Cloud Search Service

FAQs

Issue	01
Date	2024-07-01

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

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

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

Trademarks and Permissions

HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

1 General Consulting

1.1 How Does CSS Ensure Data and Service Security?

CSS uses network isolation, in addition to various host and data security measures.

- Network isolation

  The entire network is divided into two planes: service plane and management plane. The two planes are deployed and isolated physically to ensure the security of the service and management networks.

  – Service plane: This is the network plane of the cluster. It provides service channels for users and delivers data definitions, indexing, and search capabilities.

  – Management plane: This is the management console, where you manage CSS.

- Host security

  CSS provides the following security measures:

  – The VPC security group ensures the security of the hosts in a VPC.

  – Network access control lists (ACLs) allow you to control what data can enter or exit your network.

  – The internal security infrastructure (including the network firewall, intrusion detection system, and protection system) monitors all network traffic that enters or exits the VPC through an IPsec VPN.

- Data security

  CSS uses multiple replicas, cross-AZ deployment of clusters, and third-party (OBS) backup of index data to ensure the security of user data.

1.2 What Storage Options Does CSS Provide?

CSS uses EVS and local disks to store your indices. During cluster creation, you can specify the EVS disk type and specifications (the EVS disk size).

- Supported EVS disk types include common I/O, high I/O, and ultra-high I/O.

- The EVS disk size varies depending on the node specifications selected when you create a cluster.

1.3 What Is the Maximum Storage Capacity of CSS?

You can configure up to 200 nodes for a cluster (each node corresponds to an ECS). The maximum storage capacity of an ECS is the total capacity of EVS disks attached to the ECS. You can calculate the total storage capacity of CSS based on the sizes of EVS disks attached to different ECSs. The EVS disk size is determined by the node specifications selected when you create the cluster.

1.4 How Can I Manage CSS?

You can use any of the following three methods to manage CSS or to use search engine APIs. You can initiate requests based on constructed request messages.

- curl

  curl is a command-line tool used to transfer data to or from a given URL. It serves as an HTTP client that can send HTTP requests to the HTTP server and receive response messages. You can also use curl to debug APIs. For more information about curl, visit **https://curl.haxx.se/**.

- Encoding

  You can call APIs through code to assemble, send, and process request messages.

- REST client

  Both Mozilla Firefox and Google Chrome provide a graphical browser plugin, the REST client, which you can use to send and process requests.

  – For Mozilla Firefox, see **Firefox REST Client**.

  – For Google Chrome, see **Postman**.

1.5 What Can the Disk Space of a CSS Cluster Be Used For?

You can store the following logs and files:

- Log files: Elasticsearch logs
- Data files: Elasticsearch index files
- Other files: cluster configuration files
- OS: 5% storage space reserved for the OS by default

1.6 What Data Compression Algorithms Does CSS Use?

CSS supports two data compression algorithms: LZ4 (by default) and best_compression.

- **LZ4 algorithm**

LZ4 is the default compression algorithm for Elasticsearch. This algorithm can compress and decompress data quickly, but its compression ratio is low.

LZ4 scans data with a 4-byte window, which slides 1 byte forward at a time. Duplicate data is compressed. This algorithm applies to scenarios where a large amount of data to be read while a small amount of data to be written.

- **best_compression algorithm**

  This algorithm can be used when a large amount of data is written and the index storage cost is high, such as logs and time sequence analysis. This algorithm can greatly reduce the index storage cost.

Run the following command to switch the default compression algorithm (LZ4) to best_compression:

```
PUT index-1
{
    "settings": {
        "index": {
            "codec": "best_compression"
        }
    }
}
```

The LZ4 algorithm can quickly compress and decompress data while the best_compression algorithm has a higher compression and decompression ratio.

2 Billing

2.1 How Do I Unsubscribe from a CSS Cluster?

Unsubscribing from a Yearly/Monthly Cluster

1. Log in to the CSS management console.
2. On the **Clusters** page, locate the cluster you want to unsubscribe from.
3. Choose **More** > **Unsubscribe/Release** in the **Operation** column.
4. In the **Unsubscribe from Cluster** dialog box, enter the name of the cluster you want to unsubscribe from and click **OK**.

   On the displayed page, confirm the resource information and refund amount.
5. Select the unsubscription reason, select the acknowledgement check boxes, and click **Unsubscribe**.

   In the displayed confirmation dialog box, click **Unsubscribe**.

Unsubscribing from a Pay-per-Use Cluster

1. Log in to the CSS management console.
2. On the **Clusters** page, locate the cluster you want to unsubscribe from.
3. In the **Operation** column, choose **More** > **Delete**.
4. In the **Delete Cluster** dialog box, enter the name of the cluster you want to delete and click **OK**.

2.2 How Do I Renew the Yearly/Monthly Resources of CSS?

CSS resources can be renewed yearly or monthly. The renewal operations are as follows:

Renewing an existing cluster

Perform the following steps:

1. On the CSS console, choose **Clusters**.
2. In the row of a yearly/monthly cluster, choose **More** > **Renew**.
3. Select the required duration and pay for the order.

Enabling auto-renew during cluster creation

When creating a cluster, perform the following steps:

On the cluster creation page, select a required duration and select **Auto-renew**. The cluster will be automatically renewed when its subscription expires.

Figure 2-1 Enabling auto-renew

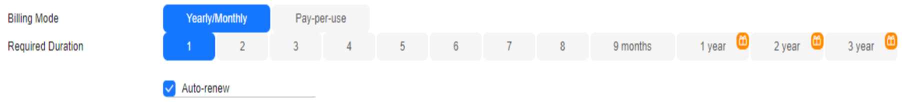

Billing Mode: Yearly/Monthly | Pay-per-use

Required Duration: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 months | 1 year | 2 year | 3 year

☑ Auto-renew

For more information about yearly/monthly renewals, see **Renewal Management**.

3 Cluster Management

3.1 Regions and AZs

3.1.1 What Are Regions and AZs?

Regions and AZs

A region or an availability zone (AZ) identifies the location of a data center. You can create resources in a specific region or an AZ.

- Regions are determined based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP, and Image Management Service (IMS), are shared within the same cloud region. Regions are classified as universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides only services of the same type or provides services only for specific tenants.
- An AZ contains one or multiple physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electrical facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers so you can build cross-AZ high-availability systems.

Figure 3-1 illustrates the relationship between regions and AZs.

