster version is earlier than 7.x, when the **Cluster Status** changes to **Available**, you need to restart all data nodes and cold data nodes in the cluster to make the new node take effect. Before the restart, the cluster may be unavailable. For details, see [Restarting a Cluster](#).

- If you added a client node and **Cluster Status** changed to **Available**, the client node has been added. You can restart data nodes and cold data nodes to shut down Cerebro and Kibana processes on the nodes.

3.8 Changing the Security Mode

After a cluster is created, its security mode can be changed using the following methods:

- [Switching from the Non-Security Mode to Security Mode](#)
- [Switching from the Security to Non-Security Mode](#)
- [Switching the Protocol of Security Clusters](#)

Context

You can create clusters in multiple security modes. For details about the differences between security modes, see [Table 3-4](#).

Table 3-4 Cluster security modes

Security Mode	Scenario	Advantage	Disadvantage
Non-Security Mode	Intranet services and test scenarios	Simple. Easy to access.	Poor security. Anyone can access such clusters.
Security Mode + HTTP Protocol	User permissions can be isolated, which is applicable to scenarios sensitive to cluster performance.	Security authentication is required for accessing such clusters, which improves cluster security. Accessing a cluster through HTTP protocol can retain the high performance of the cluster.	Cannot be accessed from the public network.
Security Mode + HTTPS Protocol	Scenarios that require high security and public network access.	Security authentication is required for accessing such clusters, which improves cluster security. HTTPS protocol allows public network to access such clusters.	The performance of clusters using HTTPS is 20% lower than that of using HTTP.

Prerequisites

- You are advised to back up data before changing the cluster security mode.
- The target cluster is available and has no tasks in progress.

Constraints

- Only clusters (whose version is 6.5.4 or later) created after November 2022 support security mode switching.
- A cluster automatically restarts when its security mode is being changed. Services are interrupted during the restart. The authentication mode for invoking the cluster will change after the restart, and client configurations need to be adjusted accordingly.
- If a cluster has already opened the Kibana session box, a session error message will be displayed after you change the cluster security mode. In this case, clear the cache and open Kibana again.

Switching from the Non-Security Mode to Security Mode

You can change a non-security cluster to a security cluster that uses HTTP or HTTPS. After a cluster's security mode is enabled, security authentication is required for accessing the cluster.

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. Choose the **Configure Security Mode** tab.
4. Enable the security mode. Enter and confirm the administrator password of the cluster.

Figure 3-3 Enabling the security mode

Security Mode If you enable this option, you will need to verify your identity to access the cluster

Administrator Username

Administrator Password

Confirm Password

HTTPS Access

5. Enable or disable **HTTPS Access**.
 - If you enable **HTTPS Access**: The HTTPS protocol is used to encrypt cluster communication and you can configure public networks to access the cluster.
 - If you disable **HTTPS Access**: The HTTP protocol is used and you cannot configure public networks to access the cluster.
6. Click **Submit**. Confirm the information and the cluster list page is displayed. The **Task Status** of the cluster is **The security mode is changing**. When the cluster status changes to **Available**, the security mode has been successfully changed.

Switching from the Security to Non-Security Mode

You can change a security cluster that uses HTTP or HTTPS to a non-security cluster. After a cluster's security mode is disabled, security authentication is no longer required for accessing the cluster.

NOTICE

- Clusters in non-security mode can be accessed without security authentication, and HTTP protocol is used to transmit data. Ensure the security of the cluster access environment and do not expose the access interface to the public network.
- During the switchover from the security mode to the non-security mode, the indexes of the original security cluster will be deleted. Back up data before disabling the security mode.
- If a security cluster has been bound to a public IP address, unbind it before changing the security mode.
- If a security cluster has enabled Kibana public network access, disable it before changing the security mode.

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. Choose the **Configure Security Mode** tab.
4. Disable the security mode.

Figure 3-4 Disabling the security mode



5. Click **Submit**. Confirm the information and the cluster list page is displayed. The **Task Status** of the cluster is **The security mode is changing**. When the cluster status changes to **Available**, the security mode has been successfully changed.

Switching the Protocol of Security Clusters

You can change the protocol of a security cluster.


NOTICE

If a security cluster has been bound to a public IP address, you need to unbind it before changing HTTPS protocol to HTTP.

1. Log in to the CSS management console.

2. In the navigation pane on the left, choose **Clusters**. On the displayed **Clusters** page, locate the target cluster and choose **More > Modify Configuration** in the **Operation** column.
3. Choose the **Configure Security Mode** tab.
4. Enable or disable **HTTPS Access**.

Figure 3-5 Configuring the protocol



- If you enable **HTTPS Access**:
HTTPS protocol is used to encrypt cluster communication and you can configure public network access.
 - If you disable **HTTPS Access**: An alarm message is displayed. Click **OK** to disable the function.
Cluster communication is no longer encrypted and the public network access function cannot be enabled.
5. Click **Submit**. Confirm the information and the cluster list page is displayed. The **Task Status** of the cluster is **The security mode is changing**. When the cluster status changes to **Available**, the security mode has been successfully changed.

4 Upgrading Versions

Same version upgrade and cross-version upgrade are supported. Same version upgrade is to upgrade the kernel patch of a cluster to fix problems or optimize performance. Cross-version upgrade is to upgrade the cluster version to enhance functions or incorporate versions.

Description

Principle

Nodes in the cluster are upgraded one by one so that services are not interrupted. The upgrade process is as follows: bring a node offline, migrate its data to another node, create a new node of the target version, and mount the NIC ports of the offline node to the new node to retain the node IP address. After a new node is added to the cluster, other nodes will be updated in the same way in sequence. If there is a large amount of data in a cluster, the upgrade duration depends on the data migration duration.

Process

1. [Pre-Upgrade Check](#)
2. [Creating a Snapshot](#)
3. [Creating an Upgrade Task](#)

Version Restrictions

The supported target version varies according to the current cluster version. For details, see [Table 4-1](#).

Table 4-1 Version restrictions

Current Version	Target Version
7.1.1	7.6.2, 7.10.2
7.6.2	7.10.2
7.9.3	7.10.2

Current Version	Target Version
<p>Note:</p> <ul style="list-style-type: none"> 7.6.2 and 7.10.2 are mainstream versions. You are advised to upgrade your clusters to these two versions. The supported target versions are displayed in the drop-down list of Target Image. Versions 5.x and 6.x cannot be upgraded to version 7.x. 	

Constraints

- A maximum of 20 clusters can be upgraded at the same time. You are advised to perform the upgrade during off-peak hours.
- Clusters that have ongoing tasks cannot be upgraded.
- Once started, an upgrade task cannot be stopped until it succeeds or fails.
- During the upgrade, nodes are replaced one by one. Requests sent to a node that is being replaced may fail. In this case, you are advised to access the cluster through the VPC Endpoint service or a dedicated load balancer.
- During the upgrade, the Kibana and Cerebro components will be rebuilt and cannot be accessed. Different Kibana versions are incompatible with each other. During the upgrade, you may fail to access Kibana due to version incompatibility. A cluster can be accessed after it is successfully upgraded.

Pre-Upgrade Check

To ensure a successful upgrade, you must check the items listed in the following table before performing an upgrade.

Table 4-2 Pre-upgrade checklist

Check Item	Check Method	Description	Normal Status
Cluster status	System check	After an upgrade task is started, the system automatically checks the cluster status. Clusters whose status is green or yellow can provide services properly and have no unallocated primary shards.	The cluster status is Available .

Check Item	Check Method	Description	Normal Status
Node quantity	System check	After an upgrade task is started, the system automatically checks the number of nodes. The total number of data nodes and cold data nodes in a cluster must be greater than or equal to 3 so that services will not be interrupted.	The total number of data nodes and cold data nodes in a cluster must be greater than or equal to 3.
Disk capacity	System check	After an upgrade task is started, the system automatically checks the disk capacity. During the upgrade, nodes are brought offline one by one and then new nodes are created. Ensure that the disk capacity of all the remaining nodes can process all data of the node that has been brought offline.	After a node is brought offline, the remaining nodes can contain all data of the cluster.
Data backup	System check	Check whether the maximum number of primary and standby shards of indexes in a cluster can be allocated to the remaining data nodes and cold data nodes. Prevent backup allocation failures after a node is brought offline during the upgrade.	Maximum number of primary and standby index shards plus 1 must be less than or equal to the total number of data nodes and cold data nodes before the upgrade.
Data backup	Manual check	Before the upgrade, back up data to prevent data loss caused by upgrade faults. When submitting an upgrade task, you can determine whether to enable the system to check for the backup of all indexes.	Check whether data has been backed up.
Resources	System check	After an upgrade task is started, the system automatically checks resources. Resources will be created during the upgrade. Ensure that resources are available.	Resources are available and sufficient.

Check Item	Check Method	Description	Normal Status
Custom plugins	System and manual check	<p>Perform this check only when custom plugins are installed in the source cluster. If a cluster has a custom plugin, upload the plugin package of the target version to a specified directory before the upgrade. During the upgrade, install the custom plugin in the new nodes. After an upgrade task is started, the system automatically checks whether the custom plugin package has been uploaded, but you need to check whether the uploaded plugin package is correct.</p> <p>NOTE An incorrect or incompatible plugin package cannot be automatically installed during the upgrade, but this does not affect the upgrade task. After the upgrade is complete, the status of the custom plugin is reset to Uploaded.</p>	The plugin package of the cluster to be upgraded has been uploaded to the plugin list.
Custom configurations	System check	During the upgrade, the system automatically synchronizes the content of the cluster configuration file elasticsearch.yml .	Clusters' custom configurations are not lost after the upgrade.
Non-standard operations	Manual check	Check whether non-standard operations are contained in the upgrade. Non-standard operations refer to manual operations that are not recorded. These operations cannot be automatically transferred during the upgrade, for example, modification of the Kibana.yml configuration file, system configuration, and route return.	Some non-standard operations are compatible. For example, the modification of a security plugin can be retained through metadata, and the modification of system configuration can be retained using images. Some non-standard operations, such as the modification of the kibana.yml file, cannot be retained, and you must back up the file in advance.

Check Item	Check Method	Description	Normal Status
Compatibility check	System and manual check	After a cross-version upgrade task is started, the system automatically checks whether the source and target versions have incompatible configurations. If a custom plugin is installed for a cluster, the version compatibility of the custom plugin needs to be manually checked.	Configurations before and after the cross-version upgrade are compatible.

Creating an Upgrade Task

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters**. On the cluster list page that is displayed, click the name of a cluster.
3. On the displayed basic cluster information page, click **Version Upgrade**.
4. On the displayed page, set upgrade parameters.

Table 4-3 Upgrade parameters

Parameter	Description
Upgrade Type	<ul style="list-style-type: none"> • Same version upgrade: Upgrade the kernel patch of the cluster. The cluster version number remains unchanged. • Cross version upgrade: Upgrade the cluster version.
Target Image	Image of the target version. When you select an image, the image name and target version details are displayed.
Agency	<p>Select an IAM agency to grant the upgrade permission to the current account.</p> <p>If no agency is available, click Create Agency to go to the IAM console and create an agency.</p> <p>NOTE The selected agency must be assigned the Tenant Administrator or VPC Administrator policy.</p>

5. After setting the parameters, click **Submit**. Determine whether to check for the backup of all indexes and click **OK** in the displayed dialog box.
6. View the upgrade task in the task list. If the task status is **Running**, you can expand the task list and click **View Progress** to view the upgrade progress.

Figure 4-1 Viewing the upgrade progress

ID	Created	Completed	Task Status	Operation
7d771	--	--	Running	Retry Terminate

Name		Task Nodes	test
ID	4545	Completed Nodes	test
Version	7.10.2	Active Nodes	test View Progress
Type	elasticsearch	Execution Times	2
Description			

If the task status is **Failed**, you can retry or terminate the task.

- Retry the task: Click **Retry** in the **Operation** column.
- Terminate the task: Click **Terminate** in the **Operation** column.

NOTICE

- Same version upgrade: If the upgrade task status is **Failed**, you can terminate the upgrade task.
- Cross version upgrade: You can stop an upgrade task only when the task status is **Failed** and no node has been upgraded.

After an upgrade task is terminated, the **Task Status** of the cluster is rolled back to the status before the upgrade, and other tasks in the cluster are not affected.

5 Importing Data to Elasticsearch

5.1 Using Logstash to Import Data to Elasticsearch

You can use Logstash to collect data and migrate collected data to Elasticsearch in CSS. This method helps you effectively obtain and manage data through Elasticsearch. Data files can be in the JSON or CSV format.

Logstash is an open-source, server-side data processing pipeline that ingests data from a multitude of sources simultaneously, transforms it, and then sends it to Elasticsearch. For details about Logstash, visit the following website: <https://www.elastic.co/guide/en/logstash/current/getting-started-with-logstash.html>

The following two scenarios are involved depending on the Logstash deployment:

- [Importing Data When Logstash Is Deployed on the External Network](#)
- [Importing Data When Logstash Is Deployed on an ECS](#)

Prerequisites

- To facilitate operations, you are advised to deploy Logstash on a host that runs the Linux operating system (OS).
- To download Logstash, visit the following website: <https://www.elastic.co/downloads/logstash-oss>

NOTE

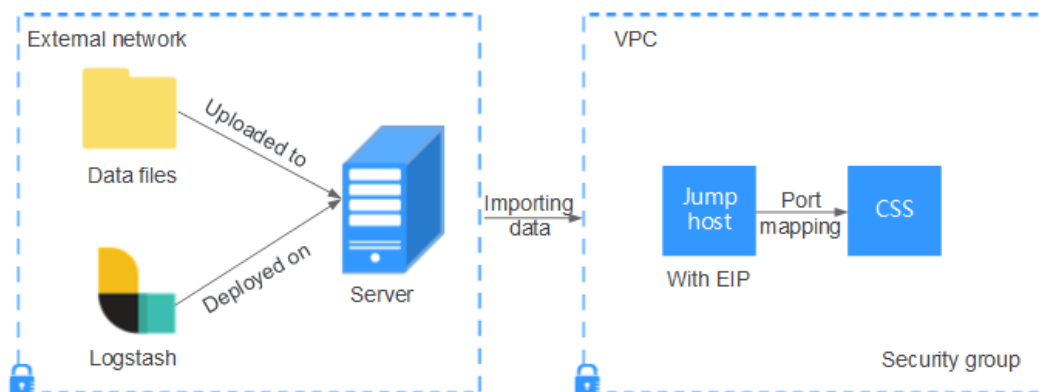
Logstash requires an OSS version same as the CSS version.

- After installing Logstash, perform the following steps to import data. For details about how to install Logstash, visit the following website: <https://www.elastic.co/guide/en/logstash/current/installing-logstash.html>
- The JDK must be installed before Logstash is installed. In Linux OS, you can run the **yum -y install java-1.8.0** command to install JDK 1.8.0. In Windows OS, you can download the required JDK version from the [official website of JDK](#), and install it by following the installation guide.
- In the [Importing Data When Logstash Is Deployed on an ECS](#) scenario, ensure that the ECS and the Elasticsearch cluster to which data is imported reside in the same VPC.

Importing Data When Logstash Is Deployed on the External Network

Figure 5-1 illustrates how data is imported when Logstash is deployed on an external network.

Figure 5-1 Importing data when Logstash is deployed on an external network



1. Create a jump host and configure it as follows:
 - The jump host is an ECS running the Linux OS and has been bound with an EIP.
 - The jump host resides in the same VPC as the CSS cluster.
 - SSH local port forwarding is configured for the jump host to forward requests from a chosen local port to port **9200** on one node of the CSS cluster.
 - Refer to [SSH documentation](#) for the local port forwarding configuration.
2. Use PuTTY to log in to the created jump host with the EIP.
3. Run the following command to perform port mapping and transfer the request sent to the port on the jump host to the target cluster:

```
ssh -g -L <Local port of the jump host:Private network address and port number of a node> -N -f root@<Private IP address of the jump host>
```

NOTE

- In the preceding command, *<Local port of the jump host>* refers to the port obtained in **1**.
- In the preceding command, *<Private network address and port number of a node>* refers to the private network address and port number of a node in the cluster. If the node is faulty, the command execution will fail. If the cluster contains multiple nodes, you can replace the value of **<private network address and port number of a node>** with the private network address and port number of any available node in the cluster. If the cluster contains only one node, restore the node and execute the command again.
- Replace *<Private IP address of the jump host>* in the preceding command with the IP address (with **Private IP**) of the created jump host in the **IP Address** column in the ECS list on the ECS management console.

For example, port **9200** on the jump host is assigned external network access permissions, the private network address and port number of the node are **192.168.0.81** and **9200**, respectively, and the private IP address of the jump host is **192.168.0.227**. You need to run the following command to perform port mapping:

```
ssh -g -L 9200:192.168.0.81:9200 -N -f root@192.168.0.227
```

- Log in to the server where Logstash is deployed and store the data files to be imported on the server.

For example, data file **access_20181029_log** needs to be imported, the file storage path is **/tmp/access_log/**, and the data file includes the following data:

 **NOTE**

Create the **access_log** folder if it does not exist.

All	Heap used for segments		18.6403	MB
All	Heap used for doc values		0.119289	MB
All	Heap used for terms		17.4095	MB
All	Heap used for norms		0.0767822	MB
All	Heap used for points		0.225246	MB
All	Heap used for stored fields		0.809448	MB
All	Segment count		101	
All	Min Throughput	index-append	66232.6	docs/s
All	Median Throughput	index-append	66735.3	docs/s
All	Max Throughput	index-append	67745.6	docs/s
All	50th percentile latency	index-append	510.261	ms

- In the server where Logstash is deployed, run the following command to create configuration file **logstash-simple.conf** in the Logstash installation directory:

```
cd /<Logstash installation directory>/
vi logstash-simple.conf
```

- Input the following content in **logstash-simple.conf**:

```
input {
  Location of data
}
filter {
  Related data processing
}
output {
  elasticsearch {
    hosts => "<IP of the jump host>:<Number of the port assigned external network access
permissions on the jump host>"
  }
}
```

- The **input** parameter indicates the data source. Set this parameter based on the actual conditions. For details about the **input** parameter and parameter usage, visit the following website: <https://www.elastic.co/guide/en/logstash/current/input-plugins.html>
- The **filter** parameter specifies the mode in which data is processed. For example, extract and process logs to convert unstructured information into structured information. For details about the **filter** parameter and parameter usage, visit the following website: <https://www.elastic.co/guide/en/logstash/current/filter-plugins.html>
- The **output** parameter indicates the destination address of the data. For details about the **output** parameter and parameter usage, visit <https://www.elastic.co/guide/en/logstash/current/output-plugins.html>. Replace *<IP address of the jump host>* with the IP address (with) of the created jump host in the **IP Address** column in the ECS list on the ECS management console. *<Number of the port assigned external network access permissions on the jump host>* is the number of the port obtained in **1**, for example, **9200**.

Consider the data files in the **/tmp/access_log/** path mentioned in **4** as an example. Assume that data import starts from data in the first row of the

data file, the filtering condition is left unspecified (indicating no data processing operations are performed), the public IP address and port number of the jump host are **192.168.0.227** and **9200**, respectively, and the name of the target index is **myindex**. Edit the configuration file as follows, and enter **:wq** to save the configuration file and exit.

```
input {
  file {
    path => "/tmp/access_log/*"
    start_position => "beginning"
  }
}
filter {
}
output {
  elasticsearch {
    hosts => "192.168.0.227:9200"
    index => "myindex"
  }
}
```

NOTE

If a license error is reported, set **ilm_enabled** to **false**.

If the cluster has the security mode enabled, you need to download a certificate first.

- a. Download a certificate on the **Basic Information** page of the cluster.
- b. Store the certificate to the server where Logstash is deployed.
- c. Modify the **logstash-simple.conf** configuration file.

Consider the data files in the **/tmp/access_log/** path mentioned in **4** as an example. Assume that data import starts from data in the first row of the data file, the filtering condition is left unspecified (indicating no data processing operations are performed), and the public IP address and port number of the jump host are **192.168.0.227** and **9200**, respectively. The name of the index for importing data is **myindex**, and the certificate is stored in **/logstash/logstash6.8/config/CloudSearchService.cer**. Edit the configuration file as follows, and enter **:wq** to save the configuration file and exit.

```
input{
  file {
    path => "/tmp/access_log/*"
    start_position => "beginning"
  }
}
filter {
}
output{
  elasticsearch{
    hosts => ["https://192.168.0.227:9200"]
    index => "myindex"
    user => "admin"
    password => "*****"
    cacert => "/logstash/logstash6.8/config/CloudSearchService.cer"
  }
}
```

NOTE

password: password for logging in to the cluster

7. Run the following command to import the data collected by Logstash to the cluster:

```
./bin/logstash -f logstash-simple.conf
```

NOTE

This command must be executed in the directory where the **logstash-simple.conf** file is stored. For example, if the **logstash-simple.conf** file is stored in **/root/logstash-7.1.1/**, go to the directory before running the command.

8. Log in to the CSS management console.
9. In the navigation pane on the left, choose **Clusters > Elasticsearch** to switch to the **Clusters** page.
10. From the cluster list, locate the row that contains the cluster to which you want to import data and click **Access Kibana** in the **Operation** column.
11. In the Kibana navigation pane on the left, choose **Dev Tools**.
12. On the **Console** page of Kibana, search for the imported data.

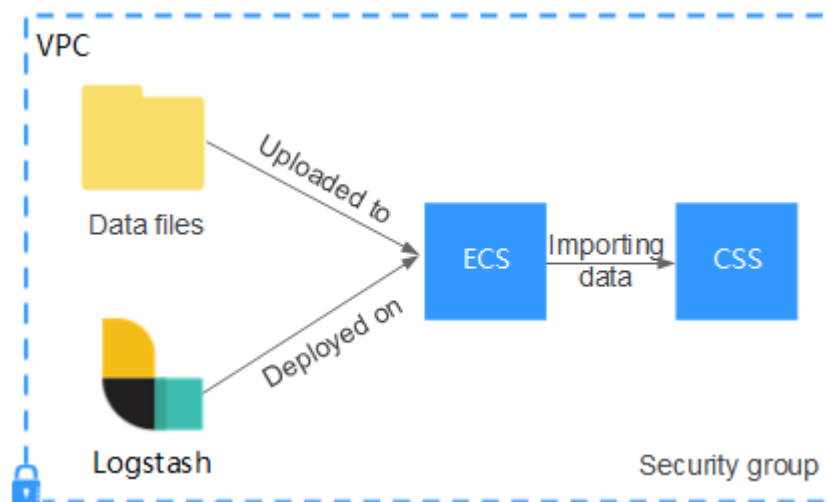
On the **Console** page of Kibana, run the following command to search for data. View the search results. If the searched data is consistent with the imported data, the data has been imported successfully.

```
GET myindex/_search
```

Importing Data When Logstash Is Deployed on an ECS

Figure 5-2 illustrates how data is imported when Logstash is deployed on an ECS that resides in the same VPC as the cluster to which data is to be imported.

Figure 5-2 Importing data when Logstash is deployed on an ECS



1. Ensure that the ECS where Logstash is deployed and the cluster to which data is to be imported reside in the same VPC, port **9200** of the ECS security group has been assigned external network access permissions, and an EIP has been bound to the ECS.
2. Use PuTTY to log in to the ECS.

For example, data file **access_20181029_log** is stored in the **/tmp/access_log/** path of the ECS, and the data file includes the following data:

All	Heap used for segments	18.6403	MB
All	Heap used for doc values	0.119289	MB
All	Heap used for terms	17.4095	MB
All	Heap used for norms	0.0767822	MB

All	Heap used for points		0.225246	MB
All	Heap used for stored fields		0.809448	MB
All	Segment count		101	
All	Min Throughput	index-append	66232.6	docs/s
All	Median Throughput	index-append	66735.3	docs/s
All	Max Throughput	index-append	67745.6	docs/s
All	50th percentile latency	index-append	510.261	ms

- Run the following command to create configuration file **logstash-simple.conf** in the Logstash installation directory:

```
cd /<Logstash installation directory>
vi logstash-simple.conf
```

Input the following content in **logstash-simple.conf**:

```
input {
  Location of data
}
filter {
  Related data processing
}
output {
  elasticsearch{
    hosts => "<Private network address and port number of the node>"
  }
}
```

- The **input** parameter indicates the data source. Set this parameter based on the actual conditions. For details about the **input** parameter and parameter usage, visit the following website: <https://www.elastic.co/guide/en/logstash/current/input-plugins.html>
- The **filter** parameter specifies the mode in which data is processed. For example, extract and process logs to convert unstructured information into structured information. For details about the **filter** parameter and parameter usage, visit the following website: <https://www.elastic.co/guide/en/logstash/current/filter-plugins.html>
- The **output** parameter indicates the destination address of the data. For details about the **output** parameter and parameter usage, visit <https://www.elastic.co/guide/en/logstash/current/output-plugins.html>. *<private network address and port number of a node>* refers to the private network address and port number of a node in the cluster.

If the cluster contains multiple nodes, you are advised to replace the value of *<Private network address and port number of a node>* with the private network addresses and port numbers of all nodes in the cluster to prevent node faults. Use commas (,) to separate the nodes' private network addresses and port numbers. The following is an example:

```
hosts => ["192.168.0.81:9200","192.168.0.24:9200"]
```

If the cluster contains only one node, the format is as follows:

```
hosts => "192.168.0.81:9200"
```

Consider the data files in the **/tmp/access_log/** path mentioned in **2** as an example. Assume that data import starts from data in the first row of the data file, the filtering condition is left unspecified (indicating no data processing operations are performed), the private network address and port number of the node in the cluster where data is to be imported are **192.168.0.81** and **9200**, respectively, and the name of the target index is **myindex**. Edit the configuration file as follows, and enter **:wq** to save the configuration file and exit.

```
input {
  file{
    path => "/tmp/access_log/*"
    start_position => "beginning"
  }
}
```

```
}  
}  
filter {  
}  
output {  
  elasticsearch {  
    hosts => "192.168.0.81:9200"  
    index => "myindex"  
  }  
}
```

4. Run the following command to import the ECS data collected by Logstash to the cluster:
`./bin/logstash -f logstash-simple.conf`
5. Log in to the CSS management console.
6. In the navigation pane on the left, choose **Clusters** > **Elasticsearch** to switch to the **Clusters** page.
7. From the cluster list, locate the row that contains the cluster to which you want to import data and click **Access Kibana** in the **Operation** column.
8. In the Kibana navigation pane on the left, choose **Dev Tools**.
9. On the **Console** page of Kibana, search for the imported data.

On the **Console** page of Kibana, run the following command to search for data. View the search results. If the searched data is consistent with the imported data, the data has been imported successfully.

```
GET myindex/_search
```

5.2 Using Kibana or APIs to Import Data to Elasticsearch

You can import data in various formats, such as JSON and CSV, to Elasticsearch in CSS by using Kibana or APIs.

Importing Data Using Kibana

Before importing data, ensure that you can use Kibana to access the cluster. The following procedure illustrates how to use the **POST** command to import data.

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters** > **Elasticsearch** to switch to the **Clusters** page.
3. Choose **Clusters** in the navigation pane. Locate the target cluster and click **Access Kibana** in the **Operation** column to log in to Kibana.
4. Click **Dev Tools** in the navigation tree on the left.
5. (Optional) On the **Console** page, run the related command to create an index for storing data and specify a custom mapping to define the data type.

If there is an available index in the cluster where you want to import data, skip this step. If there is no available index, create an index by referring to the following sample code.

For example, on the **Console** page of Kibana, run the following command to create an index named **my_store** and specify a user-defined mapping to define the data type:

Versions earlier than 7.x

```
PUT /my_store
{
  "settings": {
    "number_of_shards": 1
  },
  "mappings": {
    "products": {
      "properties": {
        "productName": {
          "type": "text"
        },
        "size": {
          "type": "keyword"
        }
      }
    }
  }
}
```

Versions later than 7.x

```
PUT /my_store
{
  "settings": {
    "number_of_shards": 1
  },
  "mappings": {
    "properties": {
      "productName": {
        "type": "text"
      },
      "size": {
        "type": "keyword"
      }
    }
  }
}
```

6. Run commands to import data. For example, run the following command to import a piece of data:

Versions earlier than 7.x

```
POST /my_store/products/_bulk
{"index":{}}
{"productName":"Latest art shirts for women in 2017 autumn","size":"L"}
```

Versions later than 7.x

```
POST /my_store/_bulk
{"index":{}}
{"productName":"Latest art shirts for women in 2017 autumn","size":"L"}
```

The command output is similar to that shown in [Figure 5-3](#). If the value of the **errors** field in the result is **false**, the data is successfully imported.

Figure 5-3 Response message

```
1 {
2   "took": 42,
3   "errors": false,
4   "items": [
5     {
6       "index": {
7         "_index": "my_store",
8         "_type": "products",
9         "_id": "AWTGbHt7BwpN-hb3LKau",
10        "_version": 1,
11        "result": "created",
12        "_shards": {
13          "total": 2,
14          "successful": 2,
15          "failed": 0
16        },
17        "created": true,
18        "status": 201
19      }
20    ]
21  }
22 }
```

Importing Data Using APIs

You can call the bulk API using the cURL command to import a JSON data file.

NOTE

You are advised to import a file smaller than 50 MB.

1. Log in to the ECS that you use to access the cluster.
2. Run the following command to import JSON data:

In the command, replace the value of *{Private network address and port number of the node}* with the private network address and port number of a node in the cluster. If the node fails to work, the command will fail to be executed. If the cluster contains multiple nodes, you can replace the value of *{Private network address and port number of the node}* with the private network address and port number of any available node in the cluster. If the cluster contains only one node, restore the node and execute the command again. **test.json** indicates the JSON file whose data is to be imported.

```
curl -X PUT "http://{Private network address and port number of the node} /_bulk" -H 'Content-Type: application/json' --data-binary @test.json
```

NOTE

The value of the **-X** parameter is a command and that of the **-H** parameter is a message header. In the preceding command, **PUT** is the value of the **-X** parameter and **'Content-Type: application/json' --data-binary @test.json** is the value of the **-H** parameter. Do not add **-k** between a parameter and its value.

Example: In this example, assume that you need to import data in the **testdata.json** file to an Elasticsearch cluster, where communication encryption is disabled and the private network address and port number of one node are

192.168.0.90 and **9200** respectively. The data in the **testdata.json** file is as follows:

Versions earlier than 7.x

```
{"index": {"_index": "my_store", "_type": "products"}}
{"productName": "Autumn new woman blouses 2019", "size": "M"}
{"index": {"_index": "my_store", "_type": "products"}}
{"productName": "Autumn new woman blouses 2019", "size": "L"}
```

Versions later than 7.x

```
{"index": {"_index": "my_store"}}
{"productName": "Autumn new woman blouse 2019", "size": "M"}
{"index": {"_index": "my_store"}}
{"productName": "Autumn new woman blouse 2019", "size": "L"}
```

Perform the following steps to import the data:

- a. Run the following command to create an index named **my_store**:

Versions earlier than 7.x

```
curl -X PUT http://192.168.0.90:9200/my_store -H 'Content-Type: application/json' -d '{
  "settings": {
    "number_of_shards": 1
  },
  "mappings": {
    "products": {
      "properties": {
        "productName": {
          "type": "text"
        },
        "size": {
          "type": "keyword"
        }
      }
    }
  }
}'
```

Versions later than 7.x

```
curl -X PUT http://192.168.0.90:9200/my_store -H 'Content-Type: application/json' -d '{
  "settings": {
    "number_of_shards": 1
  },
  "mappings": {
    "properties": {
      "productName": {
        "type": "text"
      },
      "size": {
        "type": "keyword"
      }
    }
  }
}'
```

- b. Run the following command to import the data in the **testdata.json** file:

```
curl -X PUT "http://192.168.0.90:9200/_bulk" -H 'Content-Type: application/json' --data-binary @testdata.json
```

6 Managing Elasticsearch Clusters

6.1 Cluster and Storage Capacity Statuses

On the **Dashboard** page of the CSS management console, you can view information about the status and storage capacity of existing clusters.

Table 6-1 Cluster status description

Status	Description
Available	The cluster is running properly and is providing services.
Abnormal	The cluster creation failed or the cluster is unavailable. If a cluster is in the unavailable status, you can delete the cluster or use snapshots created when the cluster is available to restore data to other clusters. However, operations such as expanding cluster capacity, accessing Kibana, creating snapshots, and restoring snapshots to the cluster are not allowed. When a cluster is in the unavailable status, data importing is not recommended to avoid data loss. You can view the cluster metrics or restart the cluster. However, the operations may fail. If the operations fail, contact technical support in a timely manner.
Processing	The cluster is being restarted, scaled, backed up, or recovered.
Creating	The cluster is being created.