Figure 3-1 Regions and AZs

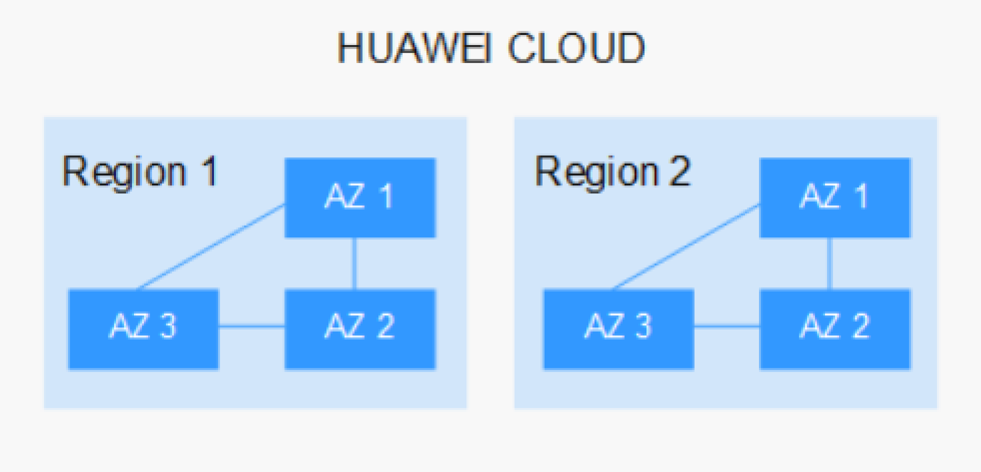

Huawei Cloud provides services in many regions worldwide. You can select a region and AZ as required. For more information, see **Huawei Cloud Global Regions**.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for low network latency and quick access.

  - If you or your target users are in the Asia Pacific area (excluding the Chinese mainland), select the **CN-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
  - If your target users are in Africa, select the **AF-Johannesburg** region.
  - If your target users are in Europe, select the **EU-Paris** region.
  - If your target users are in Latin America, select the **LA-Santiago** region.

  NOTE

  The **LA-Santiago** region is located in Chile.

- Resource price

  Resource prices may vary in different regions. For details, see **Product Pricing Details**.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- If your applications require low latency between instances, you are advised to deploy resources in the same AZ.

Regions and Endpoints

Before using an API to call resources, you will need to specify the resource region and endpoint. For details, see **Regions and Endpoints**.

3.1.2 How Do I View the AZ Where a Cluster Is Located?

You view the AZ where a cluster is located on the **Cluster Information** page.

1. Log in to the CSS management console.
2. Choose **Clusters** > **Elasticsearch**. The cluster list is displayed.
3. Click the cluster name to go to the **Cluster Information** page. In the **Configuration** area, view the **AZ** where the cluster is located.

Figure 3-2 Cluster configuration

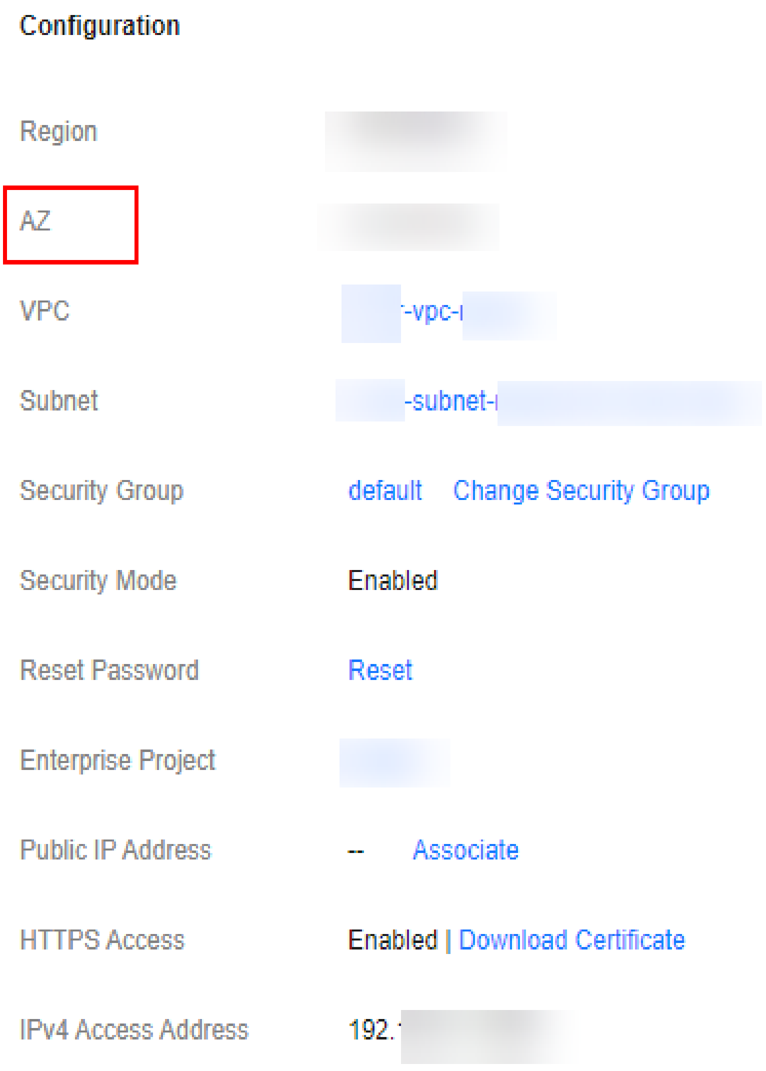

Configuration

Region	[redacted]
AZ	[redacted]
VPC	[redacted]-vpc-[redacted]
Subnet	[redacted]-subnet-[redacted]
Security Group	default Change Security Group
Security Mode	Enabled
Reset Password	Reset
Enterprise Project	[redacted]
Public IP Address	-- Associate
HTTPS Access	Enabled | Download Certificate
IPv4 Access Address	192.[redacted]

3.2 Cluster Version

3.2.1 What Is the Relationship Between the Filebeat Version and Cluster Version?

- Non-security mode: no restrictions.
- Cluster in security mode: The Filebeat OSS version must match the cluster version. For details on how to download the Filebeat OSS version, see **Past Releases of Elastic Stack Software**.

3.3 Clusters in Security Mode

3.3.1 How Do I Obtain the Security Certificate of CSS?

The security certificate (**CloudSearchService.cer**) can be downloaded only for security clusters that have enabled HTTPS access. The security certificate cannot be used in the public network environment.

1. Log in to the CSS management console.
2. In the navigation pane, choose **Clusters**. The cluster list is displayed.
3. Click the name of a cluster to go to the cluster details page.
4. On the **Configuration** page, click **Download Certificate** next to **HTTPS Access**.

Figure 3-3 Downloading a certificate

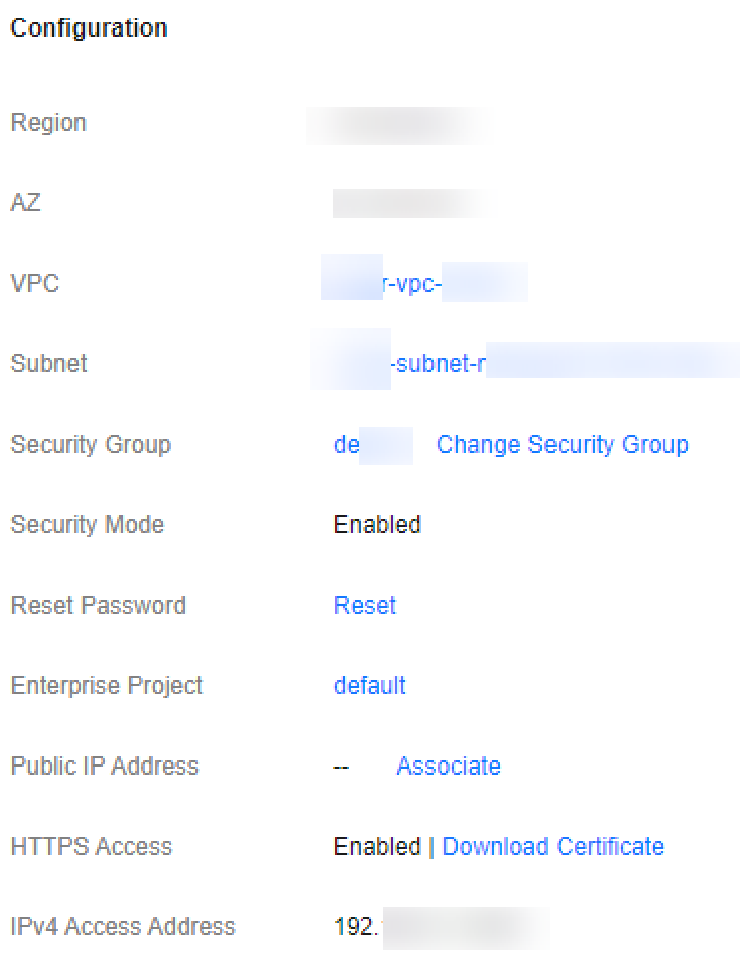

Configuration

Region	[redacted]
AZ	[redacted]
VPC	[redacted]r-vpc-[redacted]
Subnet	[redacted]-subnet-r[redacted]
Security Group	de[redacted] Change Security Group
Security Mode	Enabled
Reset Password	Reset
Enterprise Project	default
Public IP Address	-- Associate
HTTPS Access	Enabled | Download Certificate
IPv4 Access Address	192.[redacted]

3.3.2 How Do I Convert the Format of a CER Security Certificate?

The security certificate (**CloudSearchService.cer**) can be downloaded only for security clusters that have enabled HTTPS access. Most software supports certificates in the **.pem** or **.jks** format. You need to convert the format of the CSS security certificate.

- Run the following command to convert the security certificate from **.cer** to **.pem**:
  ```
  openssl x509 -inform pem -in CloudSearchService.cer –out newname.pem
  ```
- Run the following command to convert the security certificate from **.cer** to **.jks**:
  ```
  keytool -import -alias newname -keystore ./truststore.jks -file ./CloudSearchService.cer
  ```

In the preceding commands, *newname* indicates the user-defined certificate name.

After the command is executed, set the certificate password and confirm the password as prompted. Securely store the password. It will be used for accessing the cluster.

3.3.3 Can I Modify the Security Group of a CSS Cluster?

After a cluster is created, you can modify its security group.

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.
- You are advised to perform this operation during off-peak hours.
- The security group of a cluster created before February 2023 cannot be modified. You are advised to modify the security group of the cluster after **Migrating Data Through Backup and Restoration (from CSS Elasticsearch)** to a new cluster.

1. Log in to the CSS management console.
2. In the navigation pane, choose **Clusters**. The cluster list is displayed.
3. Click the name of a cluster to go to the cluster details page.
4. On the right of **Security Group**, click **Change Security Group**.

Figure 3-4 Changing a security group

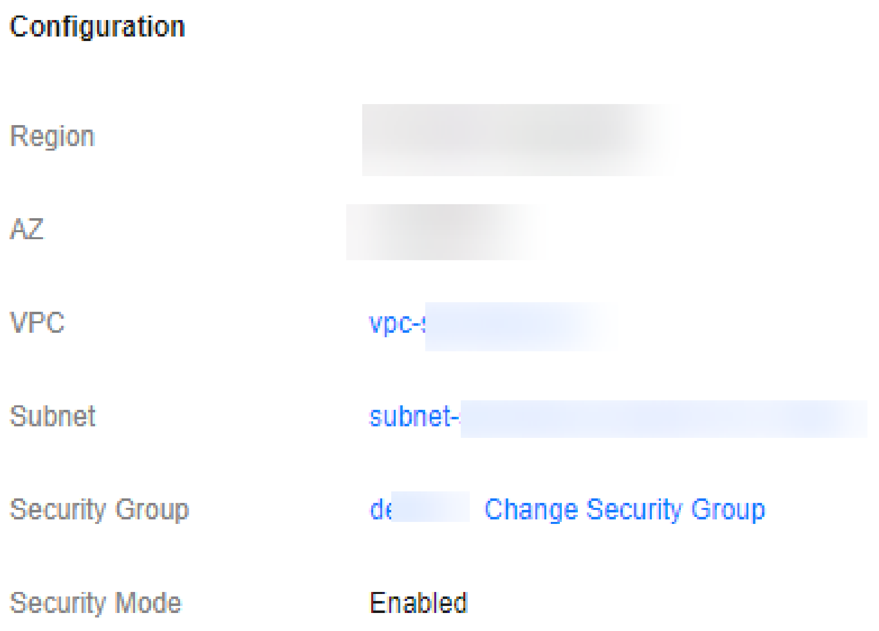

Configuration

Region

AZ

VPC vpc-

Subnet subnet-

Security Group de Change Security Group

Security Mode Enabled

5. In the **Change Security Group** dialog box, select a new security group and click **OK**.

3.4 Parameter Configuration

3.4.1 How Do I Set the search.max_buckets Parameter for an Elasticsearch Cluster?

Function

By default, CSS allows a maximum of 10,000 buckets to be returned during aggregation. If more than 10,000 buckets need to be returned, you can increase the value of **search.max_buckets**. Note that increasing the value of **search.max_buckets** also increases the cluster load and memory usage. Therefore, exercise caution when performing this operation.

Solution

Run the following command on the **Dev Tools** page of Kibana:

```
PUT _cluster/settings
{
    "persistent": {
        "search.max_buckets": 20000
    }
}
```

3.4.2 Can I Modify the TLS Algorithm of an Elasticsearch Cluster?

You can modify TLS algorithms in CSS 7.6.2 and later versions.

1. Log in to the CSS management console.
2. In the navigation pane, choose **Clusters**. The cluster list is displayed.
3. Click the name of the target cluster to go to the cluster details page.
4. Select **Parameter Configurations**, click **Edit**, expand the **Customize** parameter, and click **Add**.

   Add the **opendistro_security.ssl.http.enabled_ciphers** parameter and set it to **['TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256', 'TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384']**.

   NOTE

   If the parameter value contains multiple algorithm protocols, enclose the value with a pair of square brackets ([]). If the parameter value is a single algorithm protocol, enclose the value with a pair of single quotation marks(' ').

5. 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.

6. Return to the cluster list and choose **More** > **Restart** in the **Operation** column to restart the cluster and make the modification take effect.

3.4.3 How Do I Enable the Audit Log Function for an Elasticsearch Cluster?

Currently, CSS Elasticsearch clusters of the 7.6.2 and later versions support the audit log function. By default, this function is disabled.

NOTE

The cluster must be a security cluster.