Table 6-2 Cluster storage capacity status description

Status	Description
Normal	The storage capacity usage of all nodes in a cluster is less than 50%.
Warning	The storage capacity usage of any node in a cluster is from 50% to less than 80%.
Danger	The storage capacity usage of any node in a cluster is greater than or equal to 80%. You are advised to increase the storage space of the cluster to achieve normal data search or analysis.
Abnormal	The cluster storage capacity usage is unknown. For example, if the status of a cluster is Abnormal due to faults, the storage space status of the cluster will be Abnormal .

6.2 Introduction to the Cluster List

The cluster list displays all CSS clusters. If there are a large number of clusters, these clusters will be displayed on multiple pages. You can view clusters of all statuses from the cluster list.

Clusters are listed in chronological order by default in the cluster list, with the most recent cluster displayed at the top. [Table 6-3](#) shows the cluster parameters.



In the upper right corner of the cluster list, you can enter the name or ID of a cluster and click  to search for a cluster. You can also click  in the upper right corner to refresh the cluster list.

Table 6-3 Cluster list parameter description

Parameter	Description
Name/ID	Name and ID of a cluster. You can click a cluster name to switch to the Basic Information page. The cluster ID is automatically generated by the system and uniquely identifies a cluster.
Cluster Status	Status of a cluster. For details about the cluster status, see Cluster and Storage Capacity Statuses .
Task Status	Status of a task, such as cluster restart, cluster capacity expansion, cluster backup, and cluster restoration.
Version	Elasticsearch version of the cluster.
Created	Time when the cluster is created.
Enterprise Project	Enterprise project that a cluster belongs to.

Parameter	Description
Private Network Address	Private network address and port number of the cluster. You can use these parameters to access the cluster. If the cluster has multiple nodes, the private network addresses and port numbers of all nodes are displayed.
Billing Mode	Billing mode of a cluster.
Operation	Operations that can be performed on a cluster, including accessing Kibana, checking metrics, restarting a cluster, and deleting a cluster. If an operation is not allowed, the button is gray.

6.3 Index Backup and Restoration

You can back up index data in clusters. If data loss occurs or you want to retrieve data of a specified duration, you can restore the index data. Index backup is implemented by creating cluster snapshots. When creating a backup for the first time, you are advised to back up data of all indexes.

- **Managing Automatic Snapshot Creation:** Snapshots are automatically created at a specified time each day according to the rules you create. You can enable or disable the automatic snapshot creation function and set the automatic snapshot creation policy.
- **Manually creating a snapshot:** You can manually create a snapshot at any time to back up all data or data of specified indexes.
- **Restoring data:** You can use existing snapshots to restore the backup index data to a specified cluster.
- **Deleting a snapshot:** Delete snapshots you do not require and release resources.

 **NOTE**

- Before creating a snapshot, you need to perform basic configurations, including configuring the OBS bucket for storing snapshots, the snapshot backup path, and IAM agency used for security authentication.
- If there are available snapshots in the snapshot list when you configure the OBS bucket for storing cluster snapshots for the first time, you cannot change the bucket for snapshots that are subsequently created automatically or manually. Therefore, exercise caution when you configure the OBS bucket.
- If you want to change the OBS bucket where there are snapshots, do as follows: Disable the snapshot function, enable it, and specify a new OBS bucket. After you disable the snapshot function, you cannot use previously created snapshots to restore the cluster.
- If a cluster is in the **Unavailable** status, you can use the cluster snapshot function only to restore clusters and view existing snapshot information.
- During backup and restoration of a cluster, you can perform only certain operations, including scaling out, accessing Kibana, viewing metric, and deleting other snapshots of clusters. However, you cannot perform the following operations: restarting or deleting the cluster, deleting a snapshot that is in the **Creating** or **Restoring** status, and creating or restoring another snapshot. If a snapshot is being created or restored for a cluster, any automatic snapshot creation task initiated for the cluster will be canceled.
- The first snapshot of a cluster is a full snapshot, and subsequent snapshots are incremental snapshots. CSS snapshot files depend on each other.
- If you restore a CSS cluster snapshot to another cluster, indexes with the same name in the destination cluster will be overwritten. If the snapshot and the destination cluster use different shards, the indexes with the same name will not be overwritten.

Prerequisites

To use the function of creating or restoring snapshots, the account or IAM user logging in to the CSS management console must have both of the following permissions:

- **Tenant Administrator** for project **OBS** in region **Global service**
- **Elasticsearch Administrator** in the current region

Managing Automatic Snapshot Creation



1. In the CSS navigation pane on the left, click **Clusters**.
2. On the **Clusters** page that is displayed, click the name of the target cluster. In the navigation pane on the left, choose **Cluster Snapshots**.
3. On the displayed **Cluster Snapshots** page, click the icon to the right of **Cluster Snapshot** to enable the cluster snapshot function.
4. Enable the cluster snapshot function. OBS buckets and IAM agencies are automatically created to store snapshots. The automatically created OBS bucket and IAM agency are displayed on the page. You can also click  on the right of **Basic Configuration** to edit the configuration.

Table 6-4 Cluster snapshot parameter

Parameter	Description
OBS bucket	<p>Select an OBS bucket for storing snapshots from the drop-down list box. You can also click Create Bucket on the right to create an OBS bucket.</p> <p>The created or existing OBS bucket must meet the following requirements:</p> <ul style="list-style-type: none"> • Storage Class is Standard.
IAM Agency	<p>IAM agency authorized by the current account to CSS access or maintain data stored in the OBS bucket. You can also click Create IAM Agency on the right to create an IAM agency.</p> <p>The created or existing IAM agency must meet the following requirements:</p> <ul style="list-style-type: none"> • Agency Type must be Cloud service. • Set Cloud Service to . • The agency must have the permission for the project in .

5. Enable the automatic snapshot creation function. The **Configure Automatic Snapshot Creation** dialog box is displayed. If the automatic snapshot creation function is enabled, you can click  on the right of **Automatic Snapshot Creation** to modify the snapshot policy.
 - **Snapshot Name Prefix:** Enter a maximum of 32 characters starting with a lowercase letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed. A snapshot name consists of a snapshot name prefix and a timestamp, for example, **snapshot-2018022405925**.
 - **Time Zone:** indicates the time zone for the backup time. Specify based on the time zone.
 - **Index:** Enter an index name. You can select an index for backup. Use commas (,) to separate multiple indexes. Uppercase letters, spaces, and the following special characters are not allowed: "\<|>/? If you do not specify this parameter, data of all indexes in the cluster is backed up by default. You can use the asterisk (*) to back up data of certain indices. For example, if you enter **index***, then data of indices with the name prefix of **index** will be backed up.
Run the **GET /_cat/indices** command in Kibana to query the names of all indexes in the cluster.
 - : indicates the time when the backup starts automatically every day. You can specify this parameter only in hours and not minutes, for example, **00:00** or **01:00**. The value ranges from **00:00** to **23:00**. Select the backup time from the drop-down list box.
 - **Retention Period (days):** indicates the duration when snapshots are retained in the OBS bucket, in days. The value ranges from **1** to **90**. You can specify this parameter as required. The system automatically deletes snapshots that are retained over the specified retention period on the half hour.

6. Click **OK** to save the snapshot policy.
Snapshots that are automatically created according to the snapshot policy are displayed in the snapshot list, along with manually created snapshots. You can distinguish them by the **Snapshot Type** setting. In the upper right corner of the snapshot list, enter the keyword of the snapshot name or snapshot ID to search for the desired snapshots.
7. (Optional) Disable the automatic snapshot creation function.
After you disable the automatic snapshot creation function, the system stops automatic creation of snapshots. If the system is creating a snapshot based on the automatic snapshot creation policy and the snapshot is not yet displayed in the snapshot list, you cannot disable the automatic snapshot creation function. In this case, if you click the button next to **Automatic Snapshot Creation**, a message is displayed, indicating that you cannot disable the function. You are advised to disable the function after the system completes automatic creation of the snapshot, and the created snapshot is displayed in the snapshot list.
When disabling the automatic snapshot creation function, you can choose whether to delete the snapshots that have been automatically created by selecting in the displayed dialog box. By default, automatically created snapshots are not deleted.
 - If you do not select **Delete automated snapshots**, automatically created snapshots are not deleted when you disable the automatic snapshot creation function. You can manually delete them later. For details, see [Deleting a Snapshot](#). If you do not manually delete the automatically created snapshots and enable the automatic snapshot creation function again, then all snapshots with **Snapshot Type** set to **Automated** in the snapshot list of the cluster can only be automatically deleted by the system. Specifically, the system automatically deletes snapshots based on the snapshot policy configured when you enable the automatic snapshot creation function again. For example, if you set **Retention Period (days)** to **10**, the system will automatically delete the snapshots that have been retained for more than 10 days.
 - If you select **Delete automated snapshots**, all snapshots with **Snapshot Type** set to **Automated** in the snapshot list will be deleted when you disable the automatic snapshot creation function.

Manually Creating a Snapshot


1. In the CSS navigation pane on the left, click **Clusters**.
2. On the **Clusters** page that is displayed, click the name of the target cluster. In the navigation pane on the left, choose **Cluster Snapshots**.
3. On the displayed **Cluster Snapshots** page, click the icon to the right of **Cluster Snapshot** to enable the cluster snapshot function.
4. Enable the cluster snapshot function. OBS buckets and IAM agencies are automatically created to store snapshots. The automatically created OBS bucket and IAM agency are displayed on the page. You can also click  on the right of **Basic Configuration** to edit the configuration.

Table 6-5 Cluster snapshot parameter

Parameter	Description
OBS bucket	<p>Select an OBS bucket for storing snapshots from the drop-down list box. You can also click Create Bucket on the right to create an OBS bucket.</p> <p>The created or existing OBS bucket must meet the following requirements:</p> <ul style="list-style-type: none"> • Storage Class is Standard.
IAM Agency	<p>IAM agency authorized by the current account to CSS access or maintain data stored in the OBS bucket. You can also click Create IAM Agency on the right to create an IAM agency.</p> <p>The created or existing IAM agency must meet the following requirements:</p> <ul style="list-style-type: none"> • Agency Type must be Cloud service. • Set Cloud Service to . • The agency must have the permission for the project in .

5. After basic configurations are completed, click **Create**.
 - **Name** indicates the name of the manually created snapshot, which can contain 4 to 64 characters and must start with a lowercase letter. Only lowercase letters, digits, hyphens (-), and underscores (_) are allowed. For snapshots you create manually, you can specify the snapshot name. The system will not automatically add the time information to the snapshot name.
 - **Index**: Enter an index name. You can select an index for backup. Use commas (,) to separate multiple indexes. Uppercase letters, spaces, and the following special characters are not allowed: "\<|>/? If you do not specify this parameter, data of all indexes in the cluster is backed up by default. You can use the asterisk (*) to back up data of certain indices. For example, if you enter **index***, then data of indices with the name prefix of **index** will be backed up.
Run the **GET /_cat/indices** command in Kibana to query the names of all indexes in the cluster.
 - **Description**: indicates the description of the created snapshot. The value contains 0 to 256 characters, and certain special characters (<>) are not allowed.

6. Click **OK**.

After the snapshot is created, it will be displayed in the snapshot list. The status **Available** indicates that the snapshot is created successfully. All automatically and manually created snapshots are displayed in the snapshot list. You can distinguish them by the **Snapshot Type** setting. In the upper right corner of the snapshot list, enter the keyword of the snapshot name or snapshot ID to search for the desired snapshots.

Restoring Data

You can use snapshots whose **Snapshot Status** is **Available** to restore cluster data. The stored snapshot data can be restored to other clusters.

Restoring data will overwrite current data in clusters. Therefore, exercise caution when restoring data.

1. In the **Snapshots** area, locate the row that contains the snapshot you want to restore and click **Restore** in the **Operation** column.
2. On the **Restore** page, set restoration parameters.

Index: Enter the name of the index you want to restore. If you do not specify any index name, data of all indexes will be restored. The value can contain 0 to 1,024 characters. Uppercase letters, spaces, and certain special characters (including "<|>/?") are not allowed.

Rename Pattern: Enter a regular expression. Indexes that match the regular expression are restored. The default value **index_(.+)** indicates restoring data of all indices. The value contains 0 to 1,024 characters. Uppercase letters, spaces, and certain special characters (including "<|>/?") are not allowed.

Rename Replacement: Enter the index renaming rule. The default value **restored_index_\$1** indicates that **restored_** is added in front of the names of all restored indexes. The value can contain 0 to 1,024 characters. Uppercase letters, spaces, and certain special characters (including "<|>/?") are not allowed. You can set **Rename Replacement** only if you have specified **Rename Pattern**.

Cluster: Select the cluster that you want to restore. You can select the current cluster or others. However, you can only restore the snapshot to clusters whose status is **Available**. If the status of the current cluster is **Unavailable**, you cannot restore the snapshot to the current cluster. If you select another cluster and two or more indexes in the cluster have the same name, data of all indices with the same name as the name you specify will be overwritten. Therefore, exercise caution when you set the parameters.

3. Click **OK**. If restoration succeeds, **Task Status** of the snapshot in the snapshot list will change to **Restoration succeeded**, and the index data is generated again according to the snapshot information.

Deleting a Snapshot

If you no longer need a snapshot, delete it to release storage resources. If the automatic snapshot creation function is enabled, snapshots that are automatically created cannot be deleted manually, and the system automatically deletes these snapshots on the half hour after the time specified by **Retention Period (days)**. If you disable the automatic snapshot creation function while retaining the automated snapshots, then you can manually delete them later. If you do not manually delete the automatically created snapshots and enable the automatic snapshot creation function again, then all snapshots with **Snapshot Type** set to **Automated** in the snapshot list of the cluster can only be automatically deleted by the system.

NOTE

After a snapshot is deleted, its data cannot be restored. Exercise caution when deleting a snapshot.

1. In the snapshot list, locate the snapshot that you want to delete.
2. Click **Delete** in the **Operation** column. In the dialog box that is displayed, confirm the snapshot information and click **OK**.

6.4 Binding an Enterprise Project

You can create enterprise projects based on your organizational structure. Then you can manage resources across different regions by enterprise project, add users and user groups to enterprise projects, and grant different permissions to the users and user groups. This section describes how to bind a CSS cluster to an enterprise project and how to modify an enterprise project.

Prerequisites

Before binding an enterprise project, you have created an enterprise project.

Binding an Enterprise Project

When creating a cluster, you can bind an existing enterprise project to the cluster, or click **View Enterprise Project** to go to the enterprise project management console and create a new project or view existing projects.

Modifying an Enterprise Project

For a cluster that has been created, you can modify its enterprise project based on the site requirements.

1. Log in to the CSS management console.
2. In the navigation tree on the left, choose **Clusters > Elasticsearch**.
3. In the cluster list on the displayed page, click the target cluster name to switch to the **Cluster Information** page.
4. On the **Cluster Information** page, click the enterprise project name on the right of **Enterprise Project**. The project management page is displayed.
5. On the **Resources** tab page, select the region of the current cluster, and select **CSS for Service**. In this case, the corresponding CSS cluster is displayed in the resource list.
6. Select the cluster whose enterprise project you want to modify and click **Remove**.
7. On the **Remove Resource** page, specify **Mode** and select **Destination Enterprise Project**, and click **OK**.
8. After the resource is removed, you can view the modified enterprise project information on the **Clusters** page.

6.5 Restarting a Cluster

If a cluster becomes faulty, you can restart it to check if it can run normally.

Prerequisites

- The target cluster is not frozen and has no task in progress.
- If a cluster is available, ensure that it has stopped processing service requests (such as importing data and searching for data). Otherwise, data may be lost when the cluster is restarted. You are advised to perform this operation during off-peak hours.

Context

CSS supports quick restart and rolling restart.

Quick Restart

- All clusters support this function.
- If you select a node type for quick restart, all nodes of the selected type will be restarted together.
- If you select a node name for quick restart, only the specified node will be restarted.
- The cluster is unavailable during quick restart.

Rolling Restart

- Rolling restart is supported only when a cluster has at least three nodes (including master nodes, client nodes, and cold data nodes).
- Rolling restart can be performed only by specifying node types. If you select a node type for rolling restart, the nodes of the selected type will be restarted in sequence.
- During the rolling restart, only the nodes that are being restarted are unavailable and other nodes can run normally.
- When the data volume is large, rolling restart will take a long time.

Quick Restart

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Restart** in the **Operation** column.
3. On the **Restart Cluster** page, select **Quick Restart**.
You can quick restart nodes by **Node type** or **Node name**. If you select **Node type**, then you can select multiple node types and perform quick restart at the time. If you select **Node name**, you can perform quick restart only on one node at a time.
4. Refresh the page and check the cluster status. During the restart, the cluster status is **Processing**, and the task status is **Restarting**. If the cluster status changes to **Available**, the cluster has been restarted successfully.

Rolling Restart

1. Log in to the CSS management console.
2. In the navigation pane on the left, choose **Clusters > Elasticsearch**. On the displayed **Clusters** page, locate the target cluster and choose **More > Restart** in the **Operation** column.

3. On the **Restart Cluster** page, select **Rolling Restart**.
You can perform rolling restart by **Node type**. Select specific node types for restart.
4. Refresh the page and check the cluster status. During the restart, the cluster status is **Processing**, and the task status is **Restarting**. If the cluster status changes to **Available**, the cluster has been restarted successfully.

6.6 Migrating Cluster Data

You can migrate data from one cluster to another. In certain scenarios, for example, if you cannot get sufficient capacity by changing the specifications of the current cluster, you can create a cluster of higher specifications and migrate all data of the current cluster to the new one. Alternatively, you can merge indexes in two clusters to one cluster to meet your business needs. CSS enables you to migrate cluster data by using the index backup and restoration function, specifically, by restoring the snapshot of a cluster to the target cluster.

In this section, assume that we need to migrate the data of cluster **Es-1** to cluster **Es-2**. Cluster **Es-2** runs a version later than that of cluster **Es-1** and the number of nodes in cluster **Es-2** is greater than half of that in cluster **Es-1**.

Prerequisites

- The source and target clusters are in the same region.
- The version of the target cluster is the same as or later than that of the source cluster.
- The number of nodes in the target cluster must be greater than half of the number of nodes in the source cluster.

Suggestions

- The number of nodes in the target cluster should be no less than the number of replicas in the source cluster.
- The CPU, memory, and disk configurations of the target cluster should be no less than those of the source cluster. This will minimize service loss after migration.

Migration Duration

The number of nodes or index shards in the source and destination clusters determines how long the data migration will take. Data migration consists of two phases: data backup and restoration. The backup duration is determined by the source cluster and the restoration duration is determined by the destination cluster. The formula for calculating the total migration duration is as follows:

- If the number of index shards is greater than the number of nodes:

$$\text{Total duration (s)} = (800 \text{ GB}/40 \text{ MB}/\text{Number of source cluster nodes} + 800 \text{ GB}/40 \text{ MB}/\text{Number of destination cluster nodes}) \times \text{Number of indexes}$$

- If the number of index shards is smaller than the number of nodes:

$$\text{Total duration (s)} = (800 \text{ GB}/40 \text{ MB}/\text{Number of shards of the source cluster index} + 800 \text{ GB}/40 \text{ MB}/\text{Number of shards of the destination cluster index}) \times \text{Number of indexes}$$

 NOTE

The migration duration estimated using the formula is the minimal duration possible (if each node transmits data at the fastest speed, 40 MB/s). The actual duration also depends on factors such as the network and resources condition.

Procedure

1. On the **Clusters** page, click **Es-1**.
2. In the navigation pane on the left, choose **Cluster Snapshots**. Enable the cluster snapshot and set basic configuration. For details, see [Manually Creating a Snapshot](#).
3. Click **Create** to manually create a snapshot. In the displayed dialog box, enter the snapshot name and click **OK**.
4. In the snapshot list, locate the target snapshot and click **Restore** in the **Operation** column to restore data to cluster **Es-2**.
 - In the text box next to **Index**, enter *****, indicating that you want to restore data of all indexes in cluster **Es-1**.
 - From the **Cluster** drop-down list, select **Es-2**.Click **OK**.
5. After restoration is complete, data in cluster **Es-1** will be migrated to cluster **Es-2**.

6.7 Deleting a Cluster

You can delete clusters that you no longer need.

 NOTE

- If you delete a cluster, the cluster service data will be cleared. Exercise caution when performing this operation.
- The snapshots of a cluster stored in OBS are not deleted with the cluster.

Procedure

1. Log in to the CSS management console.
2. In the navigation tree on the left, choose **Clusters > Elasticsearch**.
3. Locate the target cluster and click **More > Delete** in the **Operation** column.
4. In the displayed dialog box, enter the name of the cluster to be deleted and click **OK**.

6.8 Managing Tags

Tags are cluster identifiers. Adding tags to clusters can help you identify and manage your cluster resources.

You can add tags to a cluster when creating the cluster or add them on the details page of the created cluster.

Managing Tags of a New Cluster

1. Log in to the CSS management console.
2. Click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.
3. On the **Create Cluster** page, set **Advanced Settings** to **Custom**. Add tags for a cluster.

You can select a predefined tag and set **Tag value** for the tag. You can click **View Predefined Tag** to switch to the TMS management console and view existing tags.

You can also create new tags by specifying **Tag key** and **Tag value**.

You can add a maximum of 10 tags for a CSS cluster. If the entered tag is incorrect, you can click **Delete** on the right of the tag to delete the tag.

Table 6-6 Naming rules for a tag key and value

Parameter	Description
Tag key	<p>Must be unique in a cluster.</p> <p>Can contain a maximum of 36 characters.</p> <p>Can only consist of digits, letters, hyphens (-), and underscores (_).</p>
Tag value	<p>Can contain a maximum of 43 characters.</p> <p>Can only consist of digits, letters, hyphens (-), and underscores (_).</p> <p>Cannot be left blank.</p>

Managing Tags of Existing Clusters

You can modify, delete, or add tags for a cluster.

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

The **Basic Information** page is displayed.

3. In the navigation pane on the left, choose the **Tags** tab. You can add, modify, or delete tags.

- View

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

- Add

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

- Modify

You can only change the value of an existing tag.

In the **Operation** column of a tag, click **Edit**. In the displayed **Edit Tag** page, enter a new tag value and click **OK**.

- Delete

In the **Operation** column of the tag, click **Delete**. After confirmation, click **Yes** on the displayed **Delete Tag** page.

Searching for Clusters by Tag

1. Log in to the CSS management console.
2. On the **Clusters** page, click **Search by Tag** in the upper right corner of the cluster list.
3. Select or enter the tag key and tag value you want to search for, and click **Add** to add the tag to the search text box.

You can select a tag key or tag value from their drop-down lists. The system returns a list of clusters that exactly match the tag key or tag value. If you enter multiple tags, the cluster that meets requirements of all the tags will be filtered.

You can add a maximum of 10 tags at one time.

4. Click **Search**.

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

6.9 Public Network Access

You can access a security cluster (clusters in version 6.5.4 or later support the security mode) that has the HTTPS access enabled through the public IP address provided by the system.

NOTE

If public network access is enabled for CSS, then EIP and bandwidth resources will be used and billed.

Configuring Public Network Access

1. Log in to the CSS management console.
2. On the **Create Cluster** page, enable **Security Mode**. Set the administrator password and enable HTTPS access.
3. Select **Automatically assign** for **Public IP Address** and set related parameters.

Table 6-7 Public network access parameters

Parameter	Description
Bandwidth	Bandwidth for accessing Kibana with the public IP address

Parameter	Description
Access Control	If you disable this function, all IP addresses can access the cluster through the public IP address. If you enable access control, only IP addresses in the whitelist can access the cluster through the public IP address.
Whitelist	IP address or IP address range allowed to access a cluster. Use commas (,) to separate multiple addresses. This parameter can be configured only when Access Control is enabled.

Managing Public Network Access

You can configure, modify, view the public network access of, or disassociate the public IP address from a cluster.

1. Log in to the CSS management console.
2. On the **Clusters** page, click the name of the target cluster. On the **Basic Information** page that is displayed, manage the public network access configurations.
 - Configuring public network access

If you did not configure the public network access during cluster creation, you can configure it on the **Basic Information** page after configuring the cluster.

Click **Associate** next to **Public IP Address**, set the access bandwidth, and click **OK**.

If the association fails, wait for several minutes and try again.
 - Modifying public network access

For a cluster for which you have configured public network access, you can click **Edit** next to **Bandwidth** to modify the bandwidth, or you can click **Set** next to **Access Control** to set the access control function and the whitelist for access.
 - Viewing public network access

On the **Basic Information** page, you can view the public IP address associated with the current cluster.
 - Disassociating a public IP address from a cluster

To disassociate the public IP address, click **Disassociate** next to **Public IP Address**.

Accessing a Cluster Through the Public IP Address

After configuring the public IP address, you can use it to access the cluster.

For example, run the following cURL commands to view the index information in the cluster. In this example, the public access IP address of one node in the cluster is **10.62.179.32** and the port number is **9200**.

- If the cluster you access does not have the security mode enabled, run the following command:

```
curl 'http://10.62.179.32:9200/_cat/indices'
```

- If the cluster you access has the security mode enabled, access the cluster using HTTPS and add the username, password and **-u** to the cURL command.

```
curl -u username:password -k 'https://10.62.179.32:9200/_cat/indices'
```

6.10 Managing Logs

CSS provides log backup and search functions to help you locate faults. You can back up cluster logs to OBS buckets and download required log files to analyze and locate faults.

Enabling Log Management

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, click the name of the target cluster. The cluster information page is displayed.
3. Click the **Logs** tab and toggle on the **Log Management** switch.
4. In the **Edit Log Backup Configuration** dialog box, set the parameters.

In the displayed dialog box, **OBS Bucket**, **Backup Path**, and **IAM Agency** are automatically created for log backup. You can change the default value by referring to [Table 6-8](#).

If the **Log Management** function has been enabled for the cluster, you can

click  on the right of **Log Backup Configuration** and modify the configuration in the displayed **Edit Log Backup Configuration** dialog box. For details, see [Table 6-8](#).

Table 6-8 Parameters for configuring log backup

Parameter	Description	Remarks
OBS Bucket	Select an OBS bucket from the drop-down list for storing logs. You can also click Create Bucket on the right to create an OBS bucket.	The OBS bucket and the cluster must be in the same region. NOTE To let an IAM user access an OBS bucket, you need to grant the GetBucketStoragePolicy, GetBucketLocation, ListBucket, and ListAllMyBuckets permissions to the user.


Parameter	Description	Remarks
Backup Path	Storage path of logs in the OBS bucket	<p>The backup path configuration rules are as follows:</p> <ul style="list-style-type: none"> • The backup path cannot contain the following characters: \:*?"<> • The backup path cannot start with a slash (/). • The backup path cannot start or end with a period (.) • The total length of the backup path cannot exceed 1,023 characters.
IAM Agency	IAM agency authorized by the current account for CSS to access or maintain data stored in the OBS bucket. You can also click Create IAM Agency on the right to create an IAM agency.	<p>The IAM agency must meet the following requirements:</p> <ul style="list-style-type: none"> • Agency Type must be Cloud service. • Set Cloud Service to Elasticsearch or CSS. • Mandatory policies: Tenant Administrator

5. Back up logs.

- Automatically backing up logs

Click the icon on the right of **Auto Backup** to enable the auto backup function.

After the automatic backup function is enabled, set the backup start time in the **Configure Auto Backup** dialog box. When the scheduled time arrives, the system will back up logs automatically.

After the **Automatic Snapshot Creation** function is enabled, you can click  on the right of the parameter to change the backup start time.

- Manually backing up logs

On the **Log Backup** tab page, click **Back Up**. On the displayed page, click **Yes** to start backup.


If **Task Status** in the log backup list is **Successful**, the backup is successful.

 **NOTE**

All logs in the cluster are copied to a specified OBS path. You can view or download log files from the path of the OBS bucket.

6. Search for logs.

On the **Log Search** page, select the target node, log type, and log level, and

click  . The search results are displayed.

When you search for logs, the latest 10,000 logs are matched. A maximum of 100 logs are displayed.

Viewing Logs

After backing up logs, you can click **Backup Path** to go to the OBS console and view the logs.

Backed up logs mainly include deprecation logs, run logs, index slow logs, and search slow logs. [Table 6-9](#) lists the storage types of the OBS bucket.

Table 6-9 Log types

Log Name	Description
clustername_deprecation.log	Deprecation log
clustername_index_indexing_slowlog.log	Search slow log
clustername_index_search_slowlog.log	Index slow log
clustername.log	Elasticsearch run log
clustername_access.log	Access log
clustername_audit.log	Audit log

6.11 Managing Plugins

CSS clusters have default plugins. You can view the default plugin information on the console or Kibana.

Viewing Plugins on the Console

1. Log in to the CSS management console.
2. In the navigation pane, choose **Clusters**. Click the target cluster name and go to the **Basic Information** page of the cluster.
3. Click the **Plugins** tab.
4. On the **Default** page, view default plugins supported by the current version.

Viewing Plugins on the Kibana

1. Log in to the CSS management console.
2. In the navigation pane, choose **Clusters**. Locate the target cluster and click **Access Kibana** in the **Operation** column to log in to Kibana.
3. Go to **Dev Tools** and run the following command to view the cluster plugin information:

```
GET _cat/plugins?v
```

The following is an example of the response body:

name	component	version
css-test-ess-esn-1-1	analysis-dynamic-synonym	7.6.2-hw-ei-css-v1.0.1
css-test-ess-esn-1-1	analysis-icu	7.6.2-hw-ei-css-v1.1.6
css-test-ess-esn-1-1	analysis-ik	7.6.2-hw-ei-css-v1.0.1
.....		

name indicates the cluster node name, **component** indicates the plugin name, and **version** indicates the plugin version.

6.12 Hot and Cold Data Storage

CSS provides you with cold data nodes. You can store data that requires query response in seconds on high-performance nodes and store data that requires query response in minutes on cold data nodes with large capacity and low specifications.

NOTE

- When creating a cluster, you need to configure nodes as data nodes. When you enable the cold data node function, data nodes become hot nodes.
- You can enable the cold data node, master node, and client node functions at the same time.
- You can increase nodes and expand storage capacity of cold data nodes. The maximum storage capacity is determined by the node specifications. Local disks do not support storage capacity expansion.

Hot and Cold Data Node Switchover

After you enable the cold data node function, the cold data node is labeled with **cold**. In addition, data nodes are labeled with **hot** and become hot nodes. You can specify indices to distribute data to cold or hot nodes.

You can configure a template to store indices on the specified cold or hot node.

The following figure shows this process. Log in to the **Kibana Console** page of the cluster, modify the template by configuring the index starting with **myindex**, and store the indices on the cold node. In this case, the **myindex*** data is stored on the cold data node by modifying the template.

Run the following command to create a template:

```
PUT _template/test
{
  "order": 1,
  "index_patterns": "myindex*",
  "settings": {
    "refresh_interval": "30s",
    "number_of_shards": "3",
    "number_of_replicas": "1",
    "routing.allocation.require.box_type": "cold"
  }
}
```

You can perform operations on the created index.

```
PUT myindex/_settings
{
  "index.routing.allocation.require.box_type": "cold"
}
```

You can cancel the configurations of hot and cold data nodes.

```
PUT myindex/_settings
{
  "index.routing.allocation.require.box_type": null
}
```

6.13 Configuring Parameters

You can modify the `elasticsearch.yml` file.

Modifying Parameter Configurations

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, click the name of the target cluster. The cluster information page is displayed.
3. Click **Parameter Configurations** and click **Edit** to modify module parameters as required.

Table 6-10 Module parameters

Module Name	Parameter	Description
Cross-domain Access	http.cors.allow-credentials	Whether to return the Access-Control-Allow-Credentials of the header during cross-domain access Value: true or false Default value: false
	http.cors.allow-origin	Origin IP address allowed for cross-domain access, for example, 122.122.122.122:9200
	http.cors.max-age	Cache duration of the browser. The cache is automatically cleared after the time range you specify. Unit: s Default value: 1,728,000
	http.cors.allow-headers	Headers allowed for cross-domain access, including X-Requested-With , Content-Type , and Content-Length . Use commas (,) and spaces to separate headers.
	http.cors.enabled	Whether to allow cross-domain access Value: true or false Default value: false
	http.cors.allow-methods	Methods allowed for cross-domain access, including OPTIONS , HEAD , GET , POST , PUT , and DELETE . Use commas (,) and spaces to separate methods.