1. Log in to the CSS management console.
2. In the navigation pane, choose **Clusters**. The cluster list is displayed.
3. Click the name of the target cluster to go to the cluster details page.
4. In the navigation pane on the left, choose **Parameter Configurations**. Click **Edit**, expand the **Customize** parameter, and click **Add**.

   Set **Key** to **opendistro_security.audit.type** and **Value** to **internal_elasticsearch**.

   Figure 3-5 Configuring a custom parameter

   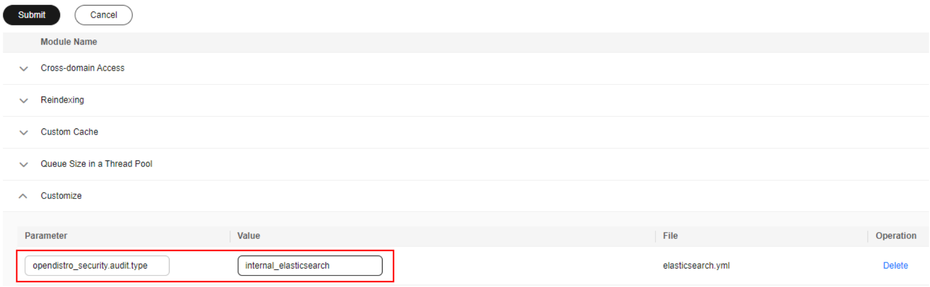

Parameter	Value	File	Operation
opendistro_security.audit.type	internal_elasticsearch	elasticsearch.yml	Delete

5. 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.
6. Return to the cluster list and choose **More** > **Restart** in the **Operation** column to restart the cluster and make the modification take effect.
7. After the cluster is restarted, click **Access Kibana** in the **Operation** column. On the displayed page, enter the username and password. The **Dev Tools** page is displayed.
8. In the **Console** page, run the **GET _cat/indices?v** command. If there are indexes related to **.*audit* index**, the audit log function is enabled.

3.4.4 How Do I Query the Index Size on OBS After the Freezing of Indexes for a CSS Cluster?

The size of indexes remains unchanged after freezing. By querying the size of frozen indexes in OBS, you obtain the size of all indexes stored on OBS.

Run the following command to obtain information about all indexes that are being frozen or have already been frozen:

```
GET _cat/freeze_indices?stage=$
```

The output is as follows (as an example only):

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

The last column of the returned result contains the index size information.

Related Questions

- **Billing for index storage on OBS**

  Fees may be incurred when you store indexes in OBS. For details, see the price of standard single-AZ storage in **OBS Price Calculator**.

- **Why can frozen indexes stored in OBS still be queried using commands?**

  Elasticsearch and OpenSearch clusters use local storage by default, and Lucene index files are stored on local disks. Lucene interacts with the underlying storage via the Directory API. Files can be read through the following API:

  ```
  public abstract IndexInput openInput(String name, IOContext context) throws IOException;
  ```

  The storage-compute decoupling feature enables interaction with OBS through the Directory API to read files stored in OBS. This is why information about frozen indexes stored in OBS can be queried using commands.

4 Open Source Search Engine Consulting

4.1 How Do I Set the Numbers of Index Copies to 0 in Batches?

1. Log in to the Kibana page of the cluster. In the navigation pane, choose **Dev Tools**.
2. Modify and run the **PUT /*/_settings{"number_of_replicas":0}** command.

   NOTE

   Do not directly run the preceding command, because the asterisk (*) may match security indexes. You are advised to specify the index required for the batch operation. Example: **PUT /test*/_settings{"number_of_replicas":0}**

4.2 Why All New Index Shards Are Allocated to the Same Node?

Possible Cause

The possible causes are as follows:

- Shards were unevenly distributed in previous index allocations, and the predominate parameter in the latest indexed shard allocation was **balance.shard**. To balance the shard distribution across nodes, the new shards were allocated to the node with only a small number of shards.
- After a new node was added to a cluster and before the automatic cluster rebalancing completes, the predominate parameter was **balance.shard**. The shards of a new index are allocated to the new node, where there are no shards yet.

The following two parameters are used to balance the shard allocation in a cluster:

cluster.routing.allocation.balance.index (default value: **0.45f**)

cluster.routing.allocation.balance.shard (default value: **0.55f**)

NOTE

- **balance.index**: A larger value indicates that all the shards of an index are more evenly distributed across nodes. For example, if an index has six shards and there are three data nodes, two shards will be distributed on each node.
- **balance.shard**: A larger value indicates that all the shards of all the indexes are more evenly distributed across nodes. For example, if index **a** has two shards, index **b** has four, and there are three data nodes, two shards will be distributed on each node.
- You can specify both **balance.index** and **balance.shard** to balance the shard allocation.

Solution

To prevent the all the shards of an index from being allocated to a single node, use either of the following methods:

1. To create an index during cluster scale-out, configure the following parameter:
   ```
   "index.routing.allocation.total_shards_per_node": 2
   ```
   That is, allow no more than two shards of an index to be allocated on each node. Determine the maximum number of shards allocated to each node based on the number of data nodes in your cluster and the number of index shards (both primary and secondary).
2. If too many shards are distributed on only a few nodes, you can move some of the shards to other nodes to balance the distribution. Run the **move** command of **POST _cluster/reroute**. The rebalance module will automatically exchange the shard with a shard on the destination node. Determine the values of **balance.index** and **balance.shard** as needed.

4.3 How Do I Create a Type Under an Index in an Elasticsearch 7.*x* Cluster?

In Elasticsearch 7.*x* and later versions, types cannot be created for indexes.

If you need to use types, add **include_type_name=true** to the command. For example:

```
PUT _template/urldialinfo_template?include_type_name=true
```

After the command is executed, the following information is displayed:

```
"#! Deprecation: [types removal] Specifying include_type_name in put index template requests is deprecated. The parameter will be removed in the next major version. "
```

4.4 How Do I Configure a Two-Replica CSS Cluster?

1. Run **GET _cat/indices?v** in Kibana to check the number of cluster replicas. If the value of **rep** is **1**, the cluster has two replicas.

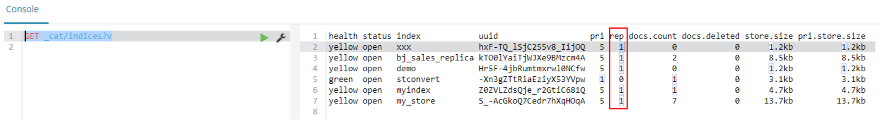

```
GET _cat/indices?v
```

```
health status index            uuid                   pri rep docs.count docs.deleted store.size pri.store.size
yellow open   xxx              hxF-TQ_lSjC255v8_IijOQ   5   1          0            0      1.2kb          1.2kb
yellow open   bj_sales_replica kTO0lYaiTjWJXe9BMzcm4A   5   1          2            0      8.5kb          8.5kb
yellow open   demo             Hr5F-4jbRumtmxrwl0NCfw   5   1          0            0      1.2kb          1.2kb
green  open   stconvert        -Xn3gZTtRiaEziyX53YVpw   1   0          1            0      3.1kb          3.1kb
yellow open   myindex          Z0ZVLZdsQje_r2GtiC681Q   5   1          1            0      4.7kb          4.7kb
yellow open   my_store         S_-AcGkoQ7Cedr7hXqHOqA   5   1          7            0     13.7kb         13.7kb
```

2. If the value of **rep** is not **1**, run the following command to set the number of replicas:

```
PUT /index/_settings
{
"number_of_replicas" : 1 //Number of replicas
}
```

NOTE

index specifies the index name. Set this parameter based on site requirements.

4.5 Can I Change the Number of Shards to Four with Two Replicas When There Is One Shard Set in the JSON File?

Once an index is created, the number of primary shards cannot be changed.

You can run the following command in Kibana to change the number of replicas:

```
PUT /indexname/_settings
{
"number_of_replicas" :1       //Number of replicas
}
```

NOTE

index specifies the index name. Set this parameter based on site requirements.

4.6 What Are the Impacts If an Elasticsearch Cluster Has Too Many Shards?

1. A large number of shards in a cluster slows down shard creation.
2. If automatic index creation is enabled, slow index creation may cause a large number of write requests to be stacked in the memory or result in a cluster breakdown.
3. If there are too many shards and you cannot properly monitor workloads, the number of records in a single shard may exceed the threshold, and write requests may be denied.

4.7 How Do I Check the Numbers of Shards and Replicas in a Cluster on the CSS Console?

1. Log in to the CSS management console.
2. On the **Clusters** page, click **Access Kibana** in the **Operation** column of a cluster.
3. Log in to Kibana and choose **Dev Tools**.

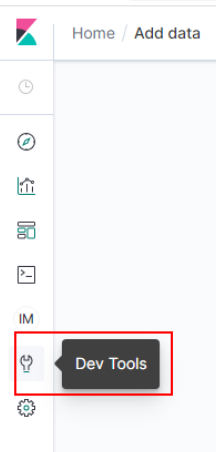

4. On the **Console** page, run the **GET _cat/indices?v** command query the number of shards and replicas in a cluster. In the following figure, the **pri** column indicates the number of index shards, and the **rep** column indicates the number of replicas. After an index is created, its **pri** value cannot be modified. Its **rep** value can be modified.

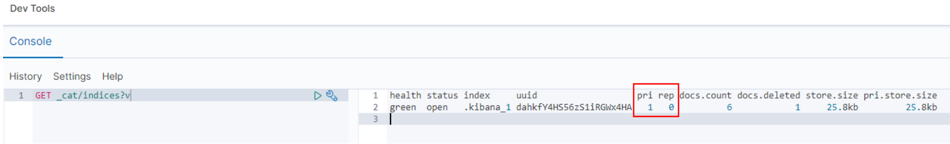

```
GET _cat/indices?v
```

```
health status index     uuid                   pri rep docs.count docs.deleted store.size pri.store.size
green  open   .kibana_1 dahkfY4HS56zS1iRGWx4HA   1   0          6            1     25.8kb         25.8kb
```

4.8 How Do I Query Index Data on Kibana in an Elasticsearch Cluster?

Run the following command to query index data through an API on Kibana:

```
GET indexname/_search
```

The returned data is shown in the following figure.

Figure 4-1 Returned data

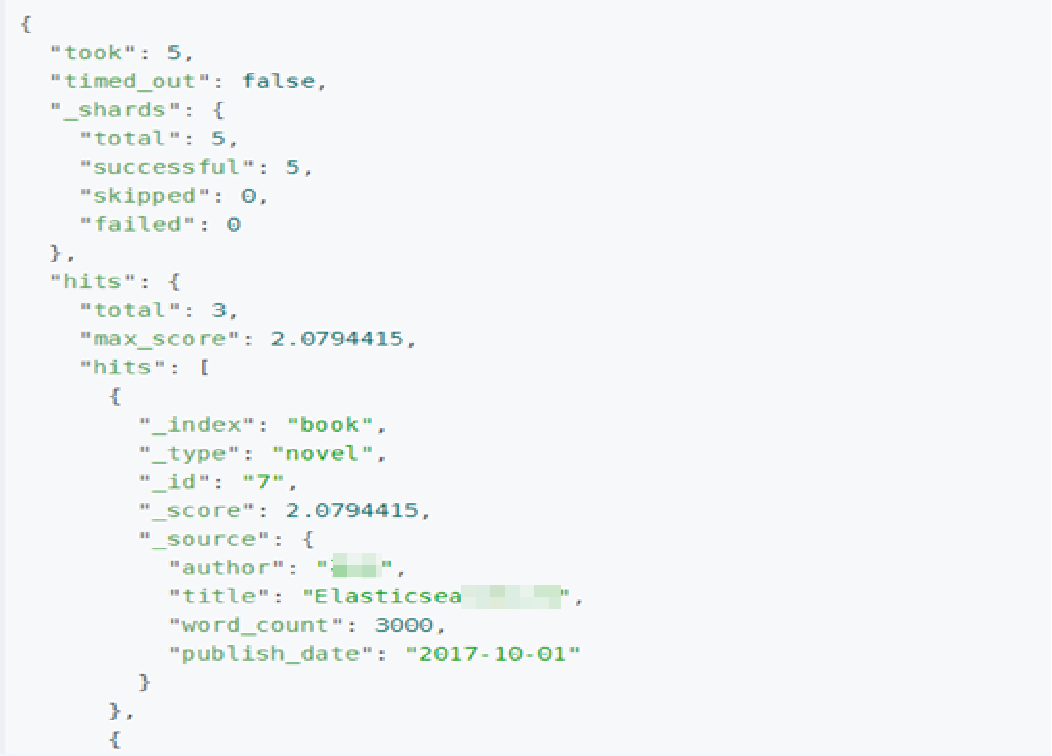

```
{
  "took": 5,
  "timed_out": false,
  "_shards": {
    "total": 5,
    "successful": 5,
    "skipped": 0,
    "failed": 0
  },
  "hits": {
    "total": 3,
    "max_score": 2.0794415,
    "hits": [
      {
        "_index": "book",
        "_type": "novel",
        "_id": "7",
        "_score": 2.0794415,
        "_source": {
          "author": "[redacted]",
          "title": "Elasticsea[redacted]",
          "word_count": 3000,
          "publish_date": "2017-10-01"
        }
      },
      {
```

Table 4-1 Parameters

Parameter	Description
took	Indicates how many milliseconds the query cost.
time_out	Indicates whether the query times out.
_shard	Data is split into five shards. All of the five shards have been searched and data is returned successfully. No query result fails to be returned. No data is skipped.
hits.total	Number of query results. Three documents are returned in this example.
max_score	Score of the returned documents. The document that is more relevant to your search criteria would have a higher score.
hits.hits	Detailed information of the returned documents

4.9 Can I Stop a CSS Cluster?

No. If you need to migrate a cluster, you can suspend the services of the old cluster and delete it after the migration is complete. You can perform the following operations:

- If the cluster version supports **flow control**, you can use **one-click traffic blocking** to block all the traffic, except the traffic of the O&M interface on the node.
- If your cluster version in use does not traffic control, you can disable the read and write of all service indexes instead. For example, if all service indexes start with **log**, run the following command on the **Dev Tools** page of Kibana:

```
PUT log*/_settings
{
  "index.blocks.read": true,
  "index.blocks.write": true,
  "index.blocks.metadata": true
}
```

4.10 Does the Value i of node.roles Indicate an Ingest Node?

Function

If the value of **node.roles** of a client node is **i**, then is this client node an ingest node?