Module Name	Parameter	Description
Reindexing	reindex.remote.whitelist	Configured for migrating data from the current cluster to the target cluster through the reindex API. The example value is 122.122.122.122:9200.
Custom Cache	indices.queries.cache.size	Cache size in the query phase Value range: 1 to 100 Unit: % Default value: 10%
Queue Size in a Thread Pool	thread_pool.bulk.queue_size	Queue size in the bulk thread pool. The value is an integer. Default value: 200
	thread_pool.write.queue_size	Queue size in the write thread pool. The value is an integer. Default value: 200
	thread_pool.force_merge.size	Queue size in the force merge thread pool. The value is an integer. Default value: 1
Customize	You can add parameters based on your needs.	Customized parameters NOTE <ul style="list-style-type: none"> Enter multiple values in the format as [value1, value2, value3...]. Separate values by commas (,) and spaces. Colons (:) are not allowed.

- After the modification is complete, click **Submit**. In the displayed **Submit Configuration** dialog box, select the box indicating "I understand that the modification will take effect after the cluster is restarted." and click **Yes**.
If the **Status** is **Succeeded** in the parameter modification list, the modification has been saved. Up to 20 modification records can be displayed.
- Return to the cluster list and choose **More > Restart** in the **Operation** column to restart the cluster and make the modification take effect.
 - You need to restart the cluster after modification, or **Configuration unupdated** will be displayed in the **Task Status** column on the **Clusters** page.
 - If you restart the cluster after the modification, and **Task Status** displays **Configuration error**, the parameter configuration file fails to be modified.

6.14 VPC Endpoint Service

If the VPC endpoint service is enabled, you can use a private domain name or node IP address generated by the endpoint to access the cluster. When the VPC

endpoint service is enabled, a VPC endpoint will be created by default. You can select **Private Domain Name Creation** as required.

⚠ CAUTION

The public IP address access and VPC endpoint service share a load balancer. If you have configured a public access whitelist, public and private IP addresses that access the cluster through VPCEP are restricted because the public IP address access shares the load balancer with the VPC endpoint service. In this case, you need to add IP address **198.19.128.0/17** to the public access whitelist to allow traffic through VPCEP.

Enabling the VPC Endpoint Service

1. Log in to the CSS management console.
2. Click **Create Cluster** in the upper right corner.
3. On the **Create Cluster** page, set **Advanced Settings** to **Custom**. Enable the VPC endpoint service.
 - **Private Domain Name Creation:** If you enable this function, the system automatically creates a private domain name for you, which you can use to access the cluster.
 - **VPC Endpoint Service Whitelist:** You can add an authorized account ID to the VPC endpoint service whitelist. Then you can access the cluster using the domain name or the node IP address.
 - You can click **Add** to add multiple accounts.
 - Click **Delete** in the **Operation** column to delete the accounts that are not allowed to access the cluster.

📖 NOTE

- If the authorized account ID is set to *, all users are allowed to access the cluster.
- You can view authorized account IDs on the **My Credentials** page.

Managing VPC Endpoint Service

You can enable the VPC endpoint service while creating a cluster, and also enable it by performing the following steps after cluster creation.


1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, click the name of the target cluster.
3. Click the **VPC Endpoint Service** tab, and turn on the button next to **VPC Endpoint Service**.

In the displayed dialog box, you can determine whether to enable the private domain name. Click **Yes** to enable the VPC endpoint service.

 **NOTE**

- If the VPC endpoint service is enabled, you can use a private domain name or node IP address generated by the VPC endpoint to access the cluster. For details, see [Accessing the Cluster Using the Private Domain Name or Node IP Address](#).
 - If you disable the VPC endpoint service, none of the users can access the cluster using the private domain name.
4. (Optional) Click **Modify** next to **VPC Endpoint Service Whitelist** to update the existing whitelist.
 5. Manage VPC endpoints.
The **VPC Endpoint Service** page displays all VPC endpoints connected to the current VPC endpoint service.

Figure 6-1 Managing VPC endpoints



VPC Endpoint ID	Status	Max. Connections	Owner	Created	Operation
2522a05c-8076-437f-aeca-988...	● Accepted	3000	c4b2d06aa08b4b7ba8439b28be65410f	Jul 08, 2022 15:15:14 GMT+08:00	Accept Reject Modification

- Click **Accept** or **Reject** in the **Operation** column to change the node status. If you reject the connection with a VPC endpoint, you cannot access the cluster through the private domain name generated by that VPC endpoint.
- Click **Modify** in the **Operation** column to change the maximum number of connections of the current node.

Accessing the Cluster Using the Private Domain Name or Node IP Address

1. Obtain the private domain name or node IP address.
Log in to the CSS console, click the target cluster name and go to the **Cluster Information** page. Click the **VPC Endpoint Service** tab and view the private domain name.
2. Run the cURL command to execute the API or call the API by using a program before accessing the cluster. For details about Elasticsearch operations and APIs, see the [Elasticsearch Reference](#).

The ECS must meet the following requirements:

- Sufficient disk space is allocated for the ECS.
- The ECS and the cluster must be in the same VPC. After enabling the VPC endpoint service, you can access the cluster from the ECS even when the cluster is not in the same VPC as the ECS.
- The security group of the ECS must be the same as that of the cluster.
If this requirement is not met, modify the ECS security group or configure the inbound and outbound rules of the ECS security group to allow the ECS security group to be accessed by all security groups of the cluster. For details, see [Configuring Security Group Rules](#).
- Configure security group rule settings of the target CSS cluster. Set **Protocol** to **TCP** and **Port Range** to **9200** or a port range including port **9200** for both the outbound and inbound directions.

For example, run the following cURL command to view the index information in the cluster. In this example, the private network address is **vpcep-7439f7f6-2c66-47d4-b5f3-790db4204b8d.region01.huaweicloud.com** and port **9200** is used to access the cluster.

- If the cluster you access does not have the security mode enabled, run the following command:
- If the cluster you access has the security mode enabled, access the cluster using HTTPS and add the username, password and **-u** to the cURL command.

6.15 Kibana Public Access

For CSS clusters that have security mode enabled, you can enable Kibana public access. After the configuration is complete, an IP address will be provided to access Kibana of this cluster over the Internet.

You can configure Kibana public access during cluster creation, or after a cluster in security mode is created.

NOTE

- You can enable **Security Mode** for clusters of version 6.5.4 and later versions.
- Kibana public access cannot be configured for clusters created in security mode before this function was rolled out (before June 2020).
- The whitelist for Kibana public network access depends on the ELB whitelist. After you updated the whitelist, the new settings take effect immediately for new connections. For existing persistent connections using the IP addresses that have been removed from the whitelist, the new settings take effect about 1 minute after these connections are stopped.

Configuring Kibana Public Access When Creating a Cluster

1. Log in to the CSS management console.
2. Click **Create Cluster** in the upper right corner. The **Create Cluster** page is displayed.
3. On the **Create Cluster** page, enable **Security Mode**.
4. Set **Advanced Settings** to **Custom**, enable **Kibana Public Access**, and set parameters.

Table 6-11 Kibana public access parameters

Parameter	Description
Bandwidth	Bandwidth for accessing Kibana with the public IP address Value range: 1 to 100 Unit: Mbit/s

Parameter	Description
Access Control	If you disable this function, all IP addresses can access Kibana through the public IP address. If you enable this function, only IP addresses or IP address in the whitelist can access Kibana through the public IP address.
Whitelist	IP address or IP address range allowed to access a cluster. Use commas (,) to separate multiple addresses. This parameter can be configured only when Access Control is enabled. You are advised to enable this function.

After the cluster is created, click the cluster name to go to the **Basic Information** page. On the **Kibana Public Access** page, you can view the Kibana public IP address.

Configuring Kibana Public Access for an Existing Cluster

You can enable, disable, modify, and view Kibana public access for an existing cluster that has security mode enabled.

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, click the name of the target cluster.
3. Click the **Kibana Public Access** tab. Turn on the **Kibana Public Access** switch to enable the Kibana public access function.
4. On the displayed page, set parameters.

Table 6-12 Kibana public access parameters

Parameter	Description
Bandwidth	Bandwidth for accessing Kibana with the public IP address Value range: 1 to 100 Unit: Mbit/s
Access Control	If you disable this function, all IP addresses can access Kibana through the public IP address. If you enable this function, only IP addresses or IP address in the whitelist can access Kibana through the public IP address.
Whitelist	IP address or IP address range allowed to access a cluster. Use commas (,) to separate multiple addresses. This parameter can be configured only when Access Control is enabled. You are advised to enable this function.

5. After you set the parameters, click **OK**.

Modifying Kibana Public Access

For clusters configured Kibana public access, you can modify its bandwidth and access control or disable this function.

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, click the name of the target cluster.
3. Click the **Kibana Public Access** tab to modify the Kibana public access configuration.
 - Modifying bandwidth
Click **Modify** on the right of **Bandwidth**. On the **Modify Bandwidth** page, modify the bandwidth and click **OK**.
 - Modifying access control
Click **Modify** on the right of **Access Control**. On the **Modify Access Control** page, set **Access Control** and **Whitelist**, and click **OK**.
 - Disabling Kibana public access
Toggle off the **Kibana Public Access** switch.

Accessing Kibana with the Public IP Address

After configuring Kibana public access, you will obtain a public IP address that you can use to access Kibana of this cluster.

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, click the name of the target cluster.
3. Click the **Kibana Public Access** tab to obtain the Kibana public IP address.
4. Use this IP address to access Kibana of this cluster through the Internet.

7 Vector Retrieval

7.1 Description

Large-scale image recognition and retrieval, video search, and personalized recommendation impose high requirements on the latency and accuracy of high-dimensional space vector retrieval. To facilitate large-scale vector search, CSS integrates the vector search feature powered by Huawei's vector search engine and the Elasticsearch plug-in mechanism.

Principles

Vector search works in a way similar to traditional search. To improve vector search performance, we need to:

- **Narrow down the matched scope**

Similar to traditional text search, vector search use indexes to accelerate the search instead of going through all data. Traditional text search uses inverted indexes to filter out irrelevant documents, whereas vector search creates indexes for vectors to bypass irrelevant vectors, narrowing down the search scope.
- **Reduce the complexity of calculating a single vector**

The vector search method can quantize and approximate high dimensional vectors first. By doing this, you can acquire a smaller and more relevant data set. Then more sophisticated algorithms are applied to this smaller data set to perform computation and sorting. This way, complex computation is performed on only part of the vectors, and efficiency is improved.

Vector search means to retrieve the k-nearest neighbors (KNN) to the query vector in a given vector data set by using a specific measurement method. Generally, we only focus on Approximate Nearest Neighbor (ANN), because a KNN search requires excessive computational resources.

Functions

Our engine integrates a variety of vector indexes, such as brute-force search, Hierarchical Navigable Small World (HNSW) graphs, product quantization, and

IVF-HNSW. It also supports multiple similarity calculation methods, such as Euclidean, inner product, cosine, and Hamming. The recall rate and retrieval performance of the engine are better than those of open-source engines. It can meet the requirements for high performance, high precision, low costs, and multi-modal computation.

The search engine also supports all the capabilities of the native Elasticsearch, including distribution, multi-replica, error recovery, snapshot, and permission control. The engine is compatible with the native Elasticsearch ecosystem, including the cluster monitoring tool Cerebro, the visualization tool Kibana, and the real-time data ingestion tool Logstash. Several client languages, such as Python, Java, Go, and C++, are supported.

Constraints

- Only clusters of version 7.6.2 and 7.10.2 support vector search.
- The vector search plug-in performs in-memory computing and requires more memory than common indexes do. You are advised to use memory-optimized computing specifications.

7.2 Cluster Planning for Vector Retrieval

Off-heap memory is used for index construction and query in vector retrieval. Therefore, the required cluster capacity is related to the index type and off-heap memory size. You can estimate the off-heap memory required by full indexing to select proper cluster specifications.

There are different methods for estimating the size of off-heap memory required by different types of indexes. The calculation formulas are as follows:

- **Graph Index**

$$mem_needs = (dim \times dim_size + neighbros \times 4) \times num + delta$$

NOTE

If you need to update indexes in real time, consider the off-heap memory overhead required for vector index construction and automatic merge. The actual size of required **mem_needs** is at least 1.5 to 2 times of the original estimation.

- **PQ Index**

$$mem_needs = frag_num \times frag_size \times num + delta$$

- **FALT and IVF Indexes**

$$mem_needs = dim \times dim_size \times num + delta$$

Table 7-1 Parameter description

Parameter	Description
dim	Vector dimensions

Parameter	Description
neighbors	Number of neighbors of a graph node. The default value is 64 .
dim_size	Number of bytes required by each dimension. The default value is four bytes in the float type.
num	Total number of vectors
delta	Metadata size. This parameter can be left blank.
frag_num	Number of vector segments during quantization and coding. If this parameter is not specified when an index is created, the value is determined by the vector dimension dim . <pre> if dim <= 256: frag_num = dim / 4 elif dim <= 512: frag_num = dim / 8 else : frag_num = 64 </pre>
frag_size	Size of the center point during quantization and coding. The default value is 1 . If the value of frag_num is greater than 256 , the value of frag_size is 2 .

These calculation methods can estimate the size of off-heap memory required by a complete vector index. To determine cluster specifications, you also need to consider the heap memory overhead of each node.

Heap memory allocation policy: The size of the heap memory of each node is half of the node physical memory, and the maximum size is **31 GB**.

For example, if you create a Graph index for the SIFT10M dataset, set **dim** to **128**, **dim_size** to **4**, **neighbors** to default value **64**, and **num** to **10 million**, the off-heap memory required by the Graph index is as follows:

$$mem_needs = (128 \times 4 + 64 \times 4) \times 10000000 \approx 7.5GB$$

Considering the overhead of heap memory, a single server with **8 vCPUs** and **16 GB memory** is recommended. If real-time write or update is required, you need to apply for larger memory.

7.3 Creating a Vector Index

Prerequisites

- A cluster of version 7.6.2 or 7.10.2 has been created by referring to [Cluster Planning for Vector Retrieval](#).
- Cluster advanced settings have been configured as required by referring to [Advanced Cluster Configurations](#).

Creating a Vector Index

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. Click **Dev Tools** in the navigation tree on the left and run the following command to create a vector index.

Create an index named **my_index** that contains a vector field **my_vector** and a text field **my_label**. The vector field creates the graph index and uses Euclidean distance to measure similarity.

```
PUT my_index
{
  "settings": {
    "index": {
      "vector": true
    }
  },
  "mappings": {
    "properties": {
      "my_vector": {
        "type": "vector",
        "dimension": 2,
        "indexing": true,
        "algorithm": "GRAPH",
        "metric": "euclidean"
      },
      "my_label": {
        "type": "text"
      }
    }
  }
}
```

Table 7-2 Parameters for creating an index

Type	Parameter	Description
Index settings parameters	vector	To use a vector index, set this parameter to true .
Field mappings parameters	type	Field type, for example, vector .
	dimension	Vector dimension. Value range: [1, 4096]

Type	Parameter	Description
	indexing	<p>Whether to enable vector index acceleration.</p> <p>The value can be:</p> <ul style="list-style-type: none">• false: disables vector index acceleration. If this parameter is set to false, vector data is written only to docvalues, and only ScriptScore and Rescore can be used for vector query.• true: enables vector index acceleration. If this parameter is set to true, an extra vector index is created. The index algorithm is specified by the algorithm field and VectorQuery can be used for data query. <p>Default value: false</p>

Type	Parameter	Description
	algorithm	<p>Index algorithm. This parameter is valid only when indexing is set to true.</p> <p>The value can be:</p> <ul style="list-style-type: none"> • FLAT: brute-force algorithm that calculates the distance between the target vector and all vectors in sequence. The algorithm relies on sheer computing power and its recall rate reaches 100%. You can use this algorithm if you require high recall accuracy. • GRAPH: Hierarchical Navigable Small Worlds (HNSW) algorithm for graph indexes. This algorithm is mainly used in scenarios where high performance and precision are required and the data records of a single shard is fewer than 10 million. • GRAPH_PQ: combination of the HNSW algorithm and the PQ algorithm. The PQ algorithm reduces the storage overhead of original vectors, so that HNSW can easily search for data among hundreds of millions of records. • IVF_GRAPH: combination of IVF and HNSW. The entire space is divided into multiple cluster centroids, which makes search much faster but slightly inaccurate. You can use this algorithm if you require high performance when searching for data among hundreds of millions of records. • IVF_GRAPH_PQ: combination of the PQ algorithm with the IVF or HNSW algorithm to further improve the system capacity and reduce the system overhead. This algorithm is applicable to scenarios where there are more than 1 billion files in shards and high retrieval performance is required. <p>Default value: GRAPH</p> <p>NOTE If IVF_GRAPH or IVF_GRAPH_PQ is specified, you need to pre-build and register a central point index. For details, see (Optional) Pre-Building and Registering a Center Point Vector.</p>
	Table 7-3	If Indexing is set to true , CSS provides optional parameters for vector search to achieve higher query performance or precision.

Type	Parameter	Description
	metric	Method of calculating the distance between vectors. The value can be: <ul style="list-style-type: none"> • euclidean: Euclidean distance • inner_product: inner product distance • cosine: cosine distance • hamming: Hamming distance Default value: euclidean

Table 7-3 Optional parameters

Type	Parameter	Description
Graph index configuration parameters	neighbors	Number of neighbors of each vector in a graph index. The default value is 64 . A larger value indicates higher query precision. A larger index results in a slower build and query speed. Value range: [10, 255]
	shrink	Cropping coefficient during HNSW build. The default value is 1.0f . Value range: (0.1, 10)
	scaling	Scaling ratio of the upper-layer graph nodes during HNSW build. The default value is 50 . Value range: (0, 128]
	efc	Queue size of the neighboring node during HNSW build. The default value is 200 . A larger value indicates a higher precision and slower build speed. Value range: (0, 100000]
	max_scan_num	Maximum number of nodes that can be scanned. The default value is 10000 . A larger value indicates a higher precision and slower indexing speed. Value range: (0, 1000000]
PQ index configuration parameters	centroid_num	Number of cluster centroids of each fragment. The default value is 255 . Value range: (0, 65535]
	fragment_num	Number of fragments. The default value is 0 . The plugin automatically sets the number of fragments based on the vector length. Value range: [0, 4096]

Importing Vector Data

Run the following command to import vector data. When writing vector data to the **my_index** index, you need to specify the vector field name and vector data.

- If the input vector data is an array of floating-point numbers separated by commas (,):

```
POST my_index/_doc
{
  "my_vector": [1.0, 2.0]
}
```

- If the input vector data is a Base64 string encoded using little endian:

For a high dimensional vector has a large number of valid bits, the Base64 encoding format is efficient for data transmission and parsing.

```
POST my_index/_doc
{
  "my_vector": "AACAPwAAAEA="
}
```

- To write a large amount of data, bulk operations are recommended.

```
POST my_index/_bulk
{"index": {}}
{"my_vector": [1.0, 2.0], "my_label": "red"}
{"index": {}}
{"my_vector": [2.0, 2.0], "my_label": "green"}
{"index": {}}
{"my_vector": [2.0, 3.0], "my_label": "red"}
```

Advanced Cluster Configurations

- When importing data offline, you are advised to set **refresh_interval** of indexes to **-1** to disable automatic index refreshing and improve batch write performance.
- You are advised to set **number_of_replicas** to **0**. After the offline data import is complete, you can modify the parameter value as needed.
- The parameters of other advanced functions as follows:

Table 7-4 Cluster parameters

Parameter	Description
native.cache.circuit_breaker.enabled	Whether to enable the circuit breaker for off-heap memory. Default value: true
native.cache.circuit_breaker.cpu.limit	Upper limit of off-heap memory usage of the vector index. For example, if the overall memory of a host is 128 GB and the heap memory occupies 31 GB, the default upper limit of the off-heap memory usage is 43.65 GB, that is, (128 - 31) x 45%. If the off-heap memory usage exceeds its upper limit, the circuit breaker will be triggered. Default value: 45%

Parameter	Description
native.cache.expire.enabled	Whether to enable the cache expiration policy. If this parameter is set to true , some cache items that have not been accessed for a long time will be cleared. Value: true or false Default value: false
native.cache.expire.time	Expiration time. Default value: 24h
native.vector.index_threads	Number of threads used for creating underlying indexes. Each shard uses multiple threads. Set a relatively small value to avoid resource preemption caused by the build queries of too many threads. Default value: 4

7.4 Querying Vectors

Standard Query

Standard vector query syntax is provided for vector fields with vector indexes. The following command will return *n* (specified by **size/topk**) data records that are most close to the query vector.

```
POST my_index/_search
{
  "size":2,
  "_source": false,
  "query": {
    "vector": {
      "my_vector": {
        "vector": [1, 1],
        "topk":2
      }
    }
  }
}
```

Table 7-5 Parameters for standard query

Parameter	Description
vector (the first one)	Indicates that the query type is VectorQuery .
my_vector	Indicates the name of the vector field you want to query.
vector (the second one)	Indicates the vector value you want to query, which can be an array or a Base64 string
topk	Same as the value of size generally.

Parameter	Description
Table 7-6	Indicates optional query parameters. You can adjust the vector index parameters to achieve higher query performance or precision.

Table 7-6 Optional query parameters

Type	Parameter	Description
Graph index configuration parameters	ef	Queue size of the neighboring node during the query. A larger value indicates a higher query precision and slower query speed. The default value is 200 . Value range: (0, 100000]
	max_scan_number	Maximum number of scanned nodes. A larger value indicates a higher query precision and slower query speed. The default value is 10000 . Value range: (0, 1000000]
IVF index configuration parameters	nprobe	Number of center points. A larger value indicates a higher query precision and slower query speed. The default value is 100 . Value range: (0, 100000]

Compound Query

Vector search can be used together with other Elasticsearch subqueries, such as Boolean query and post-filtering, for compound query.

In the following two examples, top 10 (**topk**) results closest to the query vector are queried first. **filter** retains only the results whose **my_label** field is **red**.

- Example of a Boolean query

```
POST my_index/_search
{
  "size": 10,
  "query": {
    "bool": {
      "must": {
        "vector": {
          "my_vector": {
            "vector": [1, 2],
            "topk": 10
          }
        }
      }
    },
    "filter": {
      "term": { "my_label": "red" }
    }
  }
}
```

- Example of post-filtering

```
GET my_index/_search
{
  "size": 10,
  "query": {
    "vector": {
      "my_vector": {
        "vector": [1, 2],
        "topk": 10
      }
    }
  },
  "post_filter": {
    "term": { "my_label": "red" }
  }
}
```

ScriptScore Query

You can use **script_score** to perform Nearest Neighbor Search (NSS) on vectors. The query syntax is provided below.

The pre-filtering condition can be any query. **script_score** traverses only the pre-filtered results, calculates the vector similarity, and sorts and returns the results. The performance of this query depends on the size of the intermediate result set after the pre-filtering. If the pre-filtering condition is set to **match_all**, brute-force search is performed on all data.

```
POST my_index/_search
{
  "size": 2,
  "query": {
    "script_score": {
      "query": {
        "match_all": {}
      },
      "script": {
        "source": "vector_score",
        "lang": "vector",
        "params": {
          "field": "my_vector",
          "vector": [1.0, 2.0],
          "metric": "euclidean"
        }
      }
    }
  }
}
```

Table 7-7 script_score parameters

Parameter	Description
source	Script description. Its value is vector_score if the vector similarity is used for scoring.
lang	Script syntax description. Its value is vector .
field	Vector field name
vector	Vector data to be queried

Parameter	Description
metric	Measurement method, which can be euclidean , inner_product , cosine , and hamming . Default value: euclidean

Re-Score Query

If the **GRAPH_PQ** or **IVF_GRAPH_PQ** index is used, the query results are sorted based on the asymmetric distance calculated by PQ. CSS supports re-scoring and sorting of query results to improve the recall rate.

Assuming that **my_index** is a PQ index, an example of re-scoring the query results is as follows:

```
GET my_index/_search
{
  "size": 10,
  "query": {
    "vector": {
      "my_vector": {
        "vector": [1.0, 2.0],
        "topk": 100
      }
    }
  },
  "rescore": {
    "window_size": 100,
    "vector_rescore": {
      "field": "my_vector",
      "vector": [1.0, 2.0],
      "metric": "euclidean"
    }
  }
}
```

Table 7-8 Rescore parameter description

Parameter	Description
window_size	Vector retrieval returns <i>topk</i> search results and sorts the first <i>window_size</i> results.
field	Vector field name
vector	Vector data to be queried
metric	Measurement method, which can be euclidean , inner_product , cosine , and hamming . Default value: euclidean

7.5 Optimizing the Performance of Vector Retrieval

Optimizing Write Performance

- To reduce the cost of backup, disable the backup function before data import and enable it afterwards.
- Set **refresh_interval** to **120s** or a larger value. Larger segments can reduce the vector index build overhead caused by merging.
- Increase the value of **native.vector.index_threads** (the default value is **4**) to increase the number of threads for vector index build.

```
PUT _cluster/settings
{
  "persistent": {
    "native.vector.index_threads": 8
  }
}
```

Optimizing Query Performance

- After importing data in batches, you can run the **forcemerge** command to improve the query efficiency.
- If the off-heap memory required by the vector index exceeds the circuit breaker limit, index entry swap-in and swap-out occur, which affects the query performance. In this case, you can increase the circuit breaker threshold of off-heap memory.

```
POST index_name/_forcemerge?max_num_segments=1
```

```
PUT _cluster/settings
{
  "persistent": {
    "native.cache.circuit_breaker.cpu.limit": "75%"
  }
}
```

- If the end-to-end latency is greater than the **took** value in the returned result, you can configure **_source** to reduce the **fdt** file size and reduce the **fetch** overhead.

```
PUT my_index
{
  "settings": {
    "index": {
      "vector": "true"
    },
    "index.soft_deletes.enabled": false
  },
  "mappings": {
    "_source": {
      "excludes": ["my_vector"]
    },
    "properties": {
      "my_vector": {
        "type": "vector",
        "dimension": 128,
        "indexing": true,
        "algorithm": "GRAPH",
        "metric": "euclidean"
      }
    }
  }
}
```

7.6 (Optional) Pre-Building and Registering a Center Point Vector

When you perform operations in [Creating a Vector Index](#), if `IVF_GRAPH` and `IVF_GRAPH_PQ` index algorithms are selected, you need to pre-build and register the center point vector.

Context

The vector index acceleration algorithms `IVF_GRAPH` and `IVF_GRAPH_PQ` are suitable for ultra-large-scale computing. These two algorithms allow you to narrow down the query range by dividing a vector space into subspaces through clustering or random sampling. Before pre-build, you need to obtain all center point vectors by clustering or random sampling.

Then, pre-construct and register the center point vectors to create the `GRAPH` or `GRAPH_PQ` index and register them with the Elasticsearch cluster. All nodes in the cluster can share the index file. Reuse of the center index among shards can effectively reduce the training overhead and the number of center index queries, improving the write and query performance.

Procedure

1. On the **Clusters** page, locate the target cluster, and click **Access Kibana** in the **Operation** column.
2. Click **Dev Tools** in the navigation tree on the left.
3. Create a center point index table.
 - For example, if the created index is named `my_dict`, `number_of_shards` of the index must be set to `1`. Otherwise, the index cannot be registered.
 - If you want to use the `IVF_GRAPH` index, set `algorithm` of the center point index to `GRAPH`.
 - If you want to use the `IVF_GRAPH_PQ` index, set `algorithm` of the center point index to `GRAPH_PQ`.

```
PUT my_dict
{
  "settings": {
    "index": {
      "vector": true
    },
    "number_of_shards": 1,
    "number_of_replicas": 0
  },
  "mappings": {
    "properties": {
      "my_vector": {
        "type": "vector",
        "dimension": 2,
        "indexing": true,
        "algorithm": "GRAPH",
        "metric": "euclidean"
      }
    }
  }
}
```

- Write the center point vector to the created index.
Write the center point vector obtained through sampling or clustering into the created **my_dict** index by referring to [Importing Vector Data](#).

- Call the registration API.

Register the created **my_dict** index with a **Dict** object with a globally unique identifier name (**dict_name**).

```
PUT _vector/register/my_dict
{
  "dict_name": "my_dict"
}
```

- Create an **IVF_GRAPH** or **IVF_GRAPH_PQ** index.

You do not need to specify the dimension and metric information. Simply specify the registered dictionary name.

```
PUT my_index
{
  "settings": {
    "index": {
      "vector": true
    }
  },
  "mappings": {
    "properties": {
      "my_vector": {
        "type": "vector",
        "indexing": true,
        "algorithm": "IVF_GRAPH",
        "dict_name": "my_dict",
        "offload_ivf": false
      }
    }
  }
}
```

Table 7-9 Field mappings parameters

Parameter	Description
dict_name	Specifies the name of the depended central point index. The vector dimension and measurement metric of the index are the same as those of the Dict index.
offload_ivf	Unloads the IVF inverted index implemented by the underlying index to Elasticsearch. In this way, the use of non-heap memory and the overhead of write and merge operations are reduced. However, the query performance also deteriorates. You can use the default value. Value: true or false Default value: false

7.7 Managing the Vector Index Cache

The vector retrieval engine is developed in C++ and uses off-heap memory. You can use the following APIs to manage the index cache.

- **View cache statistics.**

```
GET /_vector/stats
```

In the implementation of the vector plug-in, the vector index is the same as other types of Lucene indexes. Each segment constructs and stores an index file. During query, the index file is loaded to the non-heap memory. The plug-in uses the cache mechanism to manage the non-heap memory. You can use this API to query the non-heap memory usage, number of cache hits, and number of loading times.

- **Preload the vector index.**

```
PUT /_vector/warmup/{index_name}
```

You can use this API to preload the vector index specified by **index_name** to the off-heap memory for query.

- **Clear the cache.**

```
PUT /_vector/clear/cache
```

```
PUT /_vector/clear/cache/index_name
```

The caching mechanism limits the non-heap memory usage when vector indexes are used. When the total index size exceeds the cache size limit, index entry swap-in and swap-out occur, which affects the query performance. You can use this API to clear unnecessary index cache to ensure the query performance of hot data indexes.

7.8 Sample Code for Vector Search on a Client

Elasticsearch provides standard REST APIs and clients developed using Java, Python, and Go.

Based on the open-source dataset **SIFT1M** (<http://corpus-texmex.irisa.fr/>) and Python Elasticsearch client, this section provides a code snippet for creating a vector index, importing vector data, and querying vector data on the client.

Prerequisites

The Python dependency package has been installed on the client. If it is not installed, run the following commands to install it:

```
pip install numpy
pip install elasticsearch==7.6.0
```

Sample Code

```
import numpy as np
import time
import json

from concurrent.futures import ThreadPoolExecutor, wait
from elasticsearch import Elasticsearch
from elasticsearch import helpers

endpoint = 'http://xxx.xxx.xxx.xxx:9200/'

# Construct an Elasticsearch client object
es = Elasticsearch(endpoint)

# Index mapping information
index_mapping = ""
{
  "settings": {
```

```
"index": {
  "vector": "true"
}
},
"mappings": {
  "properties": {
    "my_vector": {
      "type": "vector",
      "dimension": 128,
      "indexing": true,
      "algorithm": "GRAPH",
      "metric": "euclidean"
    }
  }
}
}
}
"""

# Create an index.
def create_index(index_name, mapping):
    res = es.indices.create(index=index_name, ignore=400, body=mapping)
    print(res)

# Delete an index.
def delete_index(index_name):
    res = es.indices.delete(index=index_name)
    print(res)

# Refresh indexes.
def refresh_index(index_name):
    res = es.indices.refresh(index=index_name)
    print(res)

# Merge index segments.
def merge_index(index_name, seg_cnt=1):
    start = time.time()
    es.indices.forcemerge(index=index_name, max_num_segments=seg_cnt, request_timeout=36000)
    print(f" Complete the merge within {time.time() - start} seconds")

# Load vector data.
def load_vectors(file_name):
    fv = np.fromfile(file_name, dtype=np.float32)
    dim = fv.view(np.int32)[0]
    vectors = fv.reshape(-1, 1 + dim)[: , 1:]
    return vectors

# Load the ground_truth data.
def load_gts(file_name):
    fv = np.fromfile(file_name, dtype=np.int32)
    dim = fv.view(np.int32)[0]
    gts = fv.reshape(-1, 1 + dim)[: , 1:]
    return gts

def partition(ls, size):
    return [ls[i:i + size] for i in range(0, len(ls), size)]

# Write vector data.
def write_index(index_name, vec_file):
    pool = ThreadPoolExecutor(max_workers=8)
    tasks = []

    vectors = load_vectors(vec_file)
    bulk_size = 1000
```

```
partitions = partition(vectors, bulk_size)

start = time.time()
start_id = 0
for vecs in partitions:
    tasks.append(pool.submit(write_bulk, index_name, vecs, start_id))
    start_id += len(vecs)
wait(tasks)
print(f" Complete the writing within {time.time() - start} seconds")

def write_bulk(index_name, vecs, start_id):
    actions = [
        {
            "_index": index_name,
            "my_vector": vecs[j].tolist(),
            "_id": str(j + start_id)
        }
        for j in range(len(vecs))
    ]
    helpers.bulk(es, actions, request_timeout=3600)

# Query an index.
def search_index(index_name, query_file, gt_file, k):
    print("Start query! Index name: " + index_name)

    queries = load_vectors(query_file)
    gt = load_gts(gt_file)

    took = 0
    precision = []
    for idx, query in enumerate(queries):
        hits = set()
        query_json = {
            "size": k,
            "_source": False,
            "query": {
                "vector": {
                    "my_vector": {
                        "vector": query.tolist(),
                        "topk": k
                    }
                }
            }
        }
        res = es.search(index=index_name, body=json.dumps(query_json))

        for hit in res['hits']['hits']:
            hits.add(int(hit['_id']))
        precision.append(len(hits.intersection(set(gt[idx, :k]))) / k)
        took += res['took']

    print("precision: " + str(sum(precision) / len(precision)))
    print(f" Complete the retrieval within {took / 1000:.2f} seconds; average took size is {took / len(queries):.2f} ms")

if __name__ == "__main__":
    vec_file = r"./data/sift/sift_base.fvecs"
    qry_file = r"./data/sift/sift_query.fvecs"
    gt_file = r"./data/sift/sift_groundtruth.ivecs"

    index = "test"
    create_index(index, index_mapping)
    write_index(index, vec_file)
    merge_index(index)
    refresh_index(index)
```

```
search_index(index, qry_file, gt_file, 10)
```

8 Working with Kibana

8.1 Logging In to Kibana

Prerequisites

A CSS cluster has been created.

Procedure

- Logging in to the console
 - a. Log in to the CSS management console.
 - b. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column to go to the Kibana login page.
 - Non-security cluster: The Kibana console is displayed.
 - Security cluster: Enter the username and password on the login page and click **Log In** to go to the Kibana console. The default username is **admin** and the password is the one specified during cluster creation.
 - c. After the login is successful, you can access the Elasticsearch cluster through Kibana.
- Logging in using an EIP

If you have enabled Kibana public access during cluster creation, you can use the Kibana public IP address to log in to the cluster. For details, see .

8.2 Creating a User and Granting Permissions by Using Kibana

Prerequisites

The security mode has been enabled for the cluster.

Description

Table 8-1 Parameters for creating a user and assigning permissions on Kibana

Parameter	Description
Permission	Single action, for example, creating an index (for example, indices:admin/create)
Action group	A group of permissions. For example, the predefined SEARCH action group grants roles permissions to use _search and _msearchAPI .
Role	A role is a combination of permissions and action groups, including operation permissions on clusters, indexes, documents, or fields.
Backend role	(Optional) Other external roles from the backend such as LDAP/Active Directory
User	A user can send operation requests to the Elasticsearch cluster. The user has credentials such as username and password, zero or more backend roles, and zero or more custom attributes.
Role mapping	A user will be assigned a role after successful authentication. Role mapping is to map a role to a user (or a backend role). For example, the mapping from kibana_user (role) to jdoh (user) means that John Doe obtains all permissions of kibana_user after being authenticated by kibana_user . Similarly, the mapping from all_access (role) to admin (backend role) means that any user with the backend role admin (from the LDAP/Active Directory server) has all the permissions of role all_access after being authenticated. You can map each role to multiple users or backend roles.

Procedure

NOTE

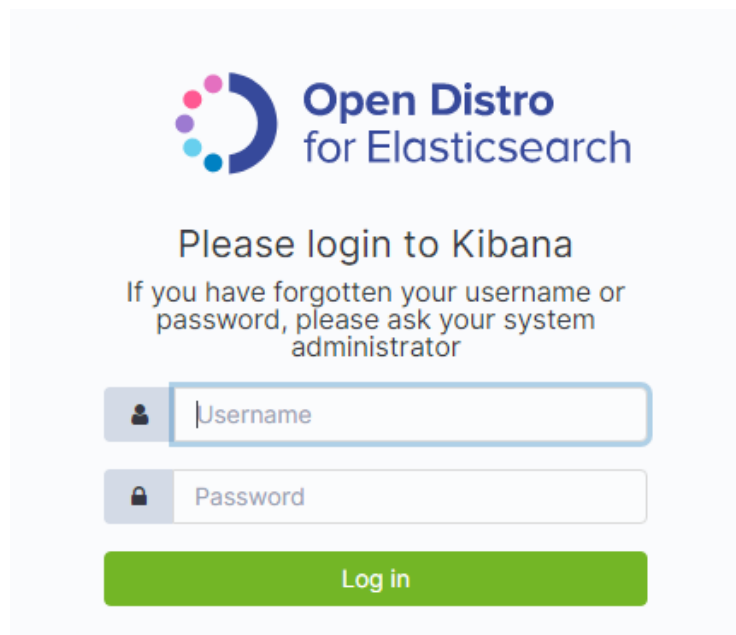
- The Kibana UI varies depending on the Kibana version, but their operations are similar. This section takes Kibana 7.6.2 as an example to describe the procedure.
- You can customize the username, role name, and tenant name in Kibana. Chinese characters are not allowed.

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.

Enter the administrator username and password to log in to Kibana.

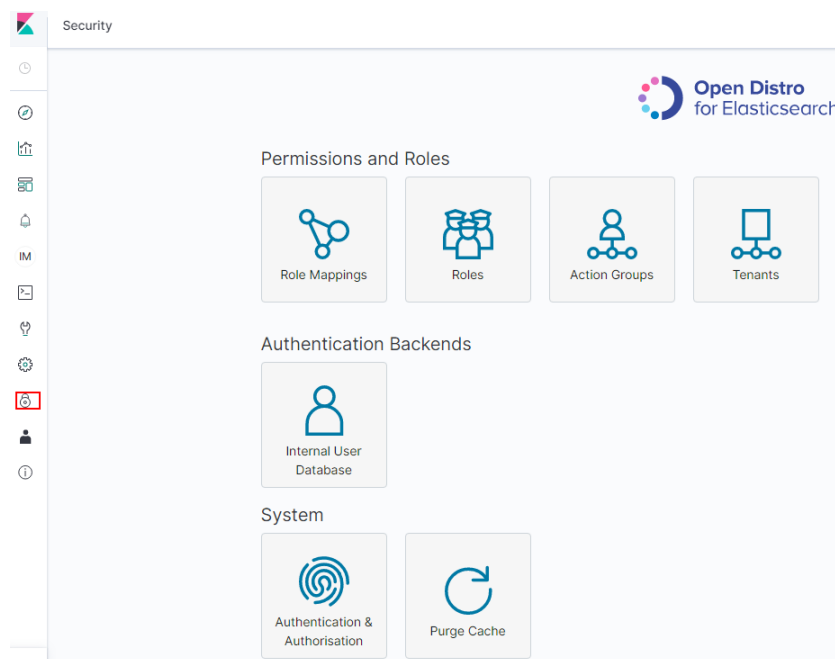
- Username: admin (default administrator account name)
- Password: Enter the administrator password you set when creating the cluster in security mode.

Figure 8-1 Login page



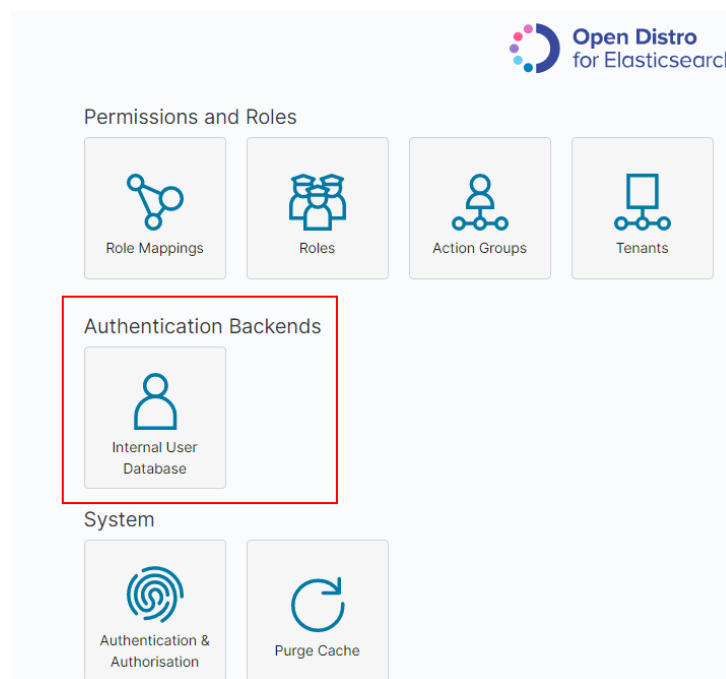
3. Click the **Security** icon on the Kibana operation page.

Figure 8-2 Going to the Security page



4. Create a user.
 - a. Choose **Authentication Backends > Internal Users Database**.

Figure 8-3 Adding a user (1)




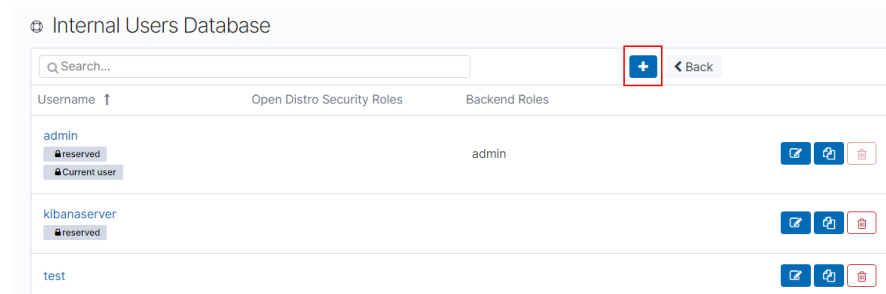
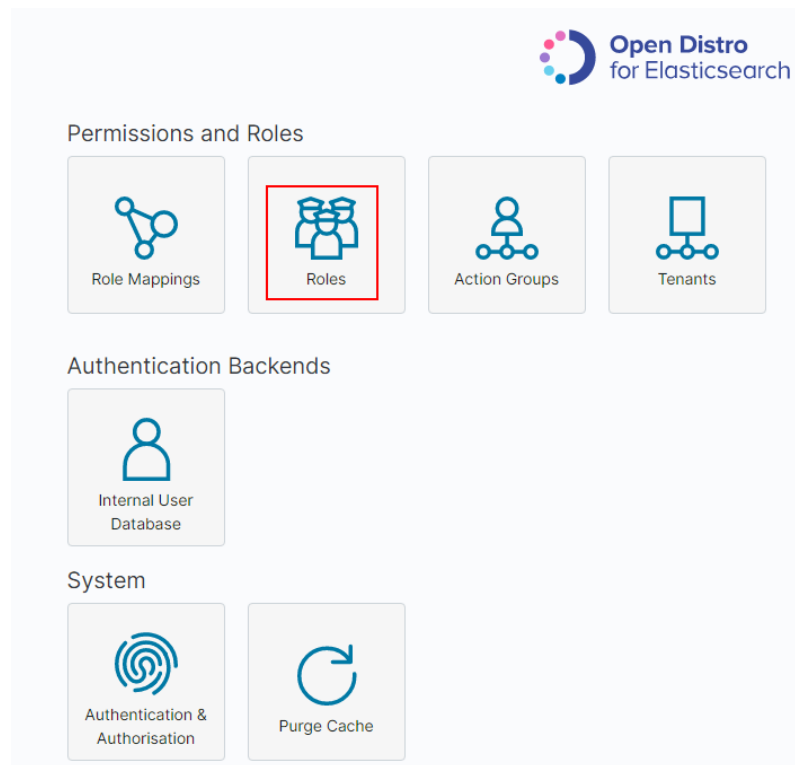
- b. On the **Internal Users Database** page, choose . The page for adding user information is displayed.

Figure 8-4 Adding a user (2)



- c. On the user creation page, configure **Username** and **Password**, and click **Submit**.
The user will be displayed in the user list.
5. Configure roles and permissions for the user.
- a. Click **Roles**.

Figure 8-5 Adding a role




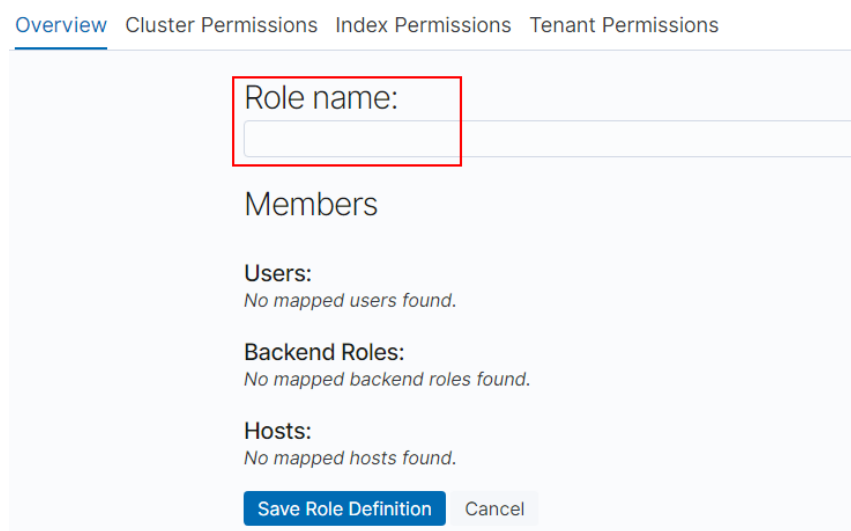
- b. On the **Open Distro Security Roles** page, click  .
 - i. Enter a role name on the **Overview** page.

Figure 8-6 Entering a role name



- ii. Configure CSS cluster permissions on the **Cluster Permissions** page. You can skip this step.
- iii. Configure index permissions on the **Index Permissions** page.
Index patterns: Set this parameter to the name of the index whose permission needs to be configured. For example, **my_store**.

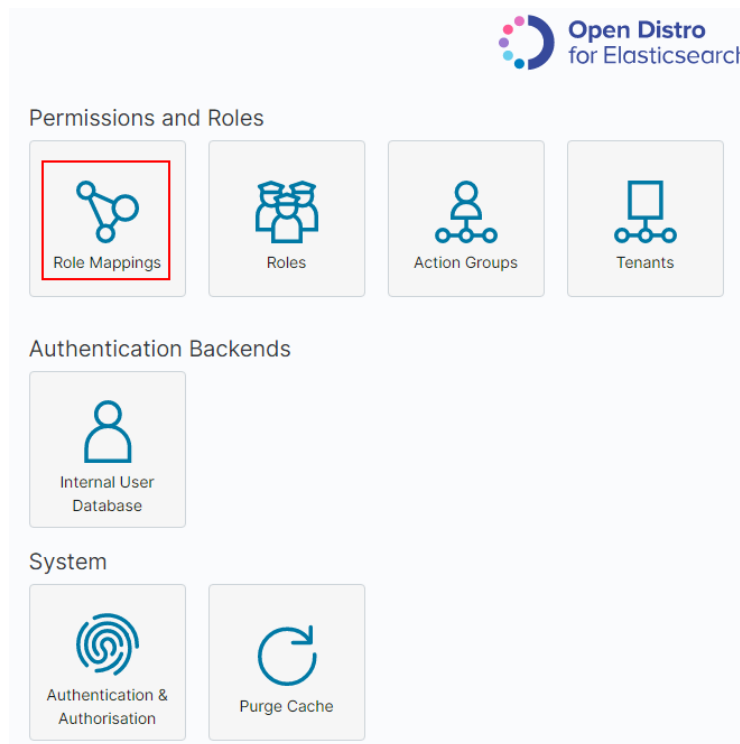
 **NOTE**

Use different names for the index and the user.

Configure **Permissions: Action Groups** as required, for example, select the read-only permission **Search**.

- iv. Retain the default settings on the **Tenant Permissions** page.
After the configuration is complete, the role will be displayed.
- 6. Assign a role to the user.
 - a. Click **Role Mappings**.

Figure 8-7 Role mapping




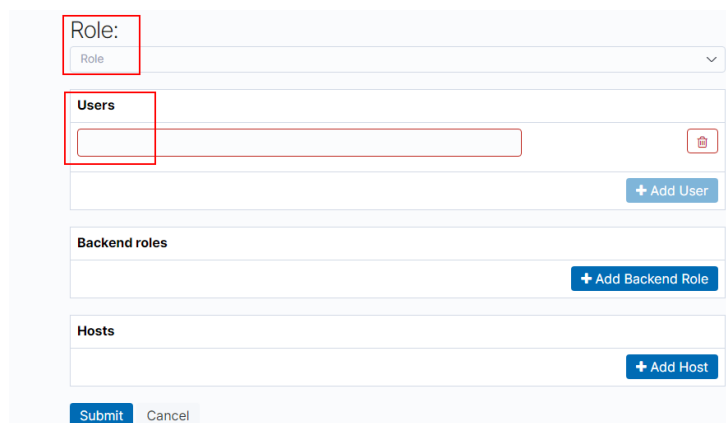
- b. Click  to add the mapping between users and roles.

Figure 8-8 Users and roles



- c. Click **Submit**.
7. Verify that the configuration takes effect in Kibana.

8.3 Managing Index Statuses

8.3.1 Creating and Managing Indexes

Clusters of version 7.6.2 or later support index status management. ISM is a plugin that allows you to automate periodic and administrative operations based on changes on the index age, index size, or number of documents. When using the ISM plug-in, you can define policies that automatically handle index rollovers or deletions based on your needs.

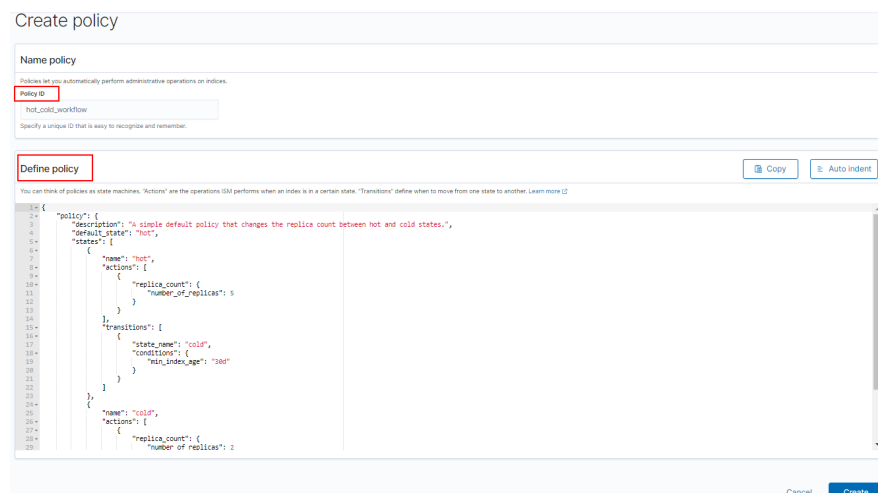
NOTE

The following procedure uses Kibana 7.6.2 as an example. The Kibana UI varies depending on the Kibana version, but their operations are similar.

Creating an Index Policy

1. Log in to Kibana and choose **IM** or **Index Management** on the left. The **Index Management** page is displayed.
2. Click **Create policy** to create an index policy.
3. Enter a policy ID in the **Policy ID** text box and enter your policy in the **Define policy** text box.

Figure 8-9 Configuring a policy



4. Click **Create**.

Attaching a Policy to an Index

You can attach a policy to one or more indexes and add the policy ID to an index template. When you create indexes using that index template pattern, the policy will be attached to all created indexes.

- **Method 1: Kibana commands**

On the **Dev Tools** page of Kibana, run the following command to associate a policy ID with an index template:

```
PUT _template/<template_name>
{
  "index_patterns": ["index_name-*"],
  "settings": {
    "opendistro.index_state_management.policy_id": "policy_id"
  }
}
```

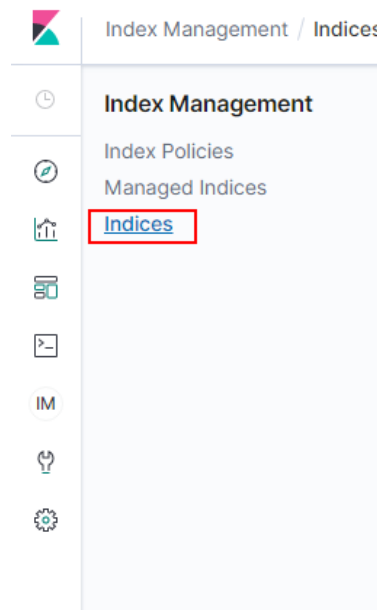
- **<template_name>**: Replace it with the name of a created index template.
- **policy_id**: Replace it with a custom policy ID.

For details about how to create an index template, see [Index Template](#).

- **Method 2: Kibana console**

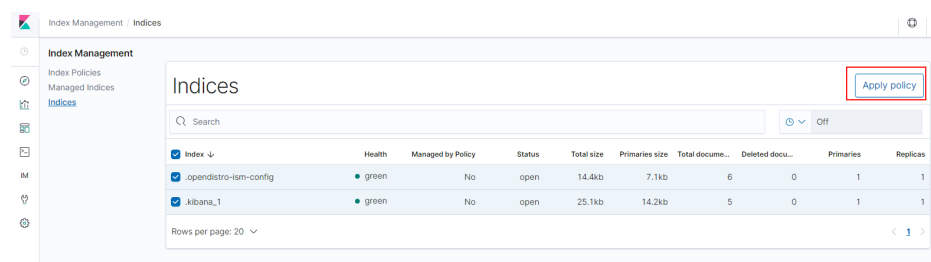
- On the **Index Management** page of Kibana, choose **Indices**.

Figure 8-10 Choosing Indices



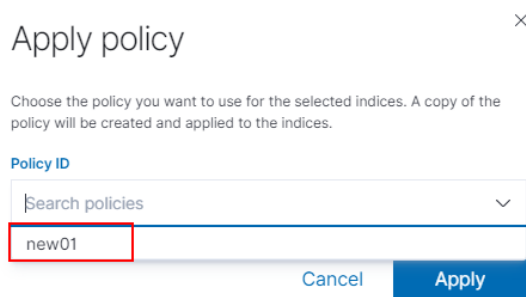
- In the **Indices** list, select the target index to which you want to attach a policy.
- Click **Apply policy** in the upper right corner.

Figure 8-11 Adding a policy



- d. Select the policy you created from the **Policy ID** drop-down list.

Figure 8-12 Selecting a policy



- e. Click **Apply**.
After you attach a policy to an index, ISM creates a job that runs every 5 minutes by default, to execute the policy, check conditions, and convert the index to different statuses.

Managing Index Policies

1. Click **Managed Indices**.
2. If you want to change the policy, click **Change policy**. For details, see [Changing Policies](#).
3. To delete a policy, select your policy, and click **Remove policy**.
4. To retry a policy, select your policy, and click **Retry policy**.

For details, see [Index State Management](#).

8.3.2 Changing Policies

You can change any managed index policy. ISM has constraints to ensure that policy changes do not break indexes.

If an index is stuck in its current status, never proceeding, and you want to update its policy immediately, make sure that the new policy includes the same status (same name, action, and order) as the old policy. In this case, ISM applies the new policy even if the policy is being executed.

If you update the policy without including an identical status, ISM updates the policy only after all actions in the current status finish executing. Alternatively, you can select a specific status in the old policy and have the new policy take effect.

To change a policy using Kibana, do the following:

1. Under **Managed Indices**, select the indices to which you want to attach the new policy.
2. Click **Change policy** in the upper right corner. The **Choose managed indices** page is displayed. Configure parameters required for changing a policy.

Table 8-2 Parameters required for changing a policy

Parameter	Description
Managed indices	Select the indexes to which you want to attach the new policy. Multiple indices can be selected.
State filters	Select an index status. When a status is selected, the new policy is attached to an index in this status.
New policy	Select a new policy.

3. After configuration is complete, click **Change**.

8.4 How Do I Connect a User-built Kibana with Elasticsearch on Huawei Cloud?

To connect the self-built Kibana with Elasticsearch on Huawei Cloud, the following conditions must be met:

1. The local environment must support access from external networks.
2. Kibana is built using ECS in the same VPC as Elasticsearch. Kibana can be accessed from the local public network.

Example of a Kibana configuration file:

Security mode:

```

elasticsearch.username: "****"
elasticsearch.password: "****"
elasticsearch.ssl.verificationMode: none
server.ssl.enabled: false
server.rewriteBasePath: false
server.port: 5601
logging.dest: /home/Ruby/log/kibana.log
pid.file: /home/Ruby/run/kibana.pid
server.host: 192.168.25.226
elasticsearch.hosts: https://10.0.0.207:9200
elasticsearch.requestHeadersWhitelist: ["securitytenant","Authorization"]
opendistro_security.multitenancy.enabled: true
opendistro_security.multitenancy.tenants.enable_global: true
opendistro_security.multitenancy.tenants.enable_private: true
opendistro_security.multitenancy.tenants.preferred: ["Private", "Global"]
opendistro_security.multitenancy.enable_filter: false
    
```

NOTE

- In security mode, the **opendistro_security_kibana** plug-in must be installed. For details, see: <https://opendistro.github.io/for-elasticsearch-docs/docs/kibana/plugins/>
- The version of the installed plug-in must be the same as that of the cluster. To check the version of the plug-in version, run the **GET _cat/plugins** command.

Non-security mode

```

server.port: 5601
logging.dest: /home/Ruby/log/kibana.log
pid.file: /home/Ruby/run/kibana.pid
server.host: 192.168.25.226
elasticsearch.hosts: http://10.0.0.207:9200
    
```

8.5 Kibana Usage Restrictions

- You can customize the username, role name, and tenant name in Kibana. Chinese characters are not allowed.
- Kibana does not support Chinese.

9 Elasticsearch SQL

For Elasticsearch 6.5.4 and later versions, Open Distro for Elasticsearch SQL lets you write queries in SQL rather than in the Elasticsearch query domain-specific language (DSL).

If you are already familiar with SQL and do not want to learn query DSL, this feature is a great option.

Basic Operations

- Kibana (recommended)
 - Log in to Kibana and send requests using request parameters or request body to **_opendistro/_sqlURI** in the **Dev Tools** page.

```
GET _opendistro/_sql?sql=select * from my-index limit 50
POST _opendistro/_sql
{
  "query": "SELECT * FROM my-index LIMIT 50"
}
```

- By default, the result is returned in the JSON structure. If you want the result to be returned in the CSV format, run the following command:

```
POST _opendistro/_sql?format=csv
{
  "query": "SELECT * FROM my-index LIMIT 50"
}
```

When data is returned in the CSV format, each row corresponds to a document and each column corresponds to a field.

- cURL commands
You can also run cURL commands in ECS to execute SQL statements.

```
curl -XPOST https://localhost:9200/_opendistro/_sql -u username:password -k -d '{"query": "SELECT * FROM kibana_sample_data_flights LIMIT 10"}' -H 'Content-Type: application/json'
```

Supported Operations

Open Distro for Elasticsearch supports the following SQL operations: statements, conditions, aggregations, include and exclude fields, common functions, joins, and show.

- Statements

Table 9-1 Statements

Statement	Example
Select	SELECT * FROM my-index
Delete	DELETE FROM my-index WHERE _id=1
Where	SELECT * FROM my-index WHERE ['field']='value'
Order by	SELECT * FROM my-index ORDER BY _id asc
Group by	SELECT * FROM my-index GROUP BY range(age, 20,30,39)
Limit	SELECT * FROM my-index LIMIT 50 (default is 200)
Union	SELECT * FROM my-index1 UNION SELECT * FROM my-index2
Minus	SELECT * FROM my-index1 MINUS SELECT * FROM my-index2

 **NOTE**

As with any complex query, large UNION and MINUS statements can strain or even crash your cluster.

- Conditions

Table 9-2 Conditions

Condition	Example
Like	SELECT * FROM my-index WHERE name LIKE 'j%'
And	SELECT * FROM my-index WHERE name LIKE 'j%' AND age > 21
Or	SELECT * FROM my-index WHERE name LIKE 'j%' OR age > 21
Count distinct	SELECT count(distinct age) FROM my-index
In	SELECT * FROM my-index WHERE name IN ('alejandro', 'carolina')
Not	SELECT * FROM my-index WHERE name NOT IN ('jane')
Between	SELECT * FROM my-index WHERE age BETWEEN 20 AND 30
Aliases	SELECT avg(age) AS Average_Age FROM my-index
Date	SELECT * FROM my-index WHERE birthday='1990-11-15'
Null	SELECT * FROM my-index WHERE name IS NULL

- Aggregations

Table 9-3 Aggregations

Aggregation	Example
avg()	SELECT avg(age) FROM my-index
count()	SELECT count(age) FROM my-index
max()	SELECT max(age) AS Highest_Age FROM my-index
min()	SELECT min(age) AS Lowest_Age FROM my-index
sum()	SELECT sum(age) AS Age_Sum FROM my-index

- Include and exclude fields

Table 9-4 Include and exclude fields

Pattern	Example
include()	SELECT include('a*'), exclude('age') FROM my-index
exclude()	SELECT exclude('*name') FROM my-index

- Functions

Table 9-5 Functions

Function	Example
floor	SELECT floor(number) AS Rounded_Down FROM my-index
trim	SELECT trim(name) FROM my-index
log	SELECT log(number) FROM my-index
log10	SELECT log10(number) FROM my-index
substring	SELECT substring(name, 2,5) FROM my-index
round	SELECT round(number) FROM my-index
sqrt	SELECT sqrt(number) FROM my-index
concat_ws	SELECT concat_ws(' ', age, height) AS combined FROM my-index

Function	Example
/	SELECT number / 100 FROM my-index
%	SELECT number % 100 FROM my-index
date_format	SELECT date_format(date, 'Y') FROM my-index

 **NOTE**

You must enable fielddata in the document mapping for most string functions to work properly.

- Joins

Table 9-6 Joins

Join	Example
Inner join	SELECT s.firstname, s.lastname, s.gender, sc.name FROM student s JOIN school sc ON sc.name = s.school_name WHERE s.age > 20
Left outer join	SELECT s.firstname, s.lastname, s.gender, sc.name FROM student s LEFT JOIN school sc ON sc.name = s.school_name
Cross join	SELECT s.firstname, s.lastname, s.gender, sc.name FROM student s CROSS JOIN school sc

For details about the restrictions, see [Joins](#).

- Show

Show commands display indexes and mappings that match an index pattern. You can use * or % for wildcards.

Table 9-7 Show

Show	Example
Show tables like	SHOW TABLES LIKE logs-*

Joins

Open Distro for Elasticsearch SQL supports inner joins, left outer joins and cross joins. Joins have the following constraints:

- You can only join two indexes.
- You must use an alias for an index (for example, people p).
- In an ON clause, you can only use the AND conditions.
- In a WHERE statement, do not combine trees that contain multiple indexes.
For example, the following statement will work:
`WHERE (a.type1 > 3 OR a.type1 < 0) AND (b.type2 > 4 OR b.type2 < -1)`
- The following statement will not work:
`WHERE (a.type1 > 3 OR b.type2 < 0) AND (a.type1 > 4 OR b.type2 < -1)`
- You cannot use GROUP BY or ORDER BY to obtain results.
- LIMIT with OFFSET (for example, LIMIT 25 OFFSET 25) is not supported.

JDBC Driver

The Java Database Connectivity (JDBC) driver allows you to integrate Open Distro for Elasticsearch with your business intelligence (BI) applications.

For details about how to download and use JAR files, see [GitHub Repositories](#).

10 Connecting a Cluster to a Dedicated Load Balancer

10.1 Scenario

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 section describes how to connect a cluster to a dedicated load balancer.

Advantages of connecting a cluster to a dedicated load balancer:

- A non-security cluster can also use capabilities of the Elastic Load Balance (ELB) service.
- You can use customized certificates for HTTPS bidirectional authentication.
- Seven-layer traffic monitoring and alarm configuration are supported, allowing you to view the cluster status at any time.

There are eight service forms for clusters in different security modes to connect to dedicated load balancers. [Table 10-1](#) describes the load balancer capabilities for the eight service forms. [Table 10-2](#) describes the configurations for the eight service forms.

NOTICE

You are not advised to connect an ELB that has bound the public network to a non-security cluster. Non-security clusters can be accessed over HTTP without security authentication. A load balancer with an EIP allows access to such clusters over the Internet, which may bring security risks.