- Are there coordinating only nodes in clusters? Are the client requests distributed to coordinating nodes?
- Are ingest nodes in idle state when there are no ingest requests?

Solution

If the value of **node.roles** of a client node is **i**, the ingest node mode is enabled.

- The coordinating only nodes of Elasticsearch are called client nodes in CSS. If a cluster has no client nodes, client requests will be distributed to all nodes.
- An ingest node functions as a set of ELK for data conversion. If there is no ingest requests, ingest nodes are not in the idle state.

4.11 How Do I Set the Default Maximum Number of Records Displayed on a Page for an Elasticsearch Cluster

Solution

- Method 1

  Open Kibana and run the following commands on the **Dev Tools** page:

  ```
  PUT _all/_settings?preserve_existing=true
  {
  "index.max_result_window" : "10000000"
  }
  ```

- Method 2

  Run the following commands in the background:

  ```
  curl –XPUT 'http://localhost:9200/_all/_setting?preserve_existing=true'-d
  {
  ```

```
"index.max_result_window":"1000000"
}
```

⚠ CAUTION

This configuration consumes memory and CPU resources. Exercise caution when setting this parameter.

4.12 How Do I Update the Lifecycle Policy of an Elasticsearch Cluster?

The lifecycle of Elasticsearch clusters is implemented using the Index State Management (ISM) of Open Distro. For details about how to configure policies related to the ISM template, see the **Open Distro documentation**.

1. When a policy is created, the system writes a record to the **.opendistro-ism-config** index. In the record, **_id** is the policy name, and the content is the policy definition.

   Figure 4-2 Writing a data record

   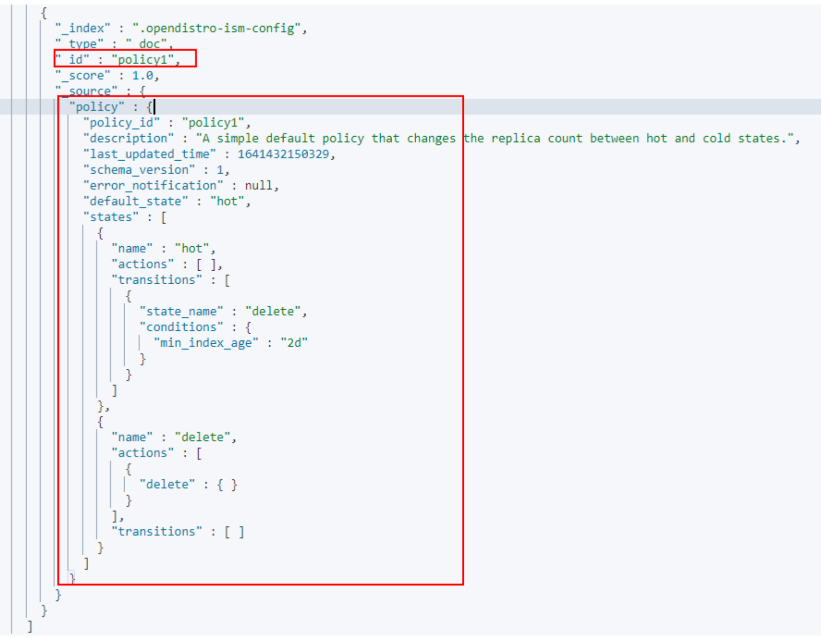

```
    {
      "_index" : ".opendistro-ism-config",
      "_type" : "_doc",
      "_id" : "policy1",
      "_score" : 1.0,
      "_source" : {
        "policy" : {
          "policy_id" : "policy1",
          "description" : "A simple default policy that changes the replica count between hot and cold states.",
          "last_updated_time" : 1641432150329,
          "schema_version" : 1,
          "error_notification" : null,
          "default_state" : "hot",
          "states" : [
            {
              "name" : "hot",
              "actions" : [ ],
              "transitions" : [
                {
                  "state_name" : "delete",
                  "conditions" : {
                    "min_index_age" : "2d"
                  }
                }
              ]
            },
            {
              "name" : "delete",
              "actions" : [
                {
                  "delete" : { }
                }
              ],
              "transitions" : [ ]
            }
          ]
        }
      }
    }
  ]
```

2. After a policy is bound to an index, the system writes another record to the **.opendistro-ism-config** index. The following figure shows the initial status of a record.

Figure 4-3 Initial data status

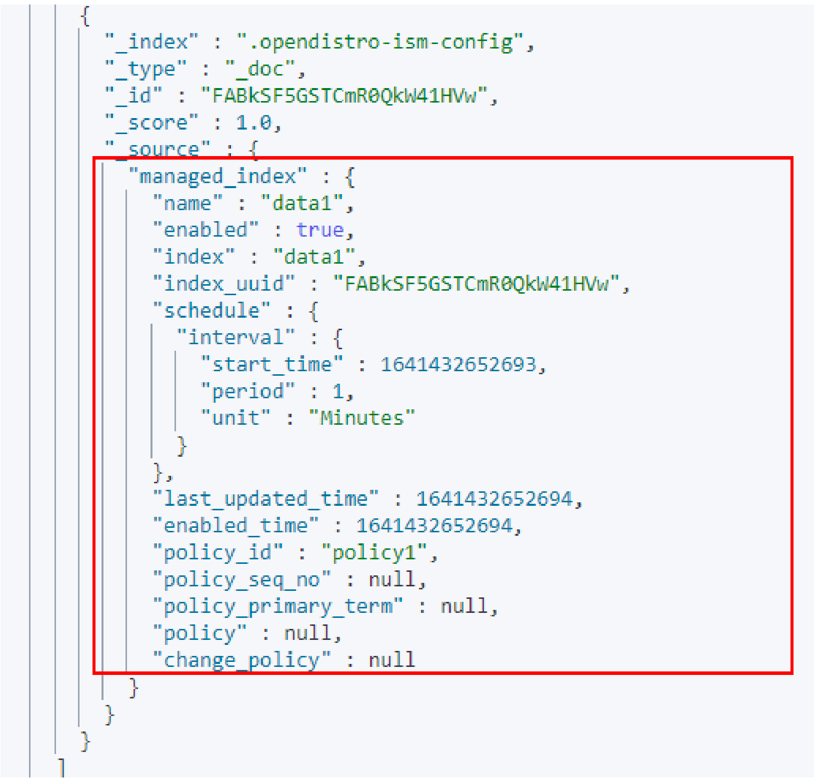

```
{
  "_index" : ".opendistro-ism-config",
  "_type" : "_doc",
  "_id" : "FABkSF5GSTCmR0QkW41HVw",
  "_score" : 1.0,
  "_source" : {
    "managed_index" : {
      "name" : "data1",
      "enabled" : true,
      "index" : "data1",
      "index_uuid" : "FABkSF5GSTCmR0QkW41HVw",
      "schedule" : {
        "interval" : {
          "start_time" : 1641432652693,
          "period" : 1,
          "unit" : "Minutes"
        }
      },
      "last_updated_time" : 1641432652694,
      "enabled_time" : 1641432652694,
      "policy_id" : "policy1",
      "policy_seq_no" : null,
      "policy_primary_term" : null,
      "policy" : null,
      "change_policy" : null
    }
  }
}
]
```

3. Run the **explain** command. Only a policy ID will be returned.

```
GET _opendistro/_ism/explain/data2
{
  "data2" : {
    "index.opendistro.index_state_management.policy_id" : "policy1"
  }
}
```

Open Distro will execute an initialization process to fill the policy content in the record. The following figure shows the initialized data.