Table 10-1 ELB capabilities for different clusters

Security Mode	Service Form Provided by ELB for External Systems	ELB Load Balancing	ELB Traffic Monitoring	ELB Two-way Authentication
Non-security	No authentication	Yes	Yes	No
	One-way authentication Two-way authentication	Yes	Yes	Yes
Security mode + HTTP	Password authentication	Yes	Yes	No
	One-way authentication + Password authentication Two-way authentication + Password authentication	Yes	Yes	Yes
Security mode + HTTPS	One-way authentication + Password authentication Two-way authentication + Password authentication	Yes	Yes	Yes

Table 10-2 Configuration for interconnecting different clusters with ELB

Security Mode	Service Form Provided by ELB for External Systems	ELB Listener			Backend Server Group		
		Frontend Protocol	Port	SSL Parsing Mode	Backend Protocol	Health Check Port	Health Check Path
Non-security	No authentication	HTTP	9200	No authentication	HTTP	9200	/
	One-way authentication	HTTPS	9200	One-way authentication	HTTP	9200	

	Two-way authentication	HTTPS	9200	Two-way authentication	HTTP	9200	
Security mode + HTTP	Password authentication	HTTP	9200	No authentication	HTTP	9200	/_opendistro/_security/health
	One-way authentication + Password authentication	HTTPS	9200	One-way authentication	HTTP	9200	
	Two-way authentication + Password authentication	HTTPS	9200	Two-way authentication	HTTP	9200	
Security mode + HTTPS	One-way authentication + Password authentication	HTTPS	9200	One-way authentication	HTTPS	9200	
	Two-way authentication + Password authentication	HTTPS	9200	Two-way authentication	HTTPS	9200	

10.2 Connecting to a Dedicated Load Balancer

This section describes how to connect a CSS cluster to a dedicated load balancer.

(Optional) Preparing a Self-signed Certificate

If the target ELB listener uses the HTTP protocol, skip this step.

Prepare and upload a self-signed certificate.

NOTE

You are advised to use a certificate purchased in Cloud Certificate Manager (CCM) or issued by an authoritative organization.

1. Log in to a Linux client where the OpenSSL tool and JDK are installed.
2. Run the following commands to create a self-signed certificate:

```
mkdir ca
mkdir server
mkdir client
```

```
#Use OpenSSL to create a CA certificate.
cd ca
#Create the OpenSSL configuration file ca_cert.conf for the CA certificate.
cat >ca_cert.conf <<EOF
[ req ]
distinguished_name = req_distinguished_name
prompt = no

[ req_distinguished_name ]
O = ELB
EOF
#Create private key file ca.key for the CA certificate.
openssl genrsa -out ca.key 2048
#Create the CSR file ca.csr for the CA certificate.
openssl req -out ca.csr -key ca.key -new -config ./ca_cert.conf
#Create a self-signed CA certificate ca.crt.
openssl x509 -req -in ca.csr -out ca.crt -sha1 -days 5000 -signkey ca.key
#Convert the CA certificate format to p12.
openssl pkcs12 -export -clcerts -in ca.crt -inkey ca.key -out ca.p12
#Convert the CA certificate format to JKS.
keytool -importkeystore -srckeystore ca.p12 -srcstoretype PKCS12 -deststoretype JKS -destkeystore
ca.jks

#Use the CA certificate to issue a server certificate.
cd ../server
#Create the OpenSSL configuration file server_cert.conf for the server certificate. Change the CN
field to the domain name or IP address of the server as required.
cat >server_cert.conf <<EOF
[ req ]
distinguished_name = req_distinguished_name
prompt = no

[ req_distinguished_name ]
O = ELB
CN = 127.0.0.1
EOF
#Create the private key file server.key for the server certificate.
openssl genrsa -out server.key 2048
#Create the CSR request file server.csr for the server certificate.
openssl req -out server.csr -key server.key -new -config ./server_cert.conf
#Use the CA certificate to issue the server certificate server.crt.
openssl x509 -req -in server.csr -out server.crt -sha1 -CAcreateserial -days 5000 -CA ../ca/ca.crt -
CAkey ../ca/ca.key
#Convert the server certificate format to p12.
openssl pkcs12 -export -clcerts -in server.crt -inkey server.key -out server.p12
#Convert the service certificate format to JKS.
keytool -importkeystore -srckeystore server.p12 -srcstoretype PKCS12 -deststoretype JKS -destkeystore
server.jks

#Use the CA certificate to issue a client certificate.
cd ../client
#Create the OpenSSL configuration file client_cert.conf for the client certificate. Change the CN field
to the domain name or IP address of the server as required.
cat >client_cert.conf <<EOF
[ req ]
distinguished_name = req_distinguished_name
prompt = no

[ req_distinguished_name ]
O = ELB
CN = 127.0.0.1
EOF
#Create private key client.key for the client certificate.
openssl genrsa -out client.key 2048
#Create the CSR file client.csr for the client certificate.
openssl req -out client.csr -key client.key -new -config ./client_cert.conf
#Use the CA certificate to issue the client certificate client.crt.
```



```
openssl x509 -req -in client.csr -out client.crt -sha1 -CAcreateserial -days 5000 -CA ../ca/ca.crt -
CAkey ../ca/ca.key
#Convert the client certificate to a p12 file that can be identified by the browser.
openssl pkcs12 -export -clcerts -in client.crt -inkey client.key -out client.p12
#Convert the client certificate format to JKS.
keytool -importkeystore -srckeystore client.p12 -srcstoretype PKCS12 -deststoretype JKS -destkeystore
client.jks
```

3. Upload the self-signed certificate. For details, see .

Creating a Dedicated Load Balancer

1. Log in to the ELB management console.
2. Create a dedicated load balancer. For details, see . [Table 10-3](#) describes the parameters required for connecting a CSS cluster with a dedicated load balancer.

Table 10-3 Parameters for interconnecting a CSS cluster with a dedicated load balancer

Parameter	Description	Example
Type	Load balancer type. Select Dedicated .	Dedicated
Billed By	Billing mode of the dedicated load balancer.	Pay-per-use
Region	Region where the CSS cluster is located.	-
IP as Backend Servers	A CSS cluster can be connected only after the cross-VPC backend is enabled.	Enabled
Network Type	Type of the network used by the load balancer to provide services for external systems.	Private IPv4 network
VPC	VPC where the load balancer works. This parameter is mandatory no matter which network type is selected. Select the VPC of the CSS cluster	-

Parameter	Description	Example
Subnet	Subnet where the load balancer is to be created. This parameter is mandatory no matter which network type is selected. Select the subnet of the CSS cluster	-
Specifications	You are advised to select Application load balancing (HTTP/HTTPS) , which provides better functions and performance.	Application load balancing (HTTP/HTTPS) Small I

Adding a Listener

1. Log in to the ELB management console.
2. Add an HTTPS listener for a dedicated load balancer. For details, see . [Table 10-4](#), [Table 10-5](#), [Table 10-6](#), and [Table 10-7](#) describe the parameters required for interconnecting a CSS cluster with a dedicated load balancer.

Table 10-4 Listener parameters

Parameter	Description	Example
Name	Listener name	listener-9200
Frontend Protocol	The protocol used by the client and listener to distribute traffic. Select a protocol as required.	HTTPS
Frontend Port	The port used by the client and listener to distribute traffic. For example, 9200. You need to specify this parameter as required.	9200
SSL Authentication	Authentication mode for the client to access the server. Select a parsing mode as required.	Two-way authentication

Parameter	Description	Example
Server Certificate	<p>The server certificate is used for SSL handshake negotiation. The certificate content and private key must be provided.</p> <p>When SSL Authentication is set to Two-way authentication, this parameter is mandatory.</p>	-
CA Certificate	<p>Also called client CA public key certificate. It is used to verify the issuer of a client certificate.</p> <p>When the HTTPS two-way authentication is enabled, an HTTPS connection can be established only when the client can provide the certificate issued by a specified CA.</p> <p>This parameter is mandatory only when the Frontend Protocol is set to HTTPS.</p>	-

Table 10-5 Parameters for backend allocation policies

Parameter	Description	Example
Backend Server Group	A group of backend servers with the same features	New
Name	Backend server group name	server_group_css
Backend Protocol	The protocol used by backend servers to receive requests	HTTPS

Parameter	Description	Example
Load Balancing Algorithm	The algorithm used by the load balancer to distribute traffic. Select an appropriate algorithm to allocate access traffic and improve load balancing capabilities.	Weighted Round Robin

 **NOTE**

Before adding a backend server, choose **IP as Backend Servers** and click **Add IP as Backend Servers** on the **Add Listener** page.

Table 10-6 Parameters for adding a backend server

Parameter	Description	Example
IP as Backend Servers	Private network address provided on the CSS console for accessing a cluster. After a cluster is scaled in or out, or a client node is added, you need to update the cross-VPC backend server IP address in a timely manner.	192.168.xxx.xxx
Backend Port	Port used by CSS to provide services for external systems. Set the value to 9200 .	9200
Weight	Weight of the weighted round robin algorithm	1

Table 10-7 Health check parameters

Parameter	Description	Example
Enable Health Check	Whether to enable health checks. You are advised to enable them.	Yes

Parameter	Description	Example
Protocol	Protocol used by the load balancer to perform health checks on backend servers. You can select HTTP, TCP, or HTTPS. Once selected, the protocol cannot be changed. Only HTTP and HTTPS are available for CSS clusters.	HTTPS
Check Path	Health check URL. <ul style="list-style-type: none"> For a CSS cluster in non-security mode, enter <code>/</code>. For a CSS cluster in security mode, enter <code>/_opendistro/_security/health</code>. 	<code>/_opendistro/_security/health</code>

- After setting the parameters, wait for 30 seconds and check the cross-VPC backend servers in the backend server group. If the health check result is **Healthy**, the listener is added successfully.

Accessing a Cluster

Run the following commands to check whether the dedicated load balancer can be connected to a cluster.

Table 10-8 Commands for accessing different clusters

Security Mode	Service Form Provided by ELB for External Systems	Curl Command for Accessing a Cluster
Non-security	No authentication	<code>curl http://IP:9200</code>
	One-way authentication	<code>curl -k --cert ./client.crt --key ./client.key https://IP:9200</code>
	Two-way authentication	<code>curl --cacert ./ca.crt --cert ./client.crt --key ./client.key https://IP:9200</code>

Security Mode	Service Form Provided by ELB for External Systems	Curl Command for Accessing a Cluster
Security mode + HTTP	Password authentication	<code>curl http://IP:9200 -u user:pwd</code>
	One-way authentication + Password authentication	<code>curl -k --cert ./client.crt --key ./client.key https://IP:9200 -u user:pwd</code>
	Two-way authentication + Password authentication	<code>curl --cacert ./ca.crt --cert ./client.crt --key ./client.key https://IP:9200 -u user:pwd</code>
Security mode + HTTPS	One-way authentication + Password authentication	<code>curl -k --cert ./client.crt --key ./client.key https://IP:9200 -u user:pwd</code>
	Two-way authentication + Password authentication	<code>curl --cacert ./ca.crt --cert ./client.crt --key ./client.key https://IP:9200 -u user:pwd</code>

Table 10-9 Variables

Variable	Description
IP	ELB IP address
user	Username for accessing the CSS cluster
pwd	Password of the user

If the Elasticsearch cluster information is returned, the connection is successful. For example, if a security cluster using the HTTPS protocol is connected to a load balancer using two-way authentication, the information shown in [Figure 10-1](#) is returned.

Figure 10-1 Accessing a cluster

```
root@xxxxxxxxxxxxx:~# curl --cacert ./ca/ca.crt --cert ./client.crt --key ./client.key https://xxxxxxxxxxxx:9200 -u admin:
{
  "name" : "css-test1-ess-esn-1-1",
  "cluster_name" : "css-test1",
  "cluster_uuid" : "nX81L1jT_2CCMBFe1bgmA",
  "version" : {
    "number" : "7.10.2",
    "build_flavor" : "oss",
    "build_type" : "tar",
    "build_hash" : "unknown",
    "build_date" : "unknown",
    "build_snapshot" : true,
    "lucene_version" : "8.7.0",
    "minimum_wire_compatibility_version" : "6.7.0",
    "minimum_index_compatibility_version" : "6.0.0-beta1"
  },
  "tagline" : "You Know, for Search"
}
```

10.3 Sample Code for Two-Way Authentication During the Access to a Cluster

This section provides the sample code for two-way authentication during the access to a cluster from a Java client.

ESSecuredClientWithCerDemo Code

```
import org.apache.commons.io.IOUtils;
import org.apache.http.auth.AuthScope;
import org.apache.http.auth.UsernamePasswordCredentials;
import org.apache.http.client.CredentialsProvider;
import org.apache.http.impl.client.BasicCredentialsProvider;
import org.apache.http.HttpHost;
import org.apache.http.nio.conn.ssl.SSLIOStrategy;
import org.elasticsearch.action.search.SearchRequest;
import org.elasticsearch.action.search.SearchResponse;
import org.elasticsearch.client.RequestOptions;
import org.elasticsearch.client.RestClient;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.client.RestHighLevelClient;
import org.elasticsearch.index.query.QueryBuilders;
import org.elasticsearch.search.SearchHit;
import org.elasticsearch.search.SearchHits;
import org.elasticsearch.search.builder.SearchSourceBuilder;
import java.io.FileInputStream;
import java.io.IOException;
import java.security.KeyStore;
import java.security.SecureRandom;
import javax.net.ssl.HostnameVerifier;
import javax.net.ssl.KeyManagerFactory;
import javax.net.ssl.SSLContext;
import javax.net.ssl.SSLSession;
import javax.net.ssl.TrustManagerFactory;
public class ESSecuredClientWithCerDemo {
    private static final String KEY_STORE_PWD = "";
    private static final String TRUST_KEY_STORE_PWD = "";
    private static final String CA_JKS_PATH = "ca.jks";
    private static final String CLIENT_JKS_PATH = "client.jks";
    private static final String ELB_ADDRESS = "127.0.0.1";
    private static final int ELB_PORT = 9200;
    private static final String CSS_USERNAME = "user";
    private static final String CSS_PWD = "";
    public static void main(String[] args) {
        // Create a client.
        RestHighLevelClient client = initESClient(ELB_ADDRESS, CSS_USERNAME, CSS_PWD);
        try {
            // Search match_all, which is equivalent to {"query":{"match_all":{"}}}.
            SearchRequest searchRequest = new SearchRequest();
            SearchSourceBuilder searchSourceBuilder = new SearchSourceBuilder();
            searchRequest.source(searchSourceBuilder);
            // query
```

```
        SearchResponse searchResponse = client.search(searchRequest, RequestOptions.DEFAULT);
        System.out.println("query result: " + searchResponse.toString());
        SearchHits hits = searchResponse.getHits();
        for (SearchHit hit : hits) {
            System.out.println(hit.getSourceAsString());
        }
        System.out.println("query success");
        Thread.sleep(2000L);
    } catch (InterruptedException | IOException e) {
        e.printStackTrace();
    } finally {
        IOUtils.closeQuietly(client);
    }
}

private static RestHighLevelClient initESClient(String clusterAddress, String userName, String password) {
    final CredentialsProvider credentialsProvider = new BasicCredentialsProvider();
    credentialsProvider.setCredentials(AuthScope.ANY, new UsernamePasswordCredentials(userName,
password));
    SSLContext ctx = null;
    try {
        KeyStore ks = getKeyStore(CLIENT_JKS_PATH, KEY_STORE_PWD, "JKS");
        KeyManagerFactory kmf = KeyManagerFactory.getInstance("SunX509");
        kmf.init(ks, SERVER_KEY_STORE_PWD.toCharArray());
        KeyStore tks = getKeyStore(CA_JKS_PATH, TRUST_KEY_STORE_PWD, "JKS");
        TrustManagerFactory tmf = TrustManagerFactory.getInstance("SunX509");
        tmf.init(tks);
        ctx = SSLContext.getInstance("SSL", "SunJSSE");
        ctx.init(kmf.getKeyManagers(), tmf.getTrustManagers(), new SecureRandom());
    } catch (Exception e) {
        e.printStackTrace();
    }
    SSLIOSessionStrategy sessionStrategy = new SSLIOSessionStrategy(ctx, new HostnameVerifier() {
        @Override
        public boolean verify(String arg0, SSLSession arg1) {
            return true;
        }
    });
    SecuredHttpClientConfigCallback httpClientConfigCallback = new
SecuredHttpClientConfigCallback(sessionStrategy,
        credentialsProvider);
    RestClientBuilder builder = RestClient.builder(new HttpHost(clusterAddress, ELB_PORT, "https"))
        .setHttpClientConfigCallback(httpClientConfigCallback);
    RestHighLevelClient client = new RestHighLevelClient(builder);
    return client;
}

private static KeyStore getKeyStore(String path, String pwd, String type) {
    KeyStore keyStore = null;
    FileInputStream is = null;
    try {
        is = new FileInputStream(path);
        keyStore = KeyStore.getInstance(type);
        keyStore.load(is, pwd.toCharArray());
    } catch (Exception e) {
        e.printStackTrace();
    } finally {
        IOUtils.closeQuietly(is);
    }
    return keyStore;
}
}
```

SecuredHttpClientConfigCallback Code

```
import org.apache.http.client.CredentialsProvider;
import org.apache.http.impl.nio.client.HttpAsyncClientBuilder;
import org.apache.http.nio.conn.ssl.SSLIOSessionStrategy;
import org.elasticsearch.client.RestClientBuilder;
import org.elasticsearch.common.Nullable;
import java.util.Objects;
```



```
class SecuredHttpClientConfigCallback implements RestClientBuilder.HttpClientConfigCallback {
    @Nullable
    private final CredentialsProvider credentialsProvider;
    /**
     * The {@link SSLIOStrategy} for all requests to enable SSL / TLS encryption.
     */
    private final SSLIOStrategy sslStrategy;
    /**
     * Create a new {@link SecuredHttpClientConfigCallback}.
     *
     * @param credentialsProvider The credential provider, if a username/password have been supplied
     * @param sslStrategy The SSL strategy, if SSL / TLS have been supplied
     * @throws NullPointerException if {@code sslStrategy} is {@code null}
     */
    SecuredHttpClientConfigCallback(final SSLIOStrategy sslStrategy,
        @Nullable final CredentialsProvider credentialsProvider) {
        this.sslStrategy = Objects.requireNonNull(sslStrategy);
        this.credentialsProvider = credentialsProvider;
    }
    /**
     * Get the {@link CredentialsProvider} that will be added to the HTTP client.
     *
     * @return Can be {@code null}.
     */
    @Nullable
    CredentialsProvider getCredentialsProvider() {
        return credentialsProvider;
    }
    /**
     * Get the {@link SSLIOStrategy} that will be added to the HTTP client.
     *
     * @return Never {@code null}.
     */
    SSLIOStrategy getSSLStrategy() {
        return sslStrategy;
    }
    /**
     * Sets the {@linkplain HttpAsyncClientBuilder#setDefaultCredentialsProvider(CredentialsProvider)
    credential provider},
     *
     * @param httpClientBuilder The client to configure.
     * @return Always {@code httpClientBuilder}.
     */
    @Override
    public HttpAsyncClientBuilder customizeHttpClient(final HttpAsyncClientBuilder httpClientBuilder) {
        // enable SSL / TLS
        httpClientBuilder.setSSLStrategy(sslStrategy);
        // enable user authentication
        if (credentialsProvider != null) {
            httpClientBuilder.setDefaultCredentialsProvider(credentialsProvider);
        }
        return httpClientBuilder;
    }
}
```

pom.xml Code

```
<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/
maven-4.0.0.xsd">
    <modelVersion>4.0.0</modelVersion>
    <groupId>1</groupId>
    <artifactId>ESClient</artifactId>
    <version>1.0-SNAPSHOT</version>
    <name>ESClient</name>

    <properties>
```

```
<maven.compiler.source>8</maven.compiler.source>
<maven.compiler.target>8</maven.compiler.target>
<project.build.sourceEncoding>UTF-8</project.build.sourceEncoding>
<elasticsearch.version>7.10.2</elasticsearch.version>
</properties>
<dependencies>
  <dependency>
    <groupId>org.elasticsearch.client</groupId>
    <artifactId>transport</artifactId>
    <version>${elasticsearch.version}</version>
  </dependency>
  <dependency>
    <groupId>org.elasticsearch</groupId>
    <artifactId>elasticsearch</artifactId>
    <version>${elasticsearch.version}</version>
  </dependency>
  <dependency>
    <groupId>org.elasticsearch.client</groupId>
    <artifactId>elasticsearch-rest-high-level-client</artifactId>
    <version>${elasticsearch.version}</version>
  </dependency>
  <dependency>
    <groupId>commons-io</groupId>
    <artifactId>commons-io</artifactId>
    <version>2.11.0</version>
  </dependency>
</dependencies>
</project>
```

11 Enhanced Features

11.1 Storage-Compute Decoupling

11.1.1 Context

You can store hot data on SSD to achieve the optimal query performance, and store historical data in OBS to reduce data storage costs.

Application Scenarios

A large volume of data is written to and stored in SSDs. If historical data is no longer updated (is turned into cold data) and its QPS decreases, you can call CSS APIs to dump hot data from SSDs to OBS buckets. This operation freezes indexes, decoupling compute from storage.

Constraints

- Currently, only versions 7.6.2 and 7.10.2 support storage-compute decoupling.
- The storage-compute decoupling feature depends on OBS. Therefore, you must comply with the restrictions on OBS bandwidth and QPS. For details, see OBS Restrictions. If these restrictions are violated, the performance of queries on OBS will deteriorate. For example, the speed of restoring shards and querying data will become slow.

11.1.2 Freezing an Index

Precautions

- Before freezing an index, ensure no data is being written to it. The index will be set to read only before being frozen, and data write will fail.
- After an index is frozen,
 - It becomes read-only.
 - The index data will be dumped to OBS. This process occupies network bandwidth.

- The query latency of a dumped index will increase. During aggregation, the latency of processing complex queries and reading a large volume of data is long.
- It cannot be unfrozen. That is, a read-only index cannot be changed to writable.
- After the freezing is complete, the index data in your local disks will be deleted.

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. Click **Dev Tools** in the navigation tree on the left.
4. Run the following command to freeze a specified index and dump it to OBS:

```
POST ${index_name}/_freeze_low_cost
```

Table 11-1 Parameter description

Parameter	Description
index_name	Name of the index to be frozen.

Information similar to the following is displayed:

```
{
  "freeze_uuid": "pdsRgUtStymVDWR_HoTGFw"
}
```

Table 11-2 Response parameter

Parameter	Description
freeze_uuid	After an index freezing request is submitted, an asynchronous job will be started. The request returns the asynchronous job ID, which can be used to query the progress of the asynchronous job.

NOTE

After an index freezing request is submitted, data cannot be written to the index. During the index freezing, query requests are not affected. After the freezing is complete, the index is closed and then opened. During this period, the index cannot be queried, and the cluster may be in the **red** status for a short time. The index is restored after being opened.

5. Run the following command to check the freezing task progress:

```
GET _freeze_low_cost_progress/${freeze_uuid}
```

Table 11-3 Parameter description

Parameter	Description
freeze_uuid	Asynchronous task ID, which is obtained in 4 .

Information similar to the following is displayed:

```
{
  "stage": "STARTED",
  "shards_stats": {
    "INIT": 0,
    "FAILURE": 0,
    "DONE": 0,
    "STARTED": 3,
    "ABORTED": 0
  },
  "indices": {
    "data1": [
      {
        "uuid": "70S-G1-tRke2jHZPlckexg",
        "index": {
          "name": "data1",
          "index_id": "4b5PHXJITLaS6AurlmfQ9A",
          "shard": 2
        },
        "start_ms": 1611972010852,
        "end_ms": -1,
        "total_time": "10.5s",
        "total_time_in_millis": 10505,
        "stage": "STARTED",
        "failure": null,
        "size": {
          "total_bytes": 3211446689,
          "finished_bytes": 222491269,
          "percent": "6.0%"
        },
        "file": {
          "total_files": 271,
          "finished_files": 12,
          "percent": "4.0%"
        },
        "rate_limit": {
          "paused_times": 1,
          "paused_nanos": 946460970
        }
      }
    ],
    {
      "uuid": "70S-G1-tRke2jHZPlckexg",
      "index": {
        "name": "data1",
        "index_id": "4b5PHXJITLaS6AurlmfQ9A",
        "shard": 0
      },
      "start_ms": 1611972010998,
      "end_ms": -1,
      "total_time": "10.3s",
      "total_time_in_millis": 10359,
      "stage": "STARTED",
      "failure": null,
      "size": {
        "total_bytes": 3221418186,
        "finished_bytes": 272347118,
        "percent": "8.0%"
      },
      "file": {
```

```

    "total_files" : 372,
    "finished_files" : 16,
    "percent" : "4.0%"
  },
  "rate_limit" : {
    "paused_times" : 5,
    "paused_nanos" : 8269016764
  }
},
{
  "uuid" : "7OS-G1-tRke2jHZPlckexg",
  "index" : {
    "name" : "data1",
    "index_id" : "4b5PHXJITLaS6AurlmfQ9A",
    "shard" : 1
  },
  "start_ms" : 1611972011021,
  "end_ms" : -1,
  "total_time" : "10.3s",
  "total_time_in_millis" : 10336,
  "stage" : "STARTED",
  "failure" : null,
  "size" : {
    "total_bytes" : 3220787498,
    "finished_bytes" : 305789614,
    "percent" : "9.0%"
  },
  "file" : {
    "total_files" : 323,
    "finished_files" : 14,
    "percent" : "4.0%"
  },
  "rate_limit" : {
    "paused_times" : 3,
    "paused_nanos" : 6057933087
  }
}
]
}
}

```

Table 11-4 Response parameters

Parameter	Description
stage	Status. Its value can be: <ul style="list-style-type: none"> ● INIT: The instance has just started or is being initialized. ● FAILURE: failed ● DONE: complete ● STARTED: started ● ABORTED: Canceled. This field is reserved.
shards_stats	Numbers of shards in each state.
indices	Index status details.

Table 11-5 Return values of **indices**

Parameter	Description
uuid	UUID of the freezing operation
index	Index and shard information
start_ms	Start time
end_ms	End time. If no end time is specified, the value -1 is displayed.
total_time	Time spent
total_time_in_millis	Time spent, in milliseconds
stage	Status of the current shard.
failure	Failure cause. If no failure occurs, null is displayed.
size.total_bytes	Size of files to be frozen, in bytes
size.finished_bytes	Frozen bytes
size.percent	Percentage of frozen bytes
file.total_bytes	Number of files to be frozen
file.finished_bytes	Number of frozen files
file.percent	Percentage of frozen files
rate_limit.paused_times	Number of times that freezing is suspended due to rate limit
rate_limit.paused_nanos	Duration of freezing task suspension due to rate limit, in nanoseconds

The following parameters are added to a frozen index. For details, see [Table 11-6](#).

Table 11-6 Frozen index parameters

Parameter	Description
index.frozen_low_cost	Whether an index is frozen. The value is true .
index.blocks.write	Whether data writing is denied in a frozen index. The value is true .
index.store.type	Storage type of an index. The value is obs .

- After an index is frozen, its data will be cached. Run the following command to check the current cache status: For details about the cache, see [Configuring Cache](#).

```
GET _frozen_stats
GET _frozen_stats/${node_id}
```

Table 11-7 Parameter description

Parameter	Description
node_id	Node ID, which can be used to obtain the cache status of a node.

Information similar to the following is displayed:

```
{
  "_nodes" : {
    "total" : 3,
    "successful" : 3,
    "failed" : 0
  },
  "cluster_name" : "css-zzz1",
  "nodes" : {
    "7uwKO38RRoaON37YsXhCYw" : {
      "name" : "css-zzz1-ess-esn-2-1",
      "transport_address" : "10.0.0.247:9300",
      "host" : "10.0.0.247",
      "ip" : "10.0.0.247",
      "block_cache" : {
        "default" : {
          "type" : "memory",
          "block_cache_capacity" : 8192,
          "block_cache_blocksize" : 8192,
          "block_cache_size" : 12,
          "block_cache_hit" : 14,
          "block_cache_miss" : 0,
          "block_cache_eviction" : 0,
          "block_cache_store_fail" : 0
        }
      }
    },
    "obs_stats" : {
      "list" : {
        "obs_list_count" : 17,
        "obs_list_ms" : 265,
        "obs_list_avg_ms" : 15
      },
      "get_meta" : {
        "obs_get_meta_count" : 79,
        "obs_get_meta_ms" : 183,
        "obs_get_meta_avg_ms" : 2
      },
      "get_obj" : {
        "obs_get_obj_count" : 12,
        "obs_get_obj_ms" : 123,
        "obs_get_obj_avg_ms" : 10
      },
      "put_obj" : {
        "obs_put_obj_count" : 12,
        "obs_put_obj_ms" : 2451,
        "obs_put_obj_avg_ms" : 204
      },
      "obs_op_total" : {
        "obs_op_total_ms" : 3022,
        "obs_op_total_count" : 120,
        "obs_op_avg_ms" : 25
      }
    },
    "reader_cache" : {
      "hit_count" : 0,
```



```

"miss_count" : 1,
"load_success_count" : 1,
"load_exception_count" : 0,
"total_load_time" : 291194714,
"eviction_count" : 0
}
},
"73EDpEqoQES749umJqxOzQ" : {
"name" : "css-zzz1-ess-esn-3-1",
"transport_address" : "10.0.0.201:9300",
"host" : "10.0.0.201",
"ip" : "10.0.0.201",
"block_cache" : {
"default" : {
"type" : "memory",
"block_cache_capacity" : 8192,
"block_cache_blocksize" : 8192,
"block_cache_size" : 12,
"block_cache_hit" : 14,
"block_cache_miss" : 0,
"block_cache_eviction" : 0,
"block_cache_store_fail" : 0
}
},
"obs_stats" : {
"list" : {
"obs_list_count" : 17,
"obs_list_ms" : 309,
"obs_list_avg_ms" : 18
},
"get_meta" : {
"obs_get_meta_count" : 79,
"obs_get_meta_ms" : 216,
"obs_get_meta_avg_ms" : 2
},
"get_obj" : {
"obs_get_obj_count" : 12,
"obs_get_obj_ms" : 140,
"obs_get_obj_avg_ms" : 11
},
"put_obj" : {
"obs_put_obj_count" : 12,
"obs_put_obj_ms" : 1081,
"obs_put_obj_avg_ms" : 90
},
"obs_op_total" : {
"obs_op_total_ms" : 1746,
"obs_op_total_count" : 120,
"obs_op_avg_ms" : 14
}
},
"reader_cache" : {
"hit_count" : 0,
"miss_count" : 1,
"load_success_count" : 1,
"load_exception_count" : 0,
"total_load_time" : 367179751,
"eviction_count" : 0
}
},
"EF8WoLCUQbqJl1Pkqo9-OA" : {
"name" : "css-zzz1-ess-esn-1-1",
"transport_address" : "10.0.0.18:9300",
"host" : "10.0.0.18",
"ip" : "10.0.0.18",
"block_cache" : {
"default" : {
"type" : "memory",
"block_cache_capacity" : 8192,

```


- Run the following command to check all the frozen indexes:

```
GET _cat/freeze_indices
```

Information similar to the following is displayed:

```
green open data2 0bNtxWDtRbOSkS4JYaUgMQ 3 0 5 0 7.9kb 7.9kb
green open data3 oYMLvw31QnyasqUNuyP6RA 3 0 51 0 23.5kb 23.5kb
```

 **NOTE**

The parameters and return values of this command are the same as those of `_cat/indices` of Elasticsearch.

11.1.3 Configuring Cache

After data is dumped to OBS, some data is cached to reduce access to OBS and improve Elasticsearch query performance. Data that is requested for the first time is obtained from OBS. The obtained data is cached in the memory. In subsequent queries, the system searches for data in the cache first. Data can be cached in memory or files.

Elasticsearch accesses different files in different modes. The cache system supports multi-level cache and uses blocks of different sizes to cache different files. For example, a large number of small blocks are used to cache .fdx and .tip files, and a small number of large blocks are used to cache .fdt files.