Figure 4-4 Initialized data

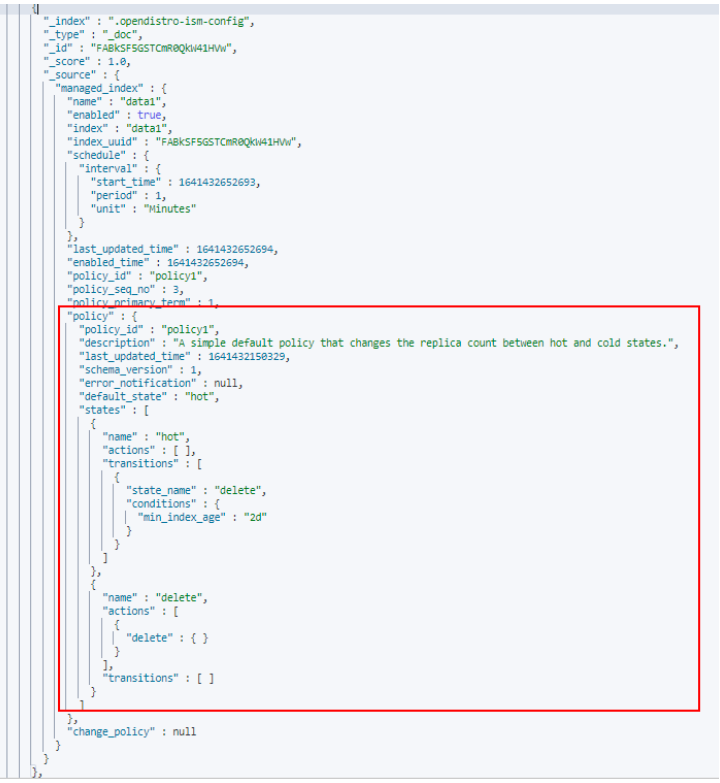

```
{
  "_index" : ".opendistro-ism-config",
  "_type" : "_doc",
  "_id" : "FABkSF5GSTCmR0QkW41HVw",
  "_score" : 1.0,
  "_source" : {
    "managed_index" : {
      "name" : "data1",
      "enabled" : true,
      "index" : "data1",
      "index_uuid" : "FABkSF5GSTCmR0QkW41HVw",
      "schedule" : {
        "interval" : {
          "start_time" : 1641432652693,
          "period" : 1,
          "unit" : "Minutes"
        }
      },
      "last_updated_time" : 1641432652694,
      "enabled_time" : 1641432652694,
      "policy_id" : "policy1",
      "policy_seq_no" : 3,
      "policy_primary_term" : 1,
      "policy" : {
        "policy_id" : "policy1",
        "description" : "A simple default policy that changes the replica count between hot and cold states.",
        "last_updated_time" : 1641432150329,
        "schema_version" : 1,
        "error_notification" : null,
        "default_state" : "hot",
        "states" : [
          {
            "name" : "hot",
            "actions" : [ ],
            "transitions" : [
              {
                "state_name" : "delete",
                "conditions" : {
                  "min_index_age" : "2d"
                }
              }
            ]
          },
          {
            "name" : "delete",
            "actions" : [
              {
                "delete" : { }
              }
            ],
            "transitions" : [ ]
          }
        ]
      },
      "change_policy" : null
    }
  }
},
```

After the initialization, **min_index_age** in the policy will be copied.

NOTE

The initialized index uses a copy of this policy. The policy update will not take effect on the index.

4. After the policy is modified, call the **change_policy** API to update the policy.

```
POST _opendistro/_ism/change_policy/data1
{
  "policy_id": "policy1"
}
```

Related Information

For details about how to create and use a lifecycle policy, see **Managing the Index Life Cycle**.

4.13 How Do I Configure the Threshold for CSS Slow Query Logs?

The slow query log settings of CSS are the same as those of Elasticsearch. You can configure slow query logs via the _settings API. For example, you can run the following command in Kibana to set the index level:

```
PUT /my_index/_settings
{
    "index.search.slowlog.threshold.query.warn": "10s",
    "index.search.slowlog.threshold.fetch.debug": "500ms",
  "index.indexing.slowlog.threshold.index.info": "5s"
}
```

- If a query takes longer than 10 seconds, a WARN log will be generated.
- If retrieval takes longer than 500 milliseconds, a DEBUG log will be generated.
- If an index takes longer than 5 seconds, an INFO log will be generated.

For details, visit the official website: https://www.elastic.co/guide/cn/elasticsearch/guide/current/logging.html

4.14 How Do I Delete Index Data?

- Automatically and Periodically Clearing Indexes

  You can create a scheduled task to call and execute the index deletion request periodically. CSS supports Open Distro Index State Management. For details, see: https://opendistro.github.io/for-elasticsearch-docs/docs/im/ism/
- Manually clear indexes.
  - Log in to Kibana and run the **DELETE / Index name** command in Dev Tools.
  - Log in to Cerebro, search for the target index name, click the index name, click **delete index**, and click **Confirm** in the displayed dialog box.