Table 11-8 Cache configurations

Parameter	Type	Description
low_cost.obs.blockcache.names	Array	The cache system supports multi-level cache for data of different access granularities. This configuration lists the names of all caches. If this parameter is not set, the system has a cache named default . To customize the configuration, ensure there is a cache named default . Default value: default
low_cost.obs.blockcache.<NAME>.type	ENUM	Cache type, which can be memory or file . If it is set to memory , certain memory will be occupied. If it is set to file , cache will be stored in disks. You are advised to use ultra-high I/O disks to improve cache performance. Default value: memory
low_cost.obs.blockcache.<NAME>.blockshift	Integer	Size of each block in the cache. Its value is the number of bytes shifted left. For example, if this parameter is set to 16 , the block size is 2¹⁶ bytes, that is, 65536 bytes (64 KB). Default value: 13 (8 KB)
low_cost.obs.blockcache.<NAME>.bank.count	Integer	Number of cache partitions. Default value: 1

Parameter	Type	Description
low_cost.obs.blockcache.<NAME>.number.blocks.perbank	Integer	Number of blocks included in each cache partition. Default value: 8192
low_cost.obs.blockcache.<NAME>.exclude.file.types	Array	Extensions of files that are not cached. If the extensions of certain files are neither in the exclude list nor in the include list, they are stored in the default cache.
low_cost.obs.blockcache.<NAME>.file.types	Array	Extensions of cached files. If the extensions of certain files are neither in the exclude list nor in the include list, they are stored in the default cache.

The following is a common cache configuration. It uses two levels of caches, **default** and **large**. The **default** cache uses 64 KB blocks and has a total of 30 x 4096 blocks. It is used to cache files except .fdt files. The **large** cache uses 2 MB blocks and contains 5 x 1000 blocks. It is used to cache .fdx, .dvd, and .tip files.

```
low_cost.obs.blockcache.names: ["default", "large"]
low_cost.obs.blockcache.default.type: file
low_cost.obs.blockcache.default.blockshift: 16
low_cost.obs.blockcache.default.number.blocks.perbank: 4096
low_cost.obs.blockcache.default.bank.count: 30
low_cost.obs.blockcache.default.exclude.file.types: [".fdt"]

low_cost.obs.blockcache.large.type: file
low_cost.obs.blockcache.large.blockshift: 21
low_cost.obs.blockcache.large.number.blocks.perbank: 1000
low_cost.obs.blockcache.large.bank.count: 5
low_cost.obs.blockcache.large.file.types: [".fdx", ".dvd", ".tip"]
```

Table 11-9 Other parameters

Parameter	Type	Description
index.frozen.obs.max_bytes_per_sec	String	Maximum rate of uploading files to OBS during freezing. It takes effect immediately after you complete configuration. Default value: 150MB
low_cost.obs.index.upload.threshold.use.multipart	String	If the file size exceeds the value of this parameter during freezing, the multipart upload function of OBS is used. Default value: 1GB

Parameter	Type	Description
index.frozen.reader.cache.expire.duration.seconds	Integer	Timeout duration. To reduce the heap memory occupied by frozen indexes, the reader caches data for a period of time after the index shard is started, and stops caching after it times out. Default value: 300s
index.frozen.reader.cache.max.size	Integer	Maximum cache size. Default value: 100

11.1.4 Enhanced Cold Data Query Performance

Context

When you query data on the **Discover** page of Kibana for the first time, all data needs to be obtained from OBS because there is no cache. If a large number of documents are returned, it takes a long time to obtain the corresponding time fields and file metadata from OBS. To accelerate queries the first time they run on the **Discover** page, you can cache data locally.

Prerequisites

This feature is available in clusters of versions 7.6.2 and 7.10.2 created after February 2023.

API for Querying Cold Data from Local Cache

This API can be used to query the cold data from local cache.

Example request:

```
GET /_frozen_stats/local_cache
GET /_frozen_stats/local_cache/{nodeId}
```

Response example:

```
{
  "_nodes" : {
    "total" : 1,
    "successful" : 1,
    "failed" : 0
  },
  "cluster_name" : "elasticsearch",
  "nodes" : {
    "6by3lPy1R3m55Dcq3liK8Q" : {
      "name" : "node-1",
      "transport_address" : "127.0.0.1:9300",
      "host" : "127.0.0.1",
      "ip" : "127.0.0.1",
      "local_cache" : {
        "get_stats" : {
          "get_total_count" : 562,
          //Total number of times data was retrieved from the local
          cold data cache.
        }
      }
    }
  }
}
```


Configuring Parameters

Configuration Item	Type	Unit	Value Range	Scope	Can Be Dynamically Modified	Description
low_cost.local_cache.max_capacity	Integer	-	The value ranges from 10 to 5000. The default value is 500 .	node	Yes	<p>Maximum number of available cold data caches on a node. Each shard corresponds to a cache object.</p> <p>NOTE</p> <ul style="list-style-type: none"> If the heap memory usage remains high, decrease the value. If the value of load_overflow_count keeps increasing rapidly, increase the value.
index.low_cost.local_cache.threshold	Integer	%	The value ranges from 0 to 100. The default value is 50 .	index	Yes	<p>Threshold for enabling the local cache of cold data.</p> <p>NOTE</p> <ul style="list-style-type: none"> If the percentage of date fields is less than the value of this parameter, the cold data of the date type will be cached locally. Otherwise, this parameter is not used. If the date fields of the current index occupy most of the data volume of the current index, you are not advised to use this function.

Configuration Item	Type	Unit	Value Range	Scope	Can Be Dynamically Modified	Description
index.low_cost.local_cache.evict_time	String	Days	The value ranges from 1d to 365d. The default value is 30d .	index	Yes	<p>Wait time before cold data is deleted from local cache. The value is determined based on index.frozen_date (time when the freezing is successful).</p> <p>NOTE</p> <ul style="list-style-type: none"> For indexes that have been frozen in old clusters and do not have index.frozen_date specified, the value of this parameter is determined based on the index creation time. You are advised to adjust the deletion time based on the disk usage to avoid high disk usage.

Modifying Parameters

- Run the following command to modify **low_cost.local_cache.max.capacity**:

```
PUT _cluster/settings
{
  "persistent": {
    "low_cost.local_cache.max.capacity":1000
  }
}
```
- Run the following command to modify **index.low_cost.local_cache.threshold**:

```
PUT es_write_pref2-00000000021/_settings
{
  "index.low_cost.local_cache.threshold":20
}
```
- Run the following command to modify **index.low_cost.local_cache.evict_time**:

```
PUT es_write_pref2-00000000021/_settings
{
  "index.low_cost.local_cache.evict_time":"7d"
}
```

11.2 Flow Control

11.2.1 Flow Control

11.2.1.1 Context

Feature Description

CSS can control traffic at the node level. You can configure the blacklist and whitelist, the maximum concurrent HTTP connections, and the maximum HTTP connections for a node. You can also configure the maximum heap memory used by specific request paths, the maximum CPU usage, and block access in one click, and collect statistics on node access IP addresses and URIs. Each function has an independent control switch, which is disabled by default. To restore default values of parameters, set them to **null**.

If flow control is enabled, requests will be blocked at the entry, which relieves the cluster pressure in high-concurrency scenario and avoids unavailability issues.

- **HTTP/HTTPS flow control:**
 - You can control client IP address access by setting IP addresses and subnets in HTTP/HTTPS blacklist or whitelist. If an IP address is in the blacklist, the client is disconnected and all its request are rejected. Whitelist rules take precedence over blacklist rules. If a client IP address exists in both the blacklist and whitelist, the client request will not be rejected.
 - HTTP/HTTPS concurrent connection flow control limits the total number of HTTP connections to a node per second.
 - HTTP/HTTPS new connection flow control limits the number of new connections to a node.
- Memory flow control limits request paths based on the node heap memory. You can configure memory flow control whitelist, global memory usage threshold, and heap memory threshold for a single path. Global memory flow control threshold takes precedence over the memory threshold of a single path. Paths in the whitelist will not be blocked in memory flow control.
- You can configure the global path whitelist for flow control as required when you need to use custom plug-ins.
- Request sampling can record the number of access requests from client IP addresses and the request paths of sampled users. Based on the statistics, you can identify the access traffic of client IP addresses and analyze the access traffic of request paths.
- Flow control provides an independent API for viewing traffic statistics and records the number of times the API is triggered. You can evaluate the flow control threshold and analyze the cluster load based on the statistics.
- Access logs record the URLs and bodies of HTTP/HTTPS requests received by nodes within a period of time. You can analyze the current traffic pressure based on the access logs.
- You can configure the node CPU usage threshold to limit the accessed traffic on a single node.
- One-click access blocking can block all the access traffic of a node, excluding the traffic from Kibana and Elasticsearch monitor APIs.

Constraints

- Currently, only versions 7.6.2 support the flow control feature.
- Flow control may affect the performance of some nodes.
- If flow control is enabled, user requests that exceed the flow control threshold will be rejected.
- Memory flow control and CPU flow control are based on request paths. The length and number of paths cannot be too large, or the cluster performance will be affected.

11.2.1.2 HTTP/HTTPS Flow Control

Context

You can run commands in Kibana to enable or disable HTTP/HTTPS flow control for your cluster. The command parameters are as follows.

Table 11-10 HTTP/HTTPS flow control parameters

Parameter	Type	Description
flowcontrol.http.enabled	Boolean	Whether to enable HTTP/HTTPS flow control. This function is disabled by default. Enabling it may affect node access performance. Value: true or false Default value: false
flowcontrol.http.allow	List<String>	IP address whitelist. It can contain multiple IP addresses and masks, or an IP address list. Use commas (,) to separate multiple values. Example: <i>xx.xx.xx.xx/24,xx.xx.xx.xx/24</i> , or <i>xx.xx.xx.xx.xx,xx.xx.xx</i> . The default value is null.
flowcontrol.http.deny	List<String>	IP address blacklist. Multiple IP addresses and masks or an IP address list can be configured. Use commas (,) to separate multiple IP addresses and masks. The default value is null.
flowcontrol.http.concurrent	Integer	Maximum concurrent HTTP/HTTPS connections. Default value: Number of available cores on a node x 400

Parameter	Type	Description
flowcontrol.http.newconnect	Integer	Maximum new connections that can be created for HTTP/HTTPS requests per second. Default value: Number of available cores on a node x 200
flowcontrol.http.warmup_period	Integer	Time required for the HTTP/HTTPS connection setup speed to reach the maximum. If flowcontrol.http.newconnect is set to 100 and flowcontrol.http.warmup_period is set to 5000ms , it indicates the system can set up 100 connections per second in 5 seconds. Value range: 0-10000 Unit: ms Default value: 0

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation pane on the left, choose **Dev Tools** and run commands to enable or disable HTTP/HTTPS flow control.

- Enabling HTTP/HTTPS flow control for a node

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.http.enabled": true,
    "flowcontrol.http.allow": ["192.168.0.1/24", "192.168.2.1/24"],
    "flowcontrol.http.deny": "192.168.1.1/24",
    "flowcontrol.http.concurrent": 1000,
    "flowcontrol.http.newconnect": 1000,
    "flowcontrol.http.warmup_period": 0
  }
}
```

NOTE

If all parameters are set to **null**, they will be restored to default values.

- Disabling HTTP/HTTPS flow control for a node

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.http.enabled": false
  }
}
```

11.2.1.3 Memory Flow Control

Context

Elasticsearch provides a circuit breaker, which will terminate requests if the memory usage exceeds its threshold. However, Elasticsearch does not check the heap memory usage when an API is called, and does not allow users to configure the threshold for a single request. In this case, memory usage can only be calculated during request processing, which may lead to frequent circuit breaking and cannot avoid heap memory waste. To solve this problem, CSS checks the heap memory usage when receiving REST requests, blocking excess API requests and protecting nodes. You can configure global memory flow control, or configure the request path and heap memory threshold for a specific request path. Before a request is processed, the system checks the configured heap memory threshold. If the threshold is exceeded, the request path will be blocked.

 **NOTE**

- Memory flow control may affect request processing performance.
- If the memory flow control is enabled, some Kibana search requests may fail.
- If memory flow control is enabled in Elasticsearch 5.5.1, `_mget` requests will be blocked and Kibana access will be abnormal. You can add `_mget` requests to the request whitelist to avoid this problem.

The following table describes memory flow control parameters.

Table 11-11 Memory flow control parameters

Parameter	Type	Description
flowcontrol.memory.enabled	Boolean	Whether to enable memory flow control. This function is disabled by default. Enabling memory flow control may slightly affect node performance. Value: true or false Default value: false

Parameter	Type	Description
flowcontrol.memory.allow_path	List<String>	<p>Request path whitelist for memory flow control.</p> <p>Whitelisted paths are blocked in memory flow control. Wildcard characters are supported. By default, query APIs controlled by the cluster are not blocked in memory flow control. This prevents the failure to query cluster information when the memory usage reaches the threshold.</p> <p>Example:</p> <ul style="list-style-type: none"> • "flowcontrol.memory.allow_path": "/index/_search", • "flowcontrol.memory.allow_path": "/index*/_search", • "flowcontrol.memory.allow_path": ["/index/_search", "/index1/_bulk"], <p>A maximum of 10 paths can be configured. A path can contain up to 32 characters.</p> <p>The default value is null.</p>
flowcontrol.memory.heap_limit	String	<p>Maximum global heap memory usage of a node. The value cannot be less than 10% of the heap memory.</p> <p>Value range: 10%–100%</p> <p>Default value: 90%</p>

Parameter	Type	Description
flowcontrol.memory.*.filter_path	String	<p>Paths under memory flow control. The default value is **, indicating all paths. If flowcontrol.memory.heap_limit is configured and flowcontrol.memory.*.filter_path is not, it indicates that all the paths, except those in the whitelist, are under control. The whitelist takes precedence over the single-path rule. If a path is specified in both flowcontrol.memory.allow_path and flowcontrol.memory.*.filter_path, the requests from the path will be allowed.</p> <p>For example, if flowcontrol.memory.allow_path and flowcontrol.memory.*.filter_path are both set to abc/_search, then abc/_search will not be under flow control.</p> <p>Maximum length: 32 characters</p>
flowcontrol.memory.*.heap_limit	String	<p>Heap memory usage threshold of request paths. If the heap memory usage exceeds the threshold, flow control will be triggered.</p> <p>Value range: 0–100% Default value: 90%</p>

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation pane on the left, choose **Dev Tools** and run commands to enable or disable memory flow control.

- Enabling memory flow control

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.memory.enabled": true,
    "flowcontrol.memory.allow_path": "/index/_search",
    "flowcontrol.memory.heap_limit": "85%"
  }
}
```

- Enabling memory flow control for a request path

Configure the heap memory usage threshold for a request path. You can configure the priorities of such threshold rules.

```
PUT /_cluster/settings
{
```

```
"persistent": {
  "flowcontrol.memory.enabled": true,
  "flowcontrol.memory": {
    "flowcontrol_search": {
      "filter_path": "index1/_search",
      "heap_limit": "50%"
    },
    "flowcontrol_bulk": {
      "filter_path": "index*/_bulk",
      "heap_limit": "50%"
    }
  }
}
```

- Deleting the memory flow control configuration of a request path

PUT /_cluster/settings

```
{
  "persistent": {
    "flowcontrol.memory.enabled": true,
    "flowcontrol.memory": {
      "flowcontrol_search": {
        "filter_path": null,
        "heap_limit": null
      }
    }
  }
}
```

- Disabling cluster memory flow control

PUT /_cluster/settings

```
{
  "persistent": {
    "flowcontrol.memory.enabled": false
  }
}
```

11.2.1.4 Global Path Whitelist for Flow Control

Context

The following table describes the global path whitelist parameters for flow control.

Table 11-12 Global path whitelist parameters for flow control

Parameter	Type	Description
flowcontrol.path.white_list	List<String>	<p>Paths that are not under flow control. These paths are not affected by memory flow control, CPU flow control, or one-click blocking; but are under IP address-based flow control.</p> <p>A maximum of 10 paths can be configured. A path can contain up to 32 characters.</p> <p>The default value is null.</p> <p>NOTE You are advised not to configure this parameter, unless required by plug-ins.</p>

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation tree on the left, choose **Dev Tools**. Run the following command to configure the global path whitelist for flow control:

```
PUT _cluster/settings
{
  "persistent": {
    "flowcontrol.path.white_list": "xxxx"
  }
}
```

11.2.1.5 Request Sampling

Context

Request sampling can record the access IP addresses, the number of accessed nodes, request paths, request URLs, and request bodies, which can be used to trace the IP addresses and paths of clients that have sent a large number of access requests.

The following table describes request sampling parameters.

Table 11-13 Request sampling parameters

Parameter	Type	Description
flowcontrol.statics.enabled	Boolean	Whether to enable request sampling. Request sampling may affect node performance. Value: true or false Default value: false
flowcontrol.statics.threshold	Integer	Number of recent access requests whose statistics are collected. The value 100 indicates that statistics will be collected on the 100 IP addresses and 100 URLs that are most frequently accessed. Minimum value: 10 Maximum value: 1000 Default value: 100
flowcontrol.statics.sample_frequency	Integer	Path sampling frequency. If this parameter is set to 100 , samples are collected from every 100 requests. Minimum value: 50 Default value: 100

 NOTE

- The IP address statistics and URL sampling statistics are cached based on their access time. If the cache space reaches the threshold (**flowcontrol.statics.threshold**), the records of the earliest access will be deleted.
- In URL sampling, an access path is uniquely identified by its URL hash.

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation pane on the left, choose **Dev Tools** and run commands to enable or disable sampling.

- Enabling sampling

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.statics.enabled": true,
    "flowcontrol.statics.threshold": 100,
    "flowcontrol.statics.sample_frequency": 50
  }
}
```

- Disabling sampling

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.statics.enabled": false
  }
}
```

11.2.1.6 Flow Control

Flow control can be implemented via an independent API.

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation pane on the left, choose **Dev Tools** and run the commands to query traffic control information.

- Check the traffic control status of all nodes.

```
GET /_nodes/stats/filter
```

- View the traffic control status of a specific node.

```
GET /_nodes/{nodeId}/stats/filter
```

{nodeId} indicates the ID of the node you want to check.

Example response:

```
{
  "_nodes" : {
    "total" : 1,
    "successful" : 1,
    "failed" : 0
  },
  "cluster_name" : "css-flowcontroller",
  "nodes" : {
    "EIBRNCMbTj6L1C-Wke-Dnw" : {
      "name" : "css-flowcontroller-ess-esn-1-1",
```

```

"host" : "10.0.0.133",
"timestamp" : 1613979513747,
"flow_control" : {
  "transport" : {
    "concurrent_req" : 0,
    "rejected_concurrent" : 0,
    "rejected_new" : 0,
    "rejected_deny" : 0
  },
  "http" : {
    "concurrent_req" : 0,
    "rejected_concurrent" : 0,
    "rejected_new" : 0,
    "rejected_deny" : 0
  },
  "memory" : {
    "memory_allow" : 41,
    "memory_rejected" : 0
  },
  "cpu" : {
    "rejected_cpu" : 0
  }
}
"ip_address" : [
  {
    "ip" : "/10.0.0.198",
    "count" : 453
  },
  {
    "ip" : "/198.19.49.1",
    "count" : 42
  }
],
"url_sample" : [
  {
    "url" : "/*/_search?pretty=true",
    "method" : "GET",
    "remote_address" : "/10.0.0.198:16763",
    "count" : 1
  }
]
}
}
}

```

In the response, the information of each node is separated. The **http** field records the numbers of concurrent connections and new connections. The **memory** records memory flow control statistics. The **ip_address** field records the recent client IP addresses that are accessed most recently. The **url_sample** field records the recent URLs that are requested most frequently. The **cpu** field records CPU flow control statistics.

Table 11-14 Response parameters

Parameter	Description
concurrent_req	Number of TCP connections of a node, which is recorded no matter whether flow control is enabled. This value is similar to the value of current_open of the GET /_nodes/stats/http API but is smaller, because whitelisted IP addresses and internal node IP addresses are not counted.

Parameter	Description
rejected_concurrent	Number of concurrent connections rejected during HTTP flow control. This value is not cleared when HTTP flow control is disabled.
rejected_new	Number of new connections rejected during HTTP flow control. This value is not cleared when HTTP flow control is disabled.
rejected_deny	Number of requests rejected based on the blacklist during HTTP flow control. This value is not cleared when HTTP flow control is disabled.
memory_allow	Number of allowed requests during memory flow control. This parameter takes effect when memory flow control is enabled, and its value is not cleared after memory flow control is disabled. The requests from the paths in the allow_path whitelist are not recorded. If allow_path is set to **, no requests are recorded.
memory_rejected	Number of rejected requests during memory flow control. This parameter takes effect when memory flow control is enabled, and its value is not cleared after memory flow control is disabled. The requests from the paths in the allow_path whitelist are not recorded. If allow_path is set to **, no requests are recorded.
rejected_cpu	Number of requests rejected when the CPU flow control threshold is exceeded. This parameter takes effect when CPU flow control is enabled, and its value is not cleared after CPU flow control is disabled.
ip_address	IP addresses and the number of requests. For details, see Table 11-15 .
url_sample	Request path sampling. The number of URLs of a request are collected based on the configured time and sampling interval. For details, see Table 11-16 .

Table 11-15 ip_address

Parameter	Description
ip	Source IP address for accessing the node.
method	Number of access requests from an IP address.

Table 11-16 url_sample

Parameter	Description
url	Request URL
method	Method corresponding to the request path
remote_address	Source IP address and port number of the request
count	How many times a path is sampled

11.2.1.7 Access Logs

Context

You can check access logs in either of the following ways:

- Enable and check access logs via an independent API. Configure the API parameters to record the access log time and size. The access log content is returned through a REST API.
- Print access logs. Your access logs are printed as files in backend logs.

Enabling the access log function may affect cluster performance.

The following table describes access log parameters.

Table 11-17 Access log parameters

Parameter	Type	Description
duration_limit	String	Duration recorded in an access log. Value range: 10 to 120 Unit: s Default value: 30
capacity_limit	String	Size of an access log. After access logging is enabled, the size of recorded requests is checked. If the size exceeds the value of this parameter, the access logging stops. Value range: 1 to 5 Unit: MB Default value: 1

NOTE

Access logging stops if either **duration_limit** or **capacity_limit** reaches the threshold.

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation pane on the left, choose **Dev Tools** and run commands to enable or disable access logs.
 - Enabling access logs for all nodes in a cluster
PUT /_access_log?duration_limit=30s&capacity_limit=1mb
 - Enabling access logs for a node in a cluster
PUT /_access_log/{nodeId}?duration_limit=30s&capacity_limit=1mb
{nodeId} indicates the ID of the node where you want to enable access logs.
4. Use APIs to check access logs.
 - API for checking the access logs of all nodes in a cluster
GET /_access_log
 - API for checking the access logs of a node in a cluster
GET /_access_log/{nodeId}
{nodeId} indicates the ID of the node where you want to enable access logs.

Example response:

```
{
  "_nodes" : {
    "total" : 1,
    "successful" : 1,
    "failed" : 0
  },
  "cluster_name" : "css-flowcontroller",
  "nodes" : {
    "8x-ZHu-wTemBQwpcGivFKg" : {
      "name" : "css-flowcontroller-ess-esn-1-1",
      "host" : "10.0.0.98",
      "count" : 2,
      "access" : [
        {
          "time" : "2021-02-23 02:09:50",
          "remote_address" : "/10.0.0.98:28191",
          "url" : "/_access/security/log?pretty",
          "method" : "GET",
          "content" : ""
        },
        {
          "time" : "2021-02-23 02:09:52",
          "remote_address" : "/10.0.0.98:28193",
          "url" : "/_access/security/log?pretty",
          "method" : "GET",
          "content" : ""
        }
      ]
    }
  }
}
```

Table 11-18 Response parameters

Parameter	Description
name	Node name

Parameter	Description
host	Node IP address
count	Number of node access requests in a statistical period
access	Details about node access requests in a statistical period For details, see Table 11-19 .

Table 11-19 access

Parameter	Description
time	Request time
remote_address	Source IP address and port number of the request
url	Original URL of the request
method	Method corresponding to the request path
content	Request content

5. Enable or disable the access log function.

All user access operation can be logged. By default, logs are recorded in the **access_log.log** file in the background. The maximum size of a log file is 250 MB, and there can be a maximum of five log files. You can back up access log files to OBS.

- Enabling access logs

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.accesslog.enabled": true
  }
}
```

- Disabling access logs

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.accesslog.enabled": false
  }
}
```

11.2.1.8 CPU Flow Control

Context

CPU flow control can be implemented based on the CPU usage of a node.

You can configure the CPU usage threshold of a node to prevent the node from breaking down due to heavy traffic. You can determine the CPU usage threshold based on the traffic threshold. If the CPU usage of a node exceeds the configured threshold, CPU flow control discards excess node requests to protect the cluster.

Traffic within the node or passing through Elasticsearch monitoring APIs are not affected.

The following table describes CPU flow control parameters.

Table 11-20 CPU flow control parameters

Parameter	Type	Description
flowcontrol.cpu.enabled	Boolean	Whether to enable CPU flow control. If this function is enabled, the node access performance may be affected. Value: true or false Default value: false
flowcontrol.cpu.percent_limit	Integer	Maximum CPU usage of a node. Value range: 0-100 Default value: 90
flowcontrol.cpu.allow_path	List	Path whitelist for CPU flow control. The paths specified in the allow_path whitelist are not under CPU flow control. The default value is null. A path can contain up to 32 characters. A maximum of 10 request paths can be configured. Wildcard characters are supported. For example, if this parameter is set to auto_*/_search , all the search requests of the indexes prefixed with auto_ are not under the flow control.
flowcontrol.cpu.*.filter_path	String	Paths under CPU flow control. Maximum length: 32 characters Example: "flowcontrol.cpu.search.filter_path": "/index/_search", "flowcontrol.cpu.search.limit": 60, The default value is ** , indicating all paths. If limit is configured and filter_path is not, it indicates that all the paths, except those in the whitelist, are under control. The whitelist takes precedence over the single-path rule. If a path is specified in both allow_path and filter_path , the requests from the path will be allowed. For example, if both filter_path and allow_path both set to abc/_search , then abc/_search will not be under flow control.

Parameter	Type	Description
flowcontrol.cpu.*.limit	Integer	CPU threshold of request paths. If the CPU usage exceeds the threshold, flow control will be triggered. Value range: 0-100 Default value: 90

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation pane on the left, choose **Dev Tools** and run commands to enable or disable memory flow control.

- Enabling CPU flow control

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.cpu.enabled": true,
    "flowcontrol.cpu.percent_limit": 80,
    "flowcontrol.cpu.allow_path": ["index/_search"]
  }
}
```

- Disabling CPU flow control

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.cpu.enabled": false
  }
}
```

11.2.1.9 One-click Traffic Blocking

You can block all traffic in one click, except the traffic that passes through O&M APIs, to handle unexpected traffic burst and quickly recover your cluster.

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation pane on the left, choose **Dev Tools** and run commands to enable or disable one-click traffic blocking.

- Enabling one-click traffic blocking

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.break.enabled": true
  }
}
```

- Disabling one-click traffic blocking

```
PUT /_cluster/settings
{
  "persistent": {
    "flowcontrol.break.enabled": false
  }
}
```



```
}  
}
```

11.3 Large Query Isolation

11.3.1 Context

The large query isolation feature allows you to separately manage large queries. You can isolate query requests that consume a large amount of memory or take a long period of time. If the heap memory usage of a node is too high, the interrupt control program will be triggered. The program will interrupt a large query based on the policies you configured and cancel the running query tasks of the query.

You can also configure a global query timeout duration. Long queries will be intercepted.

NOTE

Currently, only versions 7.6.2 support large query isolation.

11.3.2 Procedure

The large query isolation and global timeout features are disabled by default. If you enable them, the configuration will take effect immediately. Perform the following steps to configure the features:

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster, and click **Access Kibana** in the **Operation** column.
3. In the navigation pane of Kibana on the left, choose **Dev Tools**. Run the following command to enable large query isolation and global timeout features:

```
PUT _cluster/settings  
{  
  "persistent": {  
    "search.isolator.enabled": true,  
    "search.isolator.time.enabled": true  
  }  
}
```

The two features each has an independent switch and the following parameters.

Table 11-21 Parameters for large query isolation and global timeout duration

Switch	Parameter	Description
search.isolator.enabled	search.isolator.memory.task.limit search.isolator.time.management	Thresholds of a shard query task. A query task exceeding one of these thresholds is regarded as a large query task.

Switch	Parameter	Description
	search.isolator.memory.pool.limit search.isolator.memory.heap.limit search.isolator.count.limit	Resource usage thresholds in the isolation pool. If the resource usage of a query task exceeds one of these thresholds, the task will be intercepted. NOTE search.isolator.memory.heap.limit defines the limit on the heap memory consumed by write, query, and other operations of a node. If the limit is exceeded, large query tasks in the isolation pool will be interrupted.
	search.isolator.strategy search.isolator.strategy.ratio	Policy for selecting a query task in the isolation pool.
search.isolator.time.enabled	search.isolator.time.limit	Global timeout interval of query tasks.

4. Configure the large query isolation and global timeout duration separately.
 - Configure the thresholds of a shard query task. A query task exceeding one of these thresholds is regarded as a large query task.

```
PUT _cluster/settings
{
  "persistent": {
    "search.isolator.memory.task.limit": "50MB",
    "search.isolator.time.management": "10s"
  }
}
```

Table 11-22 Parameter description

Parameter	Data Type	Description
search.isolator.memory.task.limit	String	Threshold of the memory requested by a query task to perform aggregation or other operations. If the requested memory exceeds the threshold, the task will be isolated and observed. Value range: 0b to the maximum heap memory of a node Default value: 50MB NOTE You can run the following command to query the current heap memory and the maximum heap memory of a cluster: GET _cat/nodes?&h=id,ip,port,r,ramPercent,ramCurrent,heapMax,heapCurrent

Parameter	Data Type	Description
search.isolator.time.management	String	Threshold of the duration of a query. (started when cluster resources are used for query). If the duration of a query exceeds the threshold, it will be isolated and observed. Value range: \geq 0ms Default value: 10s

- Configure the resource usage thresholds in the isolation pool. If the resource usage of a query task exceeds one of these thresholds, the task will be intercepted.

PUT _cluster/settings

```
{
  "persistent": {
    "search.isolator.memory.pool.limit": "50%",
    "search.isolator.memory.heap.limit": "90%",
    "search.isolator.count.limit": 1000
  }
}
```

Table 11-23 Parameter description

Parameter	Data Type	Description
search.isolator.memory.pool.limit	String	Threshold of the heap memory percentage of the current node. If the total memory requested by large query tasks in the isolation pool exceeds the threshold, the interrupt control program will be triggered to cancel one of the tasks. Value range: 0.0 to 100.0% Default value: 50%
search.isolator.memory.heap.limit	String	Heap memory threshold of the current node. If the heap memory of the node exceeds the threshold, the interrupt control program will be triggered to cancel a large query task in the isolation pool. Value range: 0.0 to 100.0% Default value: 90%

Parameter	Data Type	Description
search.isolator.count.limit	Integer	<p>Threshold of the number of large query tasks in the current node isolation pool. If the number of observed query tasks exceeds the threshold, the interrupt control program will be triggered to stop accepting new large queries. New large query requests will be directly canceled.</p> <p>Value range: 10–50000</p> <p>Default value: 1000</p>

 NOTE

In addition to **search.isolator.memory.pool.limit** and **search.isolator.count.limit** parameters, you can configure **search.isolator.memory.task.limit** and **search.isolator.time.management** to control the number of query tasks that enter the isolation pool.

- Policy for selecting a query task in the isolation pool.

```
PUT _cluster/settings
{
  "persistent": {
    "search.isolator.strategy": "fair",
    "search.isolator.strategy.ratio": "0.5%"
  }
}
```

Parameter	Data Type	Description
search.isolator.strategy	String	<p>Policy for selecting large queries when the interrupt control program is triggered. The selected query will be interrupted.</p> <p>NOTE The large query isolation pool is checked every second until the heap memory is within the safe range.</p> <p>Values: fair, mem-first, or time-first</p> <ul style="list-style-type: none"> • mem-first: The query task that uses the most heap memory in the isolation pool is interrupted. • time-first: The query task that has been running for the longest time in the isolation pool is interrupted. • fair: If the difference between the heap memory of shard queries is smaller than <i>Maximum_heap_memory</i> x search.isolator.strategy.ratio, the query that takes the longest time should be interrupted. Otherwise, the query that uses the most heap memory is interrupted. <p>Default value: fair</p>
search.isolator.strategy.ratio	String	<p>Threshold of the fair policy. This parameter takes effect only if search.isolator.strategy is set to fair. If the difference between the memory usage of large query tasks does not exceed the threshold, the query that takes the longest time should be interrupted. If the difference between the memory usage of large query tasks exceeds the threshold, the query that uses the most memory is interrupted.</p> <p>Value range: 0.0 to 100.0% Default value: 1%</p>

- Configure the global timeout duration of query tasks.

```
PUT _cluster/settings
{
  "persistent": {
    "search.isolator.time.limit": "120s"
  }
}
```

Parameter	Data Type	Description
search.isolator.time.limit	String	Global query timeout duration. If this function is enabled, all the query tasks that exceed the specified duration will be canceled. Value range: $\geq 0\text{ms}$ Default value: 120s

11.4 Index Monitoring

11.4.1 Context

CSS monitors various metrics of the running status and change trend of cluster indexes to measure service usage and handle potential risks in a timely manner, ensuring that clusters can run stably.

During index monitoring, the **stats** information about indexes is collected and saved to the monitoring index (**monitoring-eye-css-[yyyy-mm-dd]**) of the cluster, and retained for one week by default.

Currently, only clusters of the version 7.6.2 support index monitoring.

11.4.2 Enabling Index Monitoring

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. Choose **Dev Tools** in the navigation pane on the left and run the following command to enable index monitoring:

```
PUT _cluster/settings
{
  "persistent": {
    "css.monitoring.index.enabled": "true"
  }
}
```

4. (Optional) To monitor a specific index, run the following command on the **Dev Tools** page of Kibana:

```
PUT _cluster/settings
{
  "persistent": {
    "css.monitoring.index.enabled": "true",
    "css.monitoring.index.interval": "30s",
    "css.monitoring.index.indices": ["index_name"],
    "css.monitoring.history.duration": "3d"
  }
}
```

Table 11-24 Parameter description

Parameter	Data Type	Description
css.monitoring.index.enabled	Boolean	Whether to enable index monitoring. If this parameter is set to true , the monitoring will be enabled. Default value: false
css.monitoring.index.interval	Time	Interval for collecting index monitoring data. Minimum value: 1s Default value: 10s
css.monitoring.index.indices	String	Name of an index to be monitored. By default, all indexes are monitored. You can configure specific indexes or a type of indexes to monitor. Example: <ul style="list-style-type: none"> "css.monitoring.index.indices": ["index_name"] indicates only <i>index_name</i> is monitored. "css.monitoring.index.indices": ["log_*"] indicates that only indexes starting with log_ are monitored. "css.monitoring.index.indices": ["index1", "index2"] indicates that index1 and index2 are monitored. Default value: * (indicating that all indexes are monitored)
css.monitoring.history.duration	Time	Retention period of monitoring data storage. The default period is a week. Minimum value: 1d Default value: 7d

NOTICE

Indexes starting with **monitoring-eye-css-*** are regarded as monitoring indexes and will not be monitored.

11.4.3 Checking the Index Read and Write Traffic

You can call an API to query the index read and write traffic within a period of time.

Prerequisites

A cluster has been created and **index monitoring** has been enabled.

Procedure

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster, and click **Access Kibana** in the **Operation** column.
3. Choose **Dev Tools** in the navigation pane on the left and run the following commands to query the index read and write traffic:
 - Check read and write traffic of all the indexes.
GET /_cat/monitoring
 - Check read and write traffic of a specific index.
GET /_cat/monitoring/{indexName}
{indexName} indicates the name of the index whose read and write traffic you want to check.
 - Check the read and write traffic of indexes for different periods.
GET _cat/monitoring?begin=1650099461000
GET _cat/monitoring?begin=2022-04-16T08:57:41
GET _cat/monitoring?begin=2022-04-16T08:57:41&end=2022-04-17T08:57:41

Table 11-25 Parameter description

Parameter	Mandatory	Description
begin	No	Start time (UTC time) of the monitoring data you want to view. Time format: strict_date_optional_time epoch_millis The default start time is five minutes before the current time.
end	No	End time (UTC time) of the monitoring data you want to view. Time format: strict_date_optional_time epoch_millis The default end time is the current time.

NOTE

These parameters cannot be used for system indexes, whose names start with a dot (.).

Information similar to the following is displayed:

```
index  begin          end          status pri rep init unassign docs.count docs.deleted store.size
pri.store.size delete.rate indexing.rate search.rate
test 2022-03-25T09:46:53.765Z 2022-03-25T09:51:43.767Z yellow 1 1 0 1 9 0
5.9kb 5.9kb 0/s 0/s 0/s
```

Table 11-26 Parameters in the returned information

Parameter	Description
index	Index name

Parameter	Description
begin	Start time of the monitoring data you queried.
end	End time of the monitoring data you queried.
status	Index status within the queried monitoring interval.
pri	The number of index shards within the queried monitoring interval.
rep	The number of index replicas within the queried monitoring interval.
init	The number of initialized indexes within the queried monitoring interval.
unassign	The number of unallocated indexes within the queried monitoring interval.
docs.count	The number of documents within the queried monitoring interval.
docs.deleted	The number of deleted documents within the queried monitoring interval.
store.size	Index storage size within the queried monitoring interval.
pri.store.size	Size of the primary index shard within the queried monitoring interval.
delete.rate	Number of indexes deleted per second within the queried monitoring interval.
indexing.rate	Number of indexes wrote per second within the queried monitoring interval.
search.rate	Number of indexes queried per second within the queried monitoring interval.

11.4.4 Checking Index Monitoring Information

You can check preconfigured index monitoring visualizations on the **Dashboard** and **Visualizations** pages of Kibana. You can also customize tables and charts.

Prerequisites


A cluster has been created and **index monitoring** has been enabled.

Checking Dashboard Charts

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster, and click **Access Kibana** in the **Operation** column.
3. In the navigation tree on the left, click **Dashboard**.

4. Click **[Monitoring] Index Monitoring Dashboard** to view the preconfigured dashboard.

Figure 11-1 Preconfigured dashboard charts



The preconfigured dashboard displays the number of read and write operations per second in the cluster and the top 10 indexes with the most read and write operations per second.

Table 11-27 Preconfigured charts

Chart Name	Description
[monitoring] markdown	Markdown chart, which briefly describes the dashboard content.
[monitoring] Indexing Rate (/s)	Number of documents written to a cluster per second.
[monitoring] Search Rate (/s)	Average number of queries per second in a cluster.
[monitoring] indexing rate of index for top10	Top 10 indexes with the most documents written per second.
[monitoring] search rate of index for top10	Top 10 indexes with the most queries per second.
[monitoring] total docs count	Total number of documents in a cluster.
[monitoring] total docs delete	Total number of deleted documents in a cluster.
[monitoring] total store size in bytes	Total storage occupied by documents in a cluster.
[monitoring] indices store_size for top10	Top 10 indexes that occupy the largest storage space.

Chart Name	Description
[monitoring] indices docs_count for top10	Top 10 indexes with the largest number of documents.
[monitoring] indexing time in millis of index for top10(ms)	Top 10 indexes with the longest document write latency in a unit time (ms).
[monitoring] search query time in millis of index for top10(ms)	Top 10 indexes with the longest index query time in a unit time (ms).
[monitoring] segment count of index for top10	Top 10 indexes with the largest number of index segments.
[monitoring] segment memory in bytes of index for top10	Top 10 indexes with the largest heap memory usage of index segments.

NOTICE

The index pattern of **monitoring-eye-css-*** cannot be deleted during index monitoring. Otherwise, the monitoring chart will be abnormal.

Customizing Visualizations Charts

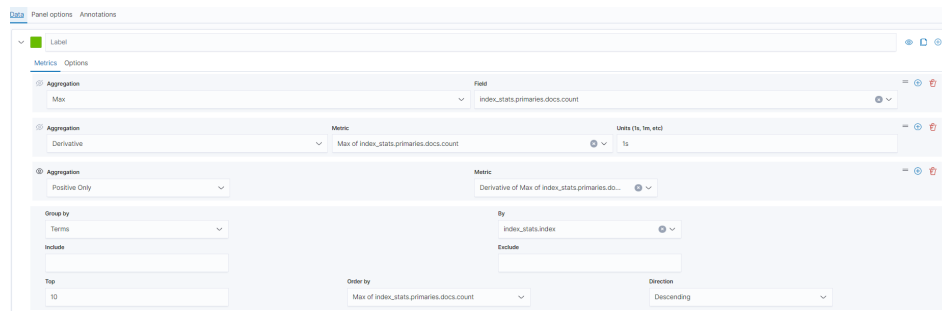
The index monitoring module periodically stores the index/stats information in the **monitoring-eyes-css** index. You can use the Kibana chart function to draw customized charts.

The following procedure describes how to check the trend of the document quantity in a chart as an example.

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. Choose **Visualize**.
4. Click **Create visualization** and select **TSVB**.
5. Set chart parameters and view the visualizations.


On the **Data** tab page, **index_stats primaries.docs.count** indicates the number of documents in the primary shard. **Derivative** indicates the difference between aggregation buckets. Set **Unit** to **1s**, visualizing network rates as "per second". Select **Positive only** to prevent negative numbers after resetting. To sort statistics by index, set **Group by** to **Terms** and **By** to **index_stats.index**. Statistics will be grouped by index name.

Figure 11-2 TSVB page



To view data in different time segments, set the aggregation interval, or the displayed data will be incomplete. On the **Panel options** tab page, set **Interval** to **1m** or **30m** to adjust the interval of **timestamp**.

Figure 11-3 Setting the interval




Importing Index Monitoring Charts

You can import or export charts on Kibana. If the index monitoring charts are not displayed, you can import the charts to Kibana again to load the monitoring view.

The following describes how to import a chart to Kibana:

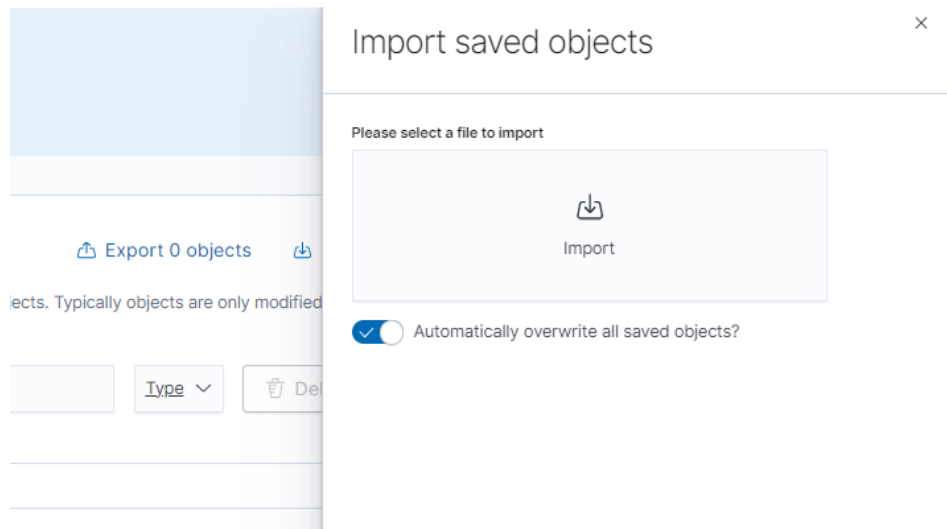
1. Create the **monitoring-kibana.ndjson** file by referring to [kibana-monitor](#).
2. Log in to Kibana and choose **Management > Stack Management > Saved objects**.

Figure 11-4 Selecting saved objects



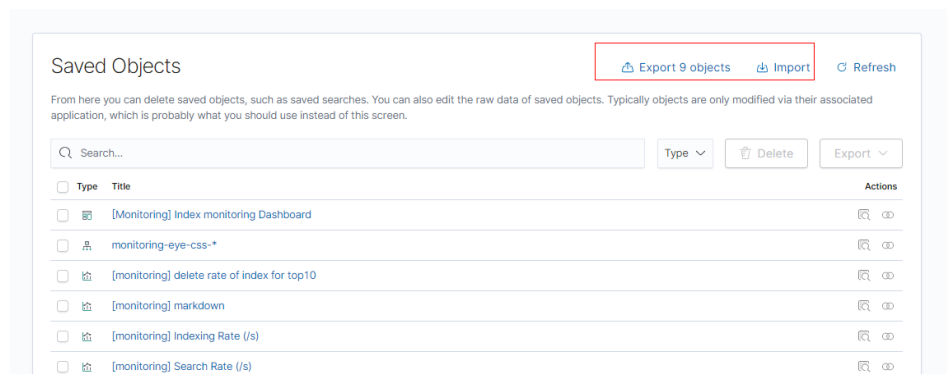
3. Click **Import** and upload the **monitoring-kibana.ndjson** file created in step 1.

Figure 11-5 Uploading a file



4. After the upload is complete, click **Done**. The index monitoring chart is successfully imported.

Figure 11-6 Successfully importing index monitoring charts



11.4.5 kibana-monitor

The configuration file content of **kibana-monitor** is as follows. You are advised to save the file as **monitoring-kibana.ndjson**.

```
{
  "attributes": {
    "description": "",
    "kibanaSavedObjectMeta": {
      "searchSourceJSON": {}
    },
    "title": "[monitoring] segment memory in bytes of index for top10",
    "uiStateJSON": {},
    "version": 1,
    "visState": {
      "title": "[monitoring] segment memory in bytes of index for top10",
      "type": "timeseries",
      "aggs": [],
      "params": {
        "id": "61ca57f0-469d-11e7-af02-69e470af7417",
        "type": "timeseries",
        "series": [
          {
            "id": "61ca57f1-469d-11e7-af02-69e470af7417",
            "color": "#68BC00",
            "split_mode": "terms",
            "split_color_mode": "kibana",
            "metrics": [
              {
                "id": "61ca57f2-469d-11e7-af02-69e470af7417",
                "type": "max",
                "field": "index_stats.total.segments.memory_in_bytes",
                "separate_axis": 0,
                "axis_position": "right",
                "formatter": "bytes",
                "chart_type": "line",
                "line_width": 1,
                "point_size": 1,
                "fill": 0.5,
                "stacked": "none",
                "label": "segments memory in bytes",
                "type": "timeseries",
                "terms_field": "index_stats.index_stats_order_by",
                "terms_order_by": "61ca57f2-469d-11e7-af02-69e470af7417",
                "time_field": "timestamp",
                "index_pattern": "monitoring-eye-css-*",
                "interval": "",
                "axis_position": "left",
                "axis_formatter": "number",
                "axis_scale": "normal",
                "show_legend": 1,
                "show_grid": 1,
                "tooltip_mode": "show_all",
                "default_index_pattern": "monitoring-eye-css-*",
                "default_timefield": "timestamp",
                "isModelInvalid": false
              }
            ],
            "id": "3ae5d820-6628-11ed-8cd7-973626cf6f70",
            "references": []
          }
        ],
        "type": "visualization",
        "updated_at": "2022-12-01T12:41:01.165Z",
        "version": "WzlwNiwYXQ=="
      }
    }
  },
  "attributes": {
    "description": "",
    "kibanaSavedObjectMeta": {
      "searchSourceJSON": {}
    },
    "title": "[monitoring] segment count of index for top10",
    "uiStateJSON": {},
    "version": 1,
    "visState": {
      "title": "[monitoring] segment count of index for top10",
      "type": "timeseries",
      "aggs": [],
      "params": {
        "id": "61ca57f3-469d-11e7-af02-69e470af7417",
        "type": "timeseries",
        "series": [
          {
            "id": "61ca57f4-469d-11e7-af02-69e470af7417",
            "color": "#68BC00",
            "split_mode": "terms",
            "split_color_mode": "kibana",
            "metrics": [
              {
                "id": "61ca57f5-469d-11e7-af02-69e470af7417",
                "type": "count",
                "field": "index_stats.index_stats_order_by",
                "separate_axis": 0,
                "axis_position": "right",
                "formatter": "number",
                "chart_type": "line",
                "line_width": 1,
                "point_size": 1,
                "fill": 0.5,
                "stacked": "none",
                "label": "segment count of index for top10",
                "type": "timeseries",
                "terms_field": "index_stats.index_stats_order_by",
                "terms_order_by": "61ca57f5-469d-11e7-af02-69e470af7417",
                "time_field": "timestamp",
                "index_pattern": "monitoring-eye-css-*",
                "interval": "",
                "axis_position": "left",
                "axis_formatter": "number",
                "axis_scale": "normal",
                "show_legend": 1,
                "show_grid": 1,
                "tooltip_mode": "show_all",
                "default_index_pattern": "monitoring-eye-css-*",
                "default_timefield": "timestamp",
                "filter": {
                  "language": "kuery",
                  "query": "monitoring-eye-css-*"
                }
              }
            ],
            "id": "3ae5d821-6628-11ed-8cd7-973626cf6f70",
            "references": []
          }
        ],
        "type": "visualization",
        "updated_at": "2022-12-01T12:41:01.165Z",
        "version": "WzlwNiwYXQ=="
      }
    }
  }
}
```

```
\"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\", \"index_pattern\":\"monitoring-eye-css-  
\"interval\":\"\", \"isModelInvalid\":false, \"series\":[ { \"axis_position\":\"right\", \"chart_type\":\"line\", \"color  
\"rgb(231,102,76,1)\", \"fill\":\"0.5\", \"formatter\":\"number\", \"id\":\"61ca57f1-469d-11e7-  
af02-69e470af7417\", \"label\":\"segment count of index for top10\", \"line_width\":1, \"metrics\":[ { \"field  
\"indices_stats.total.segments.count\", \"id\":\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max  
\"}], \"point_size\":1, \"separate_axis\":0, \"split_color_mode\":\"kibana\", \"split_mode\":\"terms\", \"stacked  
\"none\", \"terms_field\":\"index_stats.index\", \"terms_order_by\":\"61ca57f2-469d-11e7-  
af02-69e470af7417\", \"type\":\"timeseries\"}], \"show_grid\":1, \"show_legend\":1, \"time_field\":\"timestamp  
\", \"tooltip_mode\":\"show_all\", \"type\":\"timeseries\"}, { \"title\":\"[monitoring] segment count of index for  
top10\", \"type\":\"metrics\"}], \"id\":\"45d571c0-6626-11ed-8cd7-973626cf6f70\", \"references\":  
[ ], \"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlwNywyXQ==\" }  
{ \"attributes\": { \"description\":\"\", \"kibanaSavedObjectMeta\": { \"searchSourceJSON\": { } }, \"title\":\"[monitoring]  
markdown\", \"uiStateJSON\": { }, \"version\":1, \"visState\": { \"title\":\"[monitoring] markdown\", \"type  
\": \"markdown\", \"params\": { \"fontSize\":12, \"openLinksInNewTab\":false, \"markdown\": \"### Index  
Monitoring \\nThis dashboard contains default table for you to play with. You can view it, search it, and  
interact with the visualizations.\" }, \"aggs\": [ ] }, \"id\":\"b2811c70-a5f1-11ec-9a68-ada9d754c566\", \"references\":  
[ ], \"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlwOCwyXQ==\" }  
{ \"attributes\": { \"description\":\"number of document being indexing for primary and replica  
shards\", \"kibanaSavedObjectMeta\": { \"searchSourceJSON\": { } }, \"title\":\"[monitoring] Indexing Rate (/s)  
\", \"uiStateJSON\": { }, \"version\":1, \"visState\": { \"title\":\"[monitoring] Indexing Rate (/s)\", \"type\":\"metrics  
\", \"params\": { \"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\", \"type\":\"timeseries\", \"series\": [ { \"id  
\": \"61ca57f1-469d-11e7-af02-69e470af7417\", \"color\": \"rgb(0,32,188,1)\", \"split_mode\":\"everything  
\", \"metrics\": [ { \"id\":\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max\", \"field  
\": \"indices_stats_all.total.indexing.index_total\", \"unit\":\"1s\", \"id\":\"fed72db0-a5f8-11ec-  
aa10-992297d21a2e\", \"type\":\"derivative\", \"field\":\"61ca57f2-469d-11e7-af02-69e470af7417\" }, { \"unit  
\": \"\", \"id\":\"14b66420-a5f9-11ec-aa10-992297d21a2e\", \"type\":\"positive_only\", \"field\":\"fed72db0-  
a5f8-11ec-aa10-992297d21a2e\" } ], \"separate_axis\":0, \"axis_position\":\"right\", \"formatter\":\"number  
\", \"chart_type\":\"line\", \"line_width\":1, \"point_size\":1, \"fill\":\"0.5\", \"stacked\":\"none\", \"label\":\"Indexing  
Rate (/s)\", \"type\":\"timeseries\", \"split_color_mode\":\"rainbow\", \"hidden\":false }, \"time_field  
\": \"timestamp\", \"index_pattern\":\"monitoring-eye-css-\", \"interval\":\"\", \"axis_position\":\"left  
\", \"axis_formatter\":\"number\", \"axis_scale\":\"normal\", \"show_legend\":1, \"show_grid  
\":1, \"default_index_pattern\":\"monitoring-eye-css-\", \"default_timefield\":\"timestamp\", \"isModelInvalid  
\":false, \"legend_position\":\"bottom\" }, \"aggs\": [ ] }, \"id\":\"de4f8ab0-a5f8-11ec-9a68-  
ada9d754c566\", \"references\":  
[ ], \"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlwOSwyXQ==\" }  
{ \"attributes\": { \"description\":\"number of search request being executed in primary and replica  
shards\", \"kibanaSavedObjectMeta\": { \"searchSourceJSON\": { } }, \"title\":\"[monitoring] Search Rate (/s)  
\", \"uiStateJSON\": { }, \"version\":1, \"visState\": { \"title\":\"[monitoring] Search Rate (/s)\", \"type\":\"metrics  
\", \"params\": { \"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\", \"type\":\"timeseries\", \"series\": [ { \"id  
\": \"61ca57f1-469d-11e7-af02-69e470af7417\", \"color\": \"rgb(0,33,224,1)\", \"split_mode\":\"everything  
\", \"metrics\": [ { \"id\":\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max\", \"field  
\": \"indices_stats_all.total.search.query_total\", \"unit\":\"1s\", \"id\":\"b1093ac0-a5f7-11ec-  
aa10-992297d21a2e\", \"type\":\"derivative\", \"field\":\"61ca57f2-469d-11e7-af02-69e470af7417\" }, { \"unit  
\": \"\", \"id\":\"c17db930-a5f7-11ec-aa10-992297d21a2e\", \"type\":\"positive_only\", \"field\":\"b1093ac0-  
a5f7-11ec-aa10-992297d21a2e\" } ], \"separate_axis\":0, \"axis_position\":\"right\", \"formatter\":\"number  
\", \"chart_type\":\"line\", \"line_width\":1, \"point_size\":1, \"fill\":\"0.5\", \"stacked\":\"none\", \"split_color_mode  
\": \"rainbow\", \"label\":\"Search Rate (/s)\", \"type\":\"timeseries\", \"filter\": { \"query\":\"\", \"language  
\": \"kuery\" } }, \"time_field\":\"timestamp\", \"index_pattern\":\"monitoring-eye-css-\", \"interval  
\": \"\", \"axis_position\":\"left\", \"axis_formatter\":\"number\", \"axis_scale\":\"normal\", \"show_legend  
\":1, \"show_grid\":1, \"default_index_pattern\":\"monitoring-eye-css-\", \"default_timefield\":\"timestamp  
\", \"isModelInvalid\":false, \"legend_position\":\"bottom\" }, \"aggs\": [ ] }, \"id\":\"811df7a0-a5f8-11ec-9a68-  
ada9d754c566\", \"references\":  
[ ], \"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlxMCwyXQ==\" }  
{ \"attributes\": { \"description\":\"\", \"kibanaSavedObjectMeta\": { \"searchSourceJSON\": { } }, \"title\":\"[monitoring]  
total docs count\", \"uiStateJSON\": { }, \"version\":1, \"visState\": { \"title\":\"[monitoring] total docs count\", \"type  
\": \"metrics\", \"aggs\": [ ], \"params\": { \"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\", \"type\":\"timeseries  
\", \"series\": [ { \"id\":\"61ca57f1-469d-11e7-af02-69e470af7417\", \"color  
\": \"rgb(218,139,69,1)\", \"split_mode\":\"everything\", \"split_color_mode\":\"kibana\", \"metrics\": [ { \"unit  
\": \"\", \"id\":\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max\", \"field  
\": \"indices_stats_all.total.docs.count\" } ], \"separate_axis\":0, \"axis_position\":\"right\", \"formatter  
\": \"number\", \"chart_type\":\"line\", \"line_width\":1, \"point_size\":1, \"fill\":\"0.5\", \"stacked\":\"none\", \"label  
\": \"total_docs_count\", \"type\":\"timeseries\" } ], \"time_field\":\"timestamp\", \"index_pattern\":\"monitoring-  
eye-css-\", \"interval\":\"\", \"axis_position\":\"left\", \"axis_formatter\":\"number\", \"axis_scale\":\"normal  
\", \"show_legend\":1, \"show_grid\":1, \"tooltip_mode\":\"show_all\", \"default_index_pattern\":\"monitoring-  
eye-css-\", \"default_timefield\":\"timestamp\", \"isModelInvalid\":false, \"legend_position\":\"bottom  
\" } } ], \"id\":\"eea89780-664b-11ed-8cd7-973626cf6f70\", \"references\":  
[ ], \"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlxMSwyXQ==\" }  
{ \"attributes\": { \"description\":\"\", \"kibanaSavedObjectMeta\": { \"searchSourceJSON\": { } }, \"title\":\"[monitoring]  
total docs delete\", \"uiStateJSON\": { }, \"version\":1, \"visState\": { \"title\":\"[monitoring] total docs delete\", \"type
```

```
\"metrics\".\"agg\".\"params\".:{\"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\", \"type\":\"timeseries  
\".\"series\".:[{\"id\":\"61ca57f1-469d-11e7-af02-69e470af7417\", \"color  
\".\"rgba(214,191,87,1)\", \"split_mode\":\"everything\", \"split_color_mode\":\"kibana\", \"metrics\".:[{\"id  
\".\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max\", \"field  
\".\"indices_stats_all.total.docs.deleted\", \"separate_axis\":0, \"axis_position\":\"right\", \"formatter  
\".\"number\", \"chart_type\":\"line\", \"line_width\":1, \"point_size\":1, \"fill\":0.5, \"stacked\":\"none\", \"label  
\".\"total_docs_delete\", \"type\":\"timeseries\", \"hidden\":false}, {\"time_field\":\"timestamp\", \"index_pattern  
\".\"monitoring-eye-css-*, \"interval\":\"\", \"axis_position\":\"left\", \"axis_formatter\":\"number  
\".\"axis_scale\":\"normal\", \"show_legend\":1, \"show_grid\":1, \"tooltip_mode\":\"show_all  
\".\"default_index_pattern\", \"monitoring-eye-css-*, \"default_timefield\":\"timestamp\", \"isModelInvalid  
\".false, \"drop_last_bucket\":1, \"legend_position\":\"bottom  
\"}]}], \"id\":\"c7f72ae0-664f-11e7-af02-69e470af7417\", \"references\":  
[]}, {\"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlxMiwYXQ==\"}  
{\"attributes\":{\"description\":\"\", \"kibanaSavedObjectMeta\":{\"searchSourceJSON\":\"\"}, \"title\":\"[monitoring]  
total store size in bytes\", \"uiStateJSON\":\"\", \"version\":1, \"visState\":{\"\"title\":\"[monitoring] total store size in  
bytes\", \"type\":\"metrics\", \"agg\":\"\", \"params\".:{\"id\":\"61ca57f0-469d-11e7-af02-69e470af7417\", \"type  
\".\"timeseries\", \"series\".:[{\"id\":\"61ca57f1-469d-11e7-af02-69e470af7417\", \"color  
\".\"#68BC00\", \"split_mode\":\"everything\", \"split_color_mode\":\"kibana\", \"metrics\".:[{\"id  
\".\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max\", \"field  
\".\"indices_stats_all.total.store.size.in.bytes\", \"separate_axis\":0, \"axis_position\":\"right\", \"formatter  
\".\"bytes\", \"chart_type\":\"line\", \"line_width\":1, \"point_size\":1, \"fill\":0.5, \"stacked\":\"none\", \"label  
\".\"total store size in bytes\", \"type\":\"timeseries\", \"time_field\":\"timestamp\", \"index_pattern  
\".\"monitoring-eye-css-*, \"interval\":\"\", \"axis_position\":\"left\", \"axis_formatter\":\"number  
\".\"axis_scale\":\"normal\", \"show_legend\":1, \"show_grid\":1, \"tooltip_mode\":\"show_all  
\".\"default_index_pattern\", \"monitoring-eye-css-*, \"default_timefield\":\"timestamp\", \"isModelInvalid  
\".false, \"legend_position\":\"bottom\", \"background_color_rules\".:[{\"id  
\".\"7712e550-664f-11e7-af02-69e470af7417\", \"bar_color_rules\".:[{\"id  
\".\"77680a30-664f-11e7-af02-69e470af7417\", \"color\":\"#68BC00\", \"terms\".:[{\"id  
\".\"61ca57f0-469d-11e7-af02-69e470af7417\", \"type\":\"timeseries\", \"series\".:[{\"id  
\".\"61ca57f1-469d-11e7-af02-69e470af7417\", \"color\":\"#68BC00\", \"split_mode\":\"terms\", \"metrics\".:[  
[\"id\":\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max\", \"field  
\".\"index_stats.total.indexing.index_total\", \"unit\":\"1s\", \"id\":\"541ed8f0-a5ee-11ec-aa10-992297d21a2e  
\".\"type\":\"derivative\", \"field\":\"61ca57f2-469d-11e7-af02-69e470af7417\", \"unit\":\"\", \"id  
\".\"67ec1f50-a5ee-11ec-aa10-992297d21a2e\", \"type\":\"positive_only\", \"field\":\"541ed8f0-a5ee-11ec-aa10-992297d21a2e  
\".\"separate_axis\":0, \"axis_position\":\"right\", \"formatter\":\"number\", \"chart_type  
\".\"line\", \"line_width\":1, \"point_size\":1, \"fill\":0.5, \"stacked\":\"none\", \"label\":\"indexing_rate\", \"type  
\".\"timeseries\", \"split_filters\".:[{\"color\":\"#68BC00\", \"id\":\"81004200-a5ee-11ec-aa10-992297d21a2e  
\".\"filter\":{\"query\":\"\", \"language\":\"kuery\"}], \"filter\":{\"query\":\"\", \"language\":\"kuery  
\"}], \"terms_field\":\"index_stats.index\", \"terms_order_by\":\"61ca57f2-469d-11e7-  
af02-69e470af7417\", \"terms_size\":\"10\", \"terms_direction\":\"desc\", \"split_color_mode\":\"rainbow  
\"}], \"time_field\":\"timestamp\", \"index_pattern\":\"monitoring-eye-css-*, \"interval\":\"\", \"axis_position  
\".\"left\", \"axis_formatter\":\"number\", \"axis_scale\":\"normal\", \"show_legend\":1, \"show_grid  
\".\"1\", \"default_index_pattern\":\"monitoring-eye-css-*, \"default_timefield\":\"timestamp\", \"isModelInvalid  
\".false, \"tooltip_mode\":\"show_all\"}]}], \"id\":\"943b3e00-a5ef-11ec-9a68-ada9d754c566\", \"references\":  
[]}, {\"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlxNCwYXQ==\"}  
{\"attributes\":{\"description\":\"\", \"kibanaSavedObjectMeta\":{\"searchSourceJSON\":\"\"}, \"title\":\"[monitoring]  
search rate of index for top10(/s)\", \"uiStateJSON\":\"\", \"version\":1, \"visState\":{\"\"title\":\"[monitoring] search  
rate of index for top10(/s)\", \"type\":\"metrics\", \"agg\":\"\", \"params\".:{\"id\":\"61ca57f0-469d-11e7-  
af02-69e470af7417\", \"type\":\"timeseries\", \"series\".:[{\"id\":\"61ca57f1-469d-11e7-  
af02-69e470af7417\", \"color\":\"rgba(99,157,12,1)\", \"split_mode\":\"terms\", \"metrics\".:[{\"id  
\".\"61ca57f2-469d-11e7-af02-69e470af7417\", \"type\":\"max\", \"field  
\".\"index_stats.total.search.query_total\", \"unit\":\"1s\", \"id\":\"fdffad0-a5ef-11ec-aa10-992297d21a2e  
\".\"type\":\"derivative\", \"field\":\"61ca57f2-469d-11e7-af02-69e470af7417\", \"unit\":\"\", \"id  
\".\"0aaa26a0-a5f0-11ec-aa10-992297d21a2e\", \"type\":\"positive_only\", \"field\":\"fdffad0-a5ef-11ec-aa10-992297d21a2e  
\".\"separate_axis\":0, \"axis_position\":\"right\", \"formatter\":\"number\", \"chart_type  
\".\"line\", \"line_width\":1, \"point_size\":1, \"fill\":0.5, \"stacked\":\"none\", \"label\":\"search rate\", \"type  
\".\"timeseries\", \"terms_field\":\"index_stats.index\", \"terms_order_by\":\"61ca57f2-469d-11e7-  
af02-69e470af7417\", \"split_color_mode\":\"rainbow\"}], \"time_field\":\"timestamp\", \"index_pattern  
\".\"monitoring-eye-css-*, \"interval\":\"\", \"axis_position\":\"left\", \"axis_formatter\":\"number  
\".\"axis_scale\":\"normal\", \"show_legend\":1, \"show_grid\":1, \"default_index_pattern\":\"monitoring-eye-  
css-*, \"default_timefield\":\"timestamp\", \"isModelInvalid\":false, \"tooltip_mode\":\"show_all  
\"}]}], \"id\":\"ab503550-a5ef-11ec-9a68-ada9d754c566\", \"references\":  
[]}, {\"type\":\"visualization\", \"updated_at\":\"2022-12-01T12:41:01.165Z\", \"version\":\"WzlxNSwYXQ==\"}  
{\"attributes\":{\"description\":\"\", \"kibanaSavedObjectMeta\":{\"searchSourceJSON\":\"\"}, \"title\":\"[monitoring]  
indices store_size for top10\", \"uiStateJSON\":\"\", \"version\":1, \"visState\":{\"\"title\":\"[monitoring] indices
```

```
store_size for top10\","type":"metrics\","aggs":{},"params":{"id":"61ca57f0-469d-11e7-af02-69e470af7417"},"type":"timeseries"},"series":[{"id":"38474c50-a5f5-11ec-aa10-992297d21a2e","color":"#68BC00","split_mode":"terms","metrics":{"id":"38474c51-a5f5-11ec-aa10-992297d21a2e"},"type":"max"},"field":"index_stats.total.store.size_in_bytes"},"separate_axis":0,"axis_position":"right","formatter":"bytes","chart_type":"line","line_width":1,"point_size":1,"fill":0.5,"stacked":"none","label":"store_size for index"},"type":"timeseries"},"terms_field":"index_stats.index","terms_order_by":"38474c51-a5f5-11ec-aa10-992297d21a2e"},"filter":{"query":"","language":"kuery"},"split_color_mode":"rainbow"},"time_field":"timestamp"},"index_pattern":"monitoring-eye-css"},"interval":"","axis_position":"left"},"axis_formatter":"number"},"axis_scale":"normal"},"show_legend":1,"show_grid":1,"default_index_pattern":"monitoring-eye-css"},"default_timefield":"timestamp"},"isModelInvalid":false,"filter":{"query":"","language":"kuery"},"bar_color_rules":{"id":"7d9d3cb0-a5f5-11ec-aa10-992297d21a2e"},"tooltip_mode":"show_all"},"id":"c78119a0-a5f5-11ec-9a68-ada9d754c566","references":{},"type":"visualization","updated_at":"2022-12-01T12:41:01.165Z","version":"WzlxNiwYXQ=="},"attributes":{"description":"","kibanaSavedObjectMeta":{"searchSourceJSON":"","title":"[monitoring] search query time in millis of index for top10(ms)","uiStateJSON":"","version":1,"visState":{"title":"[monitoring] search query time in millis of index for top10(ms)","type":"metrics\","aggs":{},"params":{"axis_formatter":"number"},"axis_max":"","axis_min":"","axis_position":"left"},"axis_scale":"normal"},"default_index_pattern":"monitoring-eye-css"},"default_timefield":"timestamp"},"id":"61ca57f0-469d-11e7-af02-69e470af7417"},"index_pattern":"monitoring-eye-css"},"interval":"","isModelInvalid":false},"series":[{"axis_position":"right","chart_type":"line","color":"#68BC00","fill":0.5,"formatter":"number"},"id":"61ca57f1-469d-11e7-af02-69e470af7417"},"label":"index_query_time_in_millis"},"line_width":1,"metrics":{"field":"index_stats.total.search.query_time_in_millis"},"id":"61ca57f2-469d-11e7-af02-69e470af7417"},"type":"max"},"unit":"1s"},"id":"42c92b10-6645-11ed-925a-6de90846447d"},"type":"derivative"},"field":"61ca57f2-469d-11e7-af02-69e470af7417"},"point_size":1,"separate_axis":0,"split_color_mode":"kibana"},"split_mode":"terms","stacked":"none"},"terms_field":"index_stats.index"},"terms_order_by":"61ca57f2-469d-11e7-af02-69e470af7417"},"type":"timeseries"},"show_grid":1,"show_legend":1,"time_field":"timestamp"},"tooltip_mode":"show_all"},"type":"timeseries"},"background_color":null,"filter":{"query":"","language":"kuery"},"legend_position":"right"},"id":"c8109100-6627-11ed-8cd7-973626cf6f70","references":{},"type":"visualization","updated_at":"2022-12-01T12:41:01.165Z","version":"WzlxNywyXQ=="},"attributes":{"description":"","hits":0,"kibanaSavedObjectMeta":{"searchSourceJSON":{"query":{"language":"kuery"},"query":"","filter":{"optionsJSON":{"hidePanelTitles":false},"useMargins":true},"panelsJSON":[{"gridData":{"x":0,"y":0,"w":48,"h":5,"i":"971ed6c6-81b9-491b-9f08-e3ae9c382abd"},"panelIndex":"971ed6c6-81b9-491b-9f08-e3ae9c382abd"},"embeddableConfig":{"panelRefName":"panel_0"},"gridData":{"x":0,"y":5,"w":24,"h":15,"i":"5a6982e7-0c6c-4733-8a2d-e4c57cdf7397"},"panelIndex":"5a6982e7-0c6c-4733-8a2d-e4c57cdf7397"},"embeddableConfig":{"panelRefName":"panel_1"},"gridData":{"x":24,"y":5,"w":24,"h":15,"i":"662476f4-739c-4a05-858c-2ee8230cf410"},"panelIndex":"662476f4-739c-4a05-858c-2ee8230cf410"},"embeddableConfig":{"panelRefName":"panel_2"},"gridData":{"x":0,"y":20,"w":16,"h":15,"i":"d89c38e2-33f3-4592-b503-20460a6a7a57"},"panelIndex":"d89c38e2-33f3-4592-b503-20460a6a7a57"},"embeddableConfig":{"panelRefName":"panel_3"},"gridData":{"x":16,"y":20,"w":16,"h":15,"i":"1f693b49-79fa-4807-94e8-0c12f51e54f8"},"panelIndex":"1f693b49-79fa-4807-94e8-0c12f51e54f8"},"embeddableConfig":{"panelRefName":"panel_4"},"gridData":{"x":32,"y":20,"w":16,"h":15,"i":"616b143d-74e9-4dac-98ba-5849536f0fba"},"panelIndex":"616b143d-74e9-4dac-98ba-5849536f0fba"},"embeddableConfig":{"panelRefName":"panel_5"},"gridData":{"x":0,"y":35,"w":24,"h":11,"i":"cfa82f27-1b8d-49ba-a7b9-d8809d3b258c"},"panelIndex":"cfa82f27-1b8d-49ba-a7b9-d8809d3b258c"},"embeddableConfig":{"panelRefName":"panel_6"},"gridData":{"x":24,"y":35,"w":24,"h":11,"i":"135d13eb-aab6-43ca-9029-7d26e91d90e3"},"panelIndex":"135d13eb-aab6-43ca-9029-7d26e91d90e3"},"embeddableConfig":{"panelRefName":"panel_7"},"gridData":{"x":0,"y":46,"w":24,"h":11,"i":"28a77de1-9110-49e8-b273-724f880b1653"},"panelIndex":"28a77de1-9110-49e8-b273-724f880b1653"},"embeddableConfig":{"panelRefName":"panel_8"},"gridData":{"x":24,"y":46,"w":24,"h":11,"i":"80ece867-cf23-4935-bfbc-430afa51bcca"},"panelIndex":"80ece867-cf23-4935-bfbc-430afa51bcca"},"embeddableConfig":{"panelRefName":"panel_9"},"gridData":{"x":0,"y":57,"w":24,"h":11,"i":"2ba970aa-c9c4-491b-bdd3-c1b1ee9bc8d3"},"panelIndex":"2ba970aa-c9c4-491b-bdd3-c1b1ee9bc8d3"},"embeddableConfig":{"panelRefName":"panel_10"},"gridData":{"x":24,"y":57,"w":24,"h":11,"i":"f2e1b6ab-ddf7-492e-aaca-9460f11aa4aa"},"panelIndex":"f2e1b6ab-ddf7-492e-aaca-9460f11aa4aa"},"embeddableConfig":{"panelRefName":"panel_11"},"gridData":{"x":0,"y":68,"w":24,"h":11,"i":"dd14182d-d8b9-47f2-bf36-6c3a3b09586c"},"panelIndex":"dd14182d-d8b9-47f2-bf36-6c3a3b09586c"},"embeddableConfig":{"panelRefName":"panel_12"},"gridData":{"x":24,"y":68,"w":24,"h":11,"i":"a47f9333-52b7-49b7-8cac-f470cf405131"},"panelIndex":"a47f9333-52b7-49b7-8cac-f470cf405131"},"embeddableConfig":{"panelRefName":"panel_13"},"timeRestore":false,"title":"[Monitoring] Index monitoring Dashboard","version":1,"id":"524eb000-a5f2-11ec-9a68-ada9d754c566","references":{"id":"b2811c70-a5f1-11ec-9a68-ada9d754c566","name":"panel_0","type":"visualization"},"id":"de4f8ab0-a5f8-11ec-9a68-
```



```
ada9d754c566", "name": "panel_1", "type": "visualization"}, {"id": "811df7a0-a5f8-11ec-9a68-ada9d754c566", "name": "panel_2", "type": "visualization"}, {"id": "eea89780-664b-11ed-8cd7-973626cf6f70", "name": "panel_3", "type": "visualization"}, {"id": "cfbb4e20-664c-11ed-8cd7-973626cf6f70", "name": "panel_4", "type": "visualization"}, {"id": "c7f72ae0-664e-11ed-8cd7-973626cf6f70", "name": "panel_5", "type": "visualization"}, {"id": "943b3e00-a5ef-11ec-9a68-ada9d754c566", "name": "panel_6", "type": "visualization"}, {"id": "ab503550-a5ef-11ec-9a68-ada9d754c566", "name": "panel_7", "type": "visualization"}, {"id": "c78119a0-a5f5-11ec-9a68-ada9d754c566", "name": "panel_8", "type": "visualization"}, {"id": "225f6020-a5f1-11ec-9a68-ada9d754c566", "name": "panel_9", "type": "visualization"}, {"id": "17d49220-662a-11ed-8cd7-973626cf6f70", "name": "panel_10", "type": "visualization"}, {"id": "c8109100-6627-11ed-8cd7-973626cf6f70", "name": "panel_11", "type": "visualization"}, {"id": "45d571c0-6626-11ed-8cd7-973626cf6f70", "name": "panel_12", "type": "visualization"}, {"id": "3ae5d820-6628-11ed-8cd7-973626cf6f70", "name": "panel_13", "type": "visualization"}], "type": "dashboard", "updated_at": "2022-12-01T12:41:01.165Z", "version": "WzlxOCwyXQ=="}, {"exportedCount": 16, "missingRefCount": 0, "missingReferences": []}
```

11.5 Enhanced Monitoring

11.5.1 P99 Latency Monitoring

Context

The Elasticsearch community only discusses how to monitor the average latency of search requests, which cannot reflect the actual search performance of a cluster. To enhance monitoring, CSS allows you to monitor the P99 latency of search requests in clusters.

Prerequisites

Currently, only clusters of version 7.6.2 support P99 latency monitoring.

Obtaining Monitoring Information

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation tree on the left, choose **Dev Tools** and run the following command to check the P99 latency of the current cluster:

```
GET /search/stats/percentile
```

Example response:

```
{
  "overall" : {
    "1.0" : 2.0,
    "5.0" : 2.0,
    "25.0" : 6.5,
    "50.0" : 19.5,
    "75.0" : 111.0,
    "95.0" : 169.0,
    "99.0" : 169.0,
    "max" : 169.0,
    "min" : 2.0
  },
  "last_one_day" : {
    "1.0" : 2.0,
    "5.0" : 2.0,
    "25.0" : 6.5,
    "50.0" : 19.5,
    "75.0" : 111.0,
```

```
"95.0" : 169.0,
"99.0" : 169.0,
"max" : 169.0,
"min" : 2.0
},
"latest" : {
  "1.0" : 26.0,
  "5.0" : 26.0,
  "25.0" : 26.0,
  "50.0" : 26.0,
  "75.0" : 26.0,
  "95.0" : 26.0,
  "99.0" : 26.0,
  "max" : 26.0,
  "min" : 26.0
}
}
```

 **NOTE**

- In the response, **overall** indicates all the statistics that have been collected since the cluster startup, **last_one_day** indicates the statistics collected in the last day, and **latest** indicates the statistics that have been collected since the last reset.
- The calculated P99 latency is an estimation. It is more precise than the P50 latency.
- The P99 latency of a cluster is cleared and recalculated if the cluster is restarted.

Other Operations

- Define percentage.

You can run the following command to specify the percentage:

```
GET /search/stats/percentile
{
  "percents": [1, 50, 90]
}
```

- Reset the **latest** statistics.

You can run the following command to reset the **latest** statistics:

```
POST /search/stats/reset
```

Example response:

```
{
  "nodes" : {
    "css-c9c8-ess-esn-1-1" : "ok"
  }
}
```

11.5.2 HTTP Status Code Monitoring

Context

When an external system accesses Elasticsearch through the HTTP protocol, a response and the corresponding status code are returned. The open-source Elasticsearch server does not collect the status code, so users cannot monitor Elasticsearch APIs status or cluster request status. CSS allows you to monitor the HTTP status codes of clusters.

Prerequisites

Currently, only clusters of versions 7.6.2 support HTTP status code monitoring.

Obtaining Status Codes

1. Log in to the CSS management console.
2. Choose **Clusters** in the navigation pane. On the **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
3. In the navigation tree on the left, choose **Dev Tools**.
4. On the console page of **Dev Tools**, run commands based on the cluster version.

- For clusters of version 7.6.2, run the following command to obtain the status code statistics:

```
GET /_nodes/http_stats
```

Example response:

```
{
  "_nodes" : {
    "total" : 1,
    "successful" : 1,
    "failed" : 0 },
  "cluster_name" : "css-8362",
  "nodes" : {
    "F9IFdQPAPRaOJI7oL7HOxtQ" : {
      "http_code" : {
        "200" : 114,
        "201" : 5,
        "429" : 0,
        "400" : 7,
        "404" : 0,
        "405" : 0
      }
    }
  }
}
```

- For clusters of version 7.10.2, run the following command to obtain the status code statistics:

```
GET _nodes/stats/http
```

Example response:

```
{
  // ...
  "cluster_name" : "css-2985",
  "nodes" : {
    // ...
    "omvR9_W-TsGApraMApREjA" : {
      // ...
      "http" : {
        "current_open" : 4,
        "total_opened" : 37,
        "http_code" : {
          "200" : 25,
          "201" : 7,
          "429" : 0,
          "400" : 3,
          "404" : 0,
          "405" : 0
        }
      }
    }
  }
}
```

12 Monitoring

12.1 Supported Metrics

Function

This section describes CSS metrics that can be monitored by Cloud Eye as well as their namespaces and dimensions. You can use the management console or APIs provided by Cloud Eye to view the monitoring metrics and alarms generated for CSS.

Namespace

SYS.ES

Monitoring Metrics

- [Table 12-1](#) describes the monitoring metrics of CSS clusters.
- Monitored object: CSS cluster
- Monitoring period (original metric): 1 minute

NOTE

Accumulated value: The value is accumulated from the time when a node is started. After the node is restarted, the value is reset to zero and accumulated again.

Table 12-1 CSS metrics

Metric ID	Metric	Description	Value Range
status	Cluster Health Status	Health status of the monitored object	0,1,2,3 0: All primary and replica shards are allocated. Your cluster is 100% operational. 1: All primary shards are allocated, but at least one replica is missing. No data is missing, so search results will still be complete. However, high availability is compromised to some degree. If more shards disappear, you might lose data. Think of this status as a warning that should prompt investigation. 2: Data is missing and the cluster fails to work. 3: The cluster status is not obtained.
disk_util	Disk Usage	Disk usage of the monitored object. Unit: %	0-100%
max_jvm_heap_usage	Max. JVM Heap Usage	Maximum JVM heap usage of nodes in a CSS cluster. Unit: %	0-100%
max_jvm_young_gc_time	Max. JVM Young GC Duration	Maximum accumulated JVM Young GC duration of nodes in a CSS cluster. Unit: ms	≥ 0 ms

Metric ID	Metric	Description	Value Range
max_jvm_young_gc_count	Max. JVM Young GC Count	Maximum accumulated JVM Young GC count of nodes in a CSS cluster.	≥ 0
max_jvm_old_gc_time	Max. JVM Old GC Duration	Maximum accumulated JVM Old GC duration of nodes in a CSS cluster. Unit: ms	≥ 0 ms
max_jvm_old_gc_count	Max. JVM Old GC Count	Maximum accumulated JVM Old GC count of nodes in a CSS cluster.	≥ 0
total_fs_size	Total Size of File Systems	Total size of file systems in a CSS cluster. Unit: byte	≥ 0 bytes
free_fs_size	Available Size of File Systems	Available size of file systems in a CSS cluster. Unit: byte	≥ 0 bytes
max_cpu_usage	Max. CPU Usage	Maximum node CPU usage in a CSS cluster. Unit: %	0-100%
max_cpu_time_of_jvm_process	Max. CPU Time of JVM Process	Maximum accumulated CPU usage duration of node JVM processes in a CSS cluster. Unit: ms	≥ 0 ms
max_virtual_memory_size_of_jvm_process	Max. Virtual Memory Size of JVM Process	Maximum virtual memory size of node JVM processes in a CSS cluster. Unit: byte	≥ 0 bytes

Metric ID	Metric	Description	Value Range
max_current_ope ned_http_count	Current Max. Opened HTTP Connections	Maximum number of HTTP connections that are currently open for nodes in a CSS cluster.	≥ 0
max_total_opene d_http_count	Total Max. Opened HTTP Connections	Maximum number of HTTP connections that were open for nodes in a CSS cluster.	≥ 0
indices_count	Indexes	Number of indexes in a CSS cluster	≥ 0
total_shards_cou nt	Shards	Number of shards in a CSS cluster	≥ 0
primary_shards_c ount	Primary Shards	Number of primary shards in a CSS cluster	≥ 0
docs_count	Documents	Number of documents in a CSS cluster	≥ 0
docs_deleted_cou nt	Deleted Documents	Number of documents deleted in a CSS cluster	≥ 0
nodes_count	Nodes	Number of nodes in a CSS cluster	≥ 0
data_nodes_coun t	Data Nodes	Number of data nodes in a CSS cluster	≥ 0
coordinating_nod es_count	Coordinating Nodes	Number of coordinating nodes in a CSS cluster	≥ 0
master_nodes_co unt	Master Nodes	Number of master nodes in a CSS cluster	≥ 0
ingest_nodes_cou nt	Client Nodes	Number of client nodes in a CSS cluster	≥ 0

Metric ID	Metric	Description	Value Range
max_load_average	Max. Node Load	Maximum number of average queuing tasks per minute on nodes in a cluster.	≥ 0
avg_cpu_usage	Avg. CPU Usage	Average node CPU usage in a CSS cluster. Unit: %	0-100%
avg_load_average	Avg. Node Load	Average number of queuing tasks per minute on nodes in a CSS cluster.	≥ 0
avg_jvm_heap_usage	Avg. JVM Heap Usage	Average node JVM heap usage in a CSS cluster. Unit: %	0-100%
max_open_file_descriptors	Max. Open File Descriptors	Maximum number of node file descriptors that are currently open in a CSS cluster.	≥ 0
avg_open_file_descriptors	Avg. Open File Descriptors	Average number of node file descriptors that are currently open in a CSS cluster.	≥ 0
sum_max_file_descriptors	Max. Allowed File Descriptors	Maximum number of allowed node file descriptors in a CSS cluster.	≥ 0
sum_open_file_descriptors	Open File Descriptors	Number of node file descriptors that are currently open in a cluster.	≥ 0
sum_thread_pool_write_queue	Tasks in Write Queue	Number of job queues in a write thread pool	≥ 0
sum_thread_pool_search_queue	Tasks in Search Queue	Total number of queuing tasks in the search thread pools of nodes in the CSS cluster.	≥ 0

Metric ID	Metric	Description	Value Range
sum_thread_pool_force_merge_queue	Tasks in ForceMerge Queue	Total number of queuing tasks in the force merge thread pools of nodes in the CSS cluster.	≥ 0
sum_thread_pool_write_rejected	Rejected Tasks in Write Queue	Total number of rejected tasks in the write thread pools of nodes in the CSS cluster.	≥ 0
sum_thread_pool_search_rejected	Rejected Tasks in Search Queue	Total number of rejected tasks in the search thread pools of nodes in the CSS cluster.	≥ 0
sum_thread_pool_force_merge_rejected	Rejected Tasks in ForceMerge Queue	Total number of rejected tasks in the force merge thread pools of nodes in the CSS cluster.	≥ 0
max_thread_pool_search_queue	Max. Tasks in Search Queue	Maximum number of queuing tasks in the search thread pools of nodes in a CSS cluster.	≥ 0
max_thread_pool_force_merge_queue	Max. Tasks in ForceMerge Queue	Maximum number of queuing tasks in the force merge thread pools of nodes in a CSS cluster.	≥ 0
sum_thread_pool_write_threads	Size of Write Thread Pool	Total size of the write thread pools of nodes in the CSS cluster.	≥ 0
sum_thread_pool_search_threads	Size of Search Thread Pool	Total size of the search thread pools of nodes in the CSS cluster.	≥ 0
sum_thread_pool_force_merge_threads	Size of ForceMerge Thread Pool	Total size of the force merge thread pools of nodes in the CSS cluster.	≥ 0

Metric ID	Metric	Description	Value Range
avg_thread_pool_write_queue	Avg. Tasks in Write Queue	Average number of queuing tasks in the write thread pools of nodes in a CSS cluster.	≥ 0
avg_thread_pool_search_queue	Avg. Tasks in Search Queue	Average number of queuing tasks in the search thread pools of nodes in a CSS cluster.	≥ 0
avg_thread_pool_force_merge_queue	Avg. Tasks in ForceMerge Queue	Average number of queuing tasks in the force merge thread pools of nodes in the CSS cluster.	≥ 0
avg_thread_pool_search_threads	Avg. Size of Search Thread Pool	Average size of the search thread pool of a node in a CSS cluster.	≥ 0
avg_thread_pool_write_threads	Avg. Size of Write Thread Pool	Average size of the write thread pool of a node in a CSS cluster.	≥ 0
avg_thread_pool_force_merge_threads	Avg. Size of ForceMerge Thread Pool	Average size of the force merge thread pool of a node in a CSS cluster.	≥ 0
avg_thread_pool_write_rejected	Avg. Rejected Tasks in Write Queue	Average number of rejected tasks in the write thread pool of a node in a CSS cluster.	≥ 0
min_free_fs_size	Min. Available Storage Space	Minimum available storage space of nodes in a CSS cluster. Unit: byte	≥ 0 bytes
avg_jvm_old_gc_count	Avg. GCs of Old-Generation JVM	Average number of old-generation garbage collections of nodes in a CSS cluster.	≥ 0

Metric ID	Metric	Description	Value Range
avg_jvm_old_gc_time	Avg. GC Duration of Old-Generation JVM	Average old-generation garbage collection duration of nodes in a CSS cluster. Unit: ms	≥ 0 ms
avg_jvm_young_gc_count	Avg. GCs of Young-Generation JVM	Average number of young-generation garbage collections of nodes in a CSS cluster.	≥ 0
avg_jvm_young_gc_time	Avg. GC Duration of Young-Generation JVM	Average young-generation garbage collection duration of nodes in a CSS cluster. Unit: ms	≥ 0 ms
avg_max_file_descriptors	Avg. Maximum Allowed File Descriptors	Average value of the maximum number of allowed file descriptors on each node in a CSS cluster.	≥ 0
avg_mem_free_in_bytes	Avg. Available Memory	Average unused memory capacity of nodes in a CSS cluster. Unit: byte	≥ 0 bytes
avg_mem_free_percent	Avg. Available Memory Percentage	Average percentage of unused memory of nodes in a CSS cluster. Unit: %	0-100%
avg_mem_used_in_bytes	Avg. Used Memory	Average used memory of nodes in a CSS cluster. Unit: byte	≥ 0 bytes
avg_mem_used_percent	Avg. Used Memory Percentage	Average percentage of used memory of nodes in a CSS cluster. Unit: %	0-100%

Metric ID	Metric	Description	Value Range
max_mem_free_in_bytes	Max. Available Memory	Maximum unused memory of nodes in a CSS cluster. Unit: byte	≥ 0 bytes
max_mem_free_percent	Max. Available Memory Percentage	Maximum percentage of unused memory of nodes in a CSS cluster. Unit: %	0-100%
max_mem_used_in_bytes	Max. Used Memory	Maximum used memory of nodes in a CSS cluster. Unit: byte	≥ 0 bytes
max_mem_used_percent	Max. Used Memory Percentage	Maximum percentage of used memory of nodes in a CSS cluster. Unit: %	0-100%
sum_jvm_old_gc_count	Total GCs of Old-Generation JVM	Number of old-generation garbage collections of nodes in a CSS cluster.	≥ 0
sum_jvm_old_gc_time	Total GC Duration of Old-Generation JVM	Total old-generation garbage collection duration of nodes in the CSS cluster. Unit: ms	≥ 0ms
sum_jvm_young_gc_count	Total GCs of Young-Generation JVM	Number of young-generation garbage collections of nodes in a CSS cluster.	≥ 0
sum_jvm_young_gc_time	Total GC Duration of Young-Generation JVM	Total young-generation garbage collection duration of nodes in the CSS cluster. Unit: ms	≥ 0 ms
sum_current_opened_http_count	Currently Open HTTP Connections	Number of HTTP connections that are open on nodes in a CSS cluster.	≥ 0

Metric ID	Metric	Description	Value Range
sum_total_opene d_http_count	Historical Open HTTP Connections	Number of HTTP connections that were open on nodes in a CSS cluster.	≥ 0
IndexingLatency	Average Index Latency	Average time required for a shard to complete an index operation. Unit: ms	≥ 0 ms
IndexingRate	Average Index Rate	Average number of index operations per second in a cluster. Unit: s	≥ 0s
SearchLatency	Average Search Latency	Average time required for a segment to complete the search operation. Unit: ms	≥ 0 bytes
SearchRate	Average QPS	Average queries per second (QPS) in a cluster. Unit: s	≥ 0/s

Dimension

Table 12-2 Dimension description

Key	Value
cluster_id	CSS cluster

12.2 Configuring Cluster Monitoring

You can use Cloud Eye to monitor the created clusters. After configuring the cluster monitoring, you can log in to the Cloud Eye management console to view cluster metrics.

The procedure for configuring cluster monitoring:

1. **Creating Alarm Rules:** Customize alarm rules for the monitoring metrics. Once a metric exceeds the threshold, the system will notify you by sending emails or HTTP/HTTPS requests.

2. **Configuring Monitoring Metrics:** Configure monitoring metrics for a cluster or a node in the cluster.
3. **Viewing Monitoring Metrics:** View the statistics of the monitoring metrics in specific periods.

Prerequisites

- The cluster is in the **Available** or **Processing** status.
- The cluster has been running properly for more than 10 minutes.

Recommended Monitoring Metrics

- Cluster CPU and JVM usage. You are advised to configure the following monitoring metrics: average JVM heap usage, maximum JVM heap usage, average CPU usage, and maximum CPU usage.
- Cluster write and query latency and throughput. You are advised to configure the following monitoring metrics: average index latency, average index rate, average search latency, and average QPS.
- Cluster write and query queue and rejected tasks. You are advised to configure the following monitoring metrics: tasks in write queue, tasks in search queue, rejected tasks in write queue, and rejected tasks in search queue.

Creating Alarm Rules

1. Log in to the Cloud Eye console.
2. In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
3. In the **Resource Type** column, select **Cloud Search Service** as criteria to search for alarm rules that meet the requirements.

If no alarm rules are available, create one by referring to *Creating an Alarm Rule and Notification*. For details about how to set **Resource Type** and **Dimension**, see [Table 12-3](#).

Table 12-3 Alarm rule configuration parameter

Parameter	Description	Remark
Resource Type	Type of the resource that the alarm rule is created for	Select Cloud Search Service .
Dimension	Metric dimension of the selected resource type	<p>CSS supports two dimensions. Select a dimension as required.</p> <ul style="list-style-type: none"> • CSS Clusters: Alarm rules are specified by cluster. • CSS Clusters - CSS Instances: Alarm rules are specified by node in a cluster.

Configuring Monitoring Metrics

1. Create a monitoring panel. For details, see *Creating a Monitoring Panel*. If an available monitoring panel has been created, skip this step.
2. Add CSS monitoring graphs. For details, see *Adding a Graph*.

For details about how to set **Resource Type** and **Dimension**, see [Table 12-4](#).

Table 12-4 Graph configuration parameter

Parameter	Description	Remark
Resource Type	Type of the resource to be monitored	Select Cloud Search Service .
Dimension	Metric dimension	<p>CSS supports two dimensions. Select a dimension as required.</p> <ul style="list-style-type: none"> ● CSS Clusters: Monitoring is executed by cluster. ● CSS Clusters - CSS Instances: Monitoring is executed by node in a cluster.

Viewing Monitoring Metrics

1. Log in to the CSS management console.
2. Choose **Clusters**. Locate the target cluster and choose **More > View Metric** in the **Operation** column.
3. Select a time range.
4. View the monitoring metrics.

13 Auditing

13.1 Key Operations Recorded by CTS

With CTS, you can record operations associated with CSS for later query, audit, and backtrack operations.

Prerequisites

CTS has been enabled.

Key Operations Recorded by CTS



Table 13-1 Key operations recorded by CTS

Operation	Resource Type	Event Name
Creating a cluster	cluster	createCluster
Deleting a cluster	cluster	deleteCluster
Expanding the cluster capacity	cluster	roleExtendCluster
Restarting a cluster	cluster	rebootCluster
Performing basic configurations for a cluster snapshot	cluster	updateSnapshotPolicy
Setting the automatic snapshot creation policy	cluster	updateAutoSnapshotPolicy
Manually creating a snapshot	snapshot	createSnapshot
Restoring a snapshot	snapshot	restoreSnapshot
Deleting a snapshot	snapshot	deleteSnapshot

13.2 Viewing Audit Logs

After you enable CTS, it starts recording operations related to CSS. The CTS management console stores the last seven days of operation records. This section describes how to query the last seven days of operation records on the CTS management console.

Procedure

1. Log in to the CTS management console.
2. Click  in the upper left corner and select a region.
3. In the navigation pane on the left, click **Trace List**.
4. You can use filters to query traces. The following four filter criteria are available:
 - **Trace Source, Resource Type, and Search By**
Select a filter criterion from the drop-down list.
When you select **Trace name** for **Search By**, select a specific trace name.
When you select **Resource ID** for **Search By**, enter a specific resource ID.
When you select **Resource name** for **Search By**, select or enter a specific resource name.
 - **Operator**: Select a specific operator (at user level rather than tenant level).
 - **Trace Status**: Available options include **All trace statuses, normal, warning, and incident**. You can only select one of them.
 - **Time Range**: You can query traces generated during any time range of the last seven days.
5. Click  on the left of a trace to expand its details.
6. Click **View Trace** in the **Operation** column. In the displayed **View Trace** dialog box, the trace structure details are displayed.
For details about the key fields in the CTS trace structure, see the *Cloud Trace Service User Guide*.

A Change History

Released On	Description
2024-09-06	Added: Upgrading Versions
2023-06-14	Added: Connecting a Cluster to a Dedicated Load Balancer
2023-03-28	Added: VPC Endpoint Service
2023-03-15	<ul style="list-style-type: none"> • Added: <ul style="list-style-type: none"> - Viewing Cluster Information - Replacing a Specified Node - Adding Master/Client Nodes - Changing the Security Mode - Public Network Access - Kibana Public Access - Enhanced Features • Updated the procedure description in: <ul style="list-style-type: none"> - Creating an Elasticsearch Cluster in Security Mode - Creating an Elasticsearch Cluster in Non-Security Mode - Scaling Out a Cluster - Changing Specifications - Removing Specified Nodes - Index Backup and Restoration - Binding an Enterprise Project - Deleting a Cluster
2022-08-02	This is the first official release.