Figure 4-5 Deleting an index from Cerebro

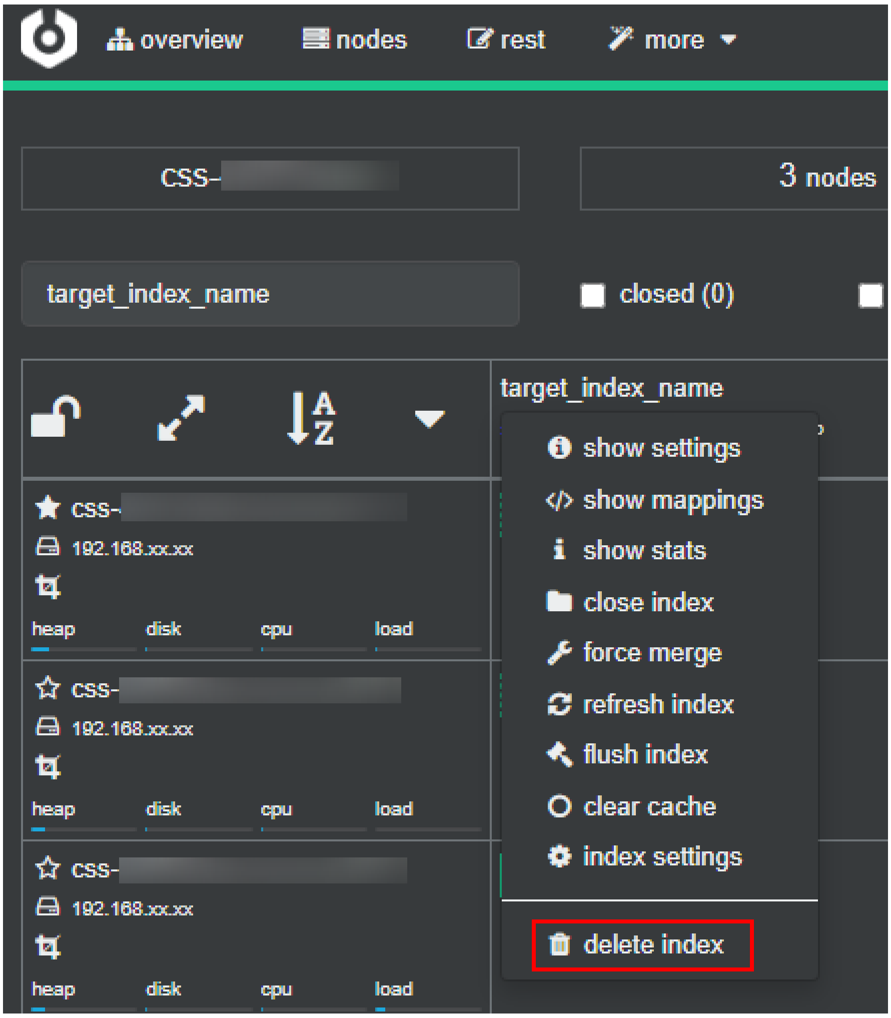

– Log in to the ECS and run the following command to delete a single index data record:

curl -XDELETE http://IP:9200/*Index_name*

Run the following command to delete all Logstash data of a day. For example, delete all data on June 19, 2017:

For a cluster in non-security mode: **curl -XDELETE 'http://IP:9200/logstash-2017.06.19*'**

For a cluster in security mode: **curl -XDELETE -u username:password 'https://IP:9200/logstash-2017.06.19' -k**

NOTE

- **username**: username of the administrator. The default value is **admin**.
- **password**: the password set during cluster creation
- **IP**: the IP address of any node in the cluster

4.15 How Do I Clear the Cache of a CSS Cluster?

- **Clear the fielddata**

During aggregation and sorting, data are converted to the fielddata structure, which occupies a large amount of memory.

a. Run the following commands on Kibana to check the memory occupied by index fielddata:

```
DELETE /_search/scroll
{
"scroll_id" :
"DXF1ZXJ5QW5kRmV0Y2gBAAAAAAAAAD4WYm9laVYtZndUQlNsdDcwakFMNjU1QQ=="
}
```

b. If the memory usage of fielddata is too high, you can run the following command to clear fielddata:

```
POST /test/_cache/clear?fielddata=true
```

In the preceding command, *test* indicates the name of the index whose fielddata occupies a large amount of memory.

- **Clear segments**

  The FST structure of each segment is loaded to the memory and will not be cleared. If the number of index segments is too large, the memory usage will be high. You are advised to periodically clear the segments.

  a. Run the following command on Kibana to check the number of segments and their memory usage on each node:

  ```
  GET /_cat/nodes?v&h=segments.count,segments.memory&s=segments.memory:desc
  ```

  b. If the memory usage of segments is too high, you can delete or disable unnecessary indexes, or periodically combine indexes that are not updated.

- **Clear the cache**

  Run the following command on Kibana to clear the cache:

  ```
  POST _cache/clear
  ```

4.16 Why Does the Disk Usage Increase After the delete_by_query Command Was Executed to Delete Data?

Running the **delete_by_query** command can only add a deletion mark to the target data instead of really deleting it. When you search for data, all data is searched and the data with the deletion mark is filtered out.

The space occupied by an index with the deletion mark will not be released immediately after you call the disk deletion API. The disk space is released only when the segment merge is performed next time.

Querying the data with deletion mark occupies disk space. In this case, the disk usage increases when you run the disk deletion commands.

5 Cluster Plugin Usage

5.1 Can I Install Search Guard on CSS?

CSS does not currently support installation of Search Guard.

CSS provides clusters in security mode, which have the same functions as Search Guard. For details about clusters in security mode, see **Clusters in Security Mode**.

5.2 Failed to Execute the Native script dotProduct of the Elasticsearch Cluster

Possible Cause

The native Elasticsearch vector function is provided by the x-pack plugin, but has not been integrated in CSS. The native **script dotProduct** cannot be executed in the Elasticsearch cluster.

Solution

You are advised to use the vector search function of CSS. Based on the vector search engine and the Elasticsearch plug-in mechanism, CSS efficiently integrates the vector search capability featuring high-performance, high-precision, low-cost, and multi-modal. For more information, see **Vector Retrieval**.

NOTE

The vector retrieval function is supported by clusters of versions 7.6.2 and 7.10.2.

6 Cluster Access/Cluster Connection

6.1 Can I Build a Kibana or Cerebro to Access CSS Clusters?

You can build a Kibana or Cerebro to access CSS clusters.

- For details about building a Kibana to access CSS clusters, see **How Do I Connect the User-Built Kibana to Elasticsearch on CSS?**.
- If you have built a Cerebro to access CSS clusters, just start your Cerebro and enter the internal IP address of the target cluster.
  - To access a cluster in security mode, enter: https://***Internal_IP_address***:9200.
  - To access a cluster in non-security mode, enter http://***Intranet_IP_address***:9200.

6.2 Do Ports 9200 and 9300 Both Open?

Yes. Port 9200 is used by external systems to access CSS clusters, and port 9300 is used for communication between nodes.

The methods for accessing port 9300 are as follows:

- If your client is in the same VPC and subnet with the CSS cluster, you can access it directly.
- If your client is in the same VPC with but different subnet from the CSS cluster, apply for a route separately.
- If your client is in the different VPCs and subnets from the CSS cluster, create a VPC peering connection to enable communication between the two VPCs, and then apply for routes to connect the two subnets.

6.3 How Do I Use a NAT Gateway to Access CSS from the Internet?

Perform the following operations:

1.**Obtaining CSS Information**

2.**Configuring a NAT Gateway**

3.**Modifying Security Group Rules**

4.**Accessing CSS from the Internet**

⚠ CAUTION

If your CSS clusters do not have the security mode enabled, do not access CSS through the NAT gateway. Otherwise, the cluster data will be exposed to the Internet.

Obtaining CSS Information

Step 1 Log in to the CSS management console.

Step 2 On the **Clusters** page, click the name of a cluster. The **Basic Information** page is displayed by default.

Step 3 In the **Configuration Information** area, view the **Private Network Address**, **VPC**, and **Subnet** information.

Figure 6-1 Required information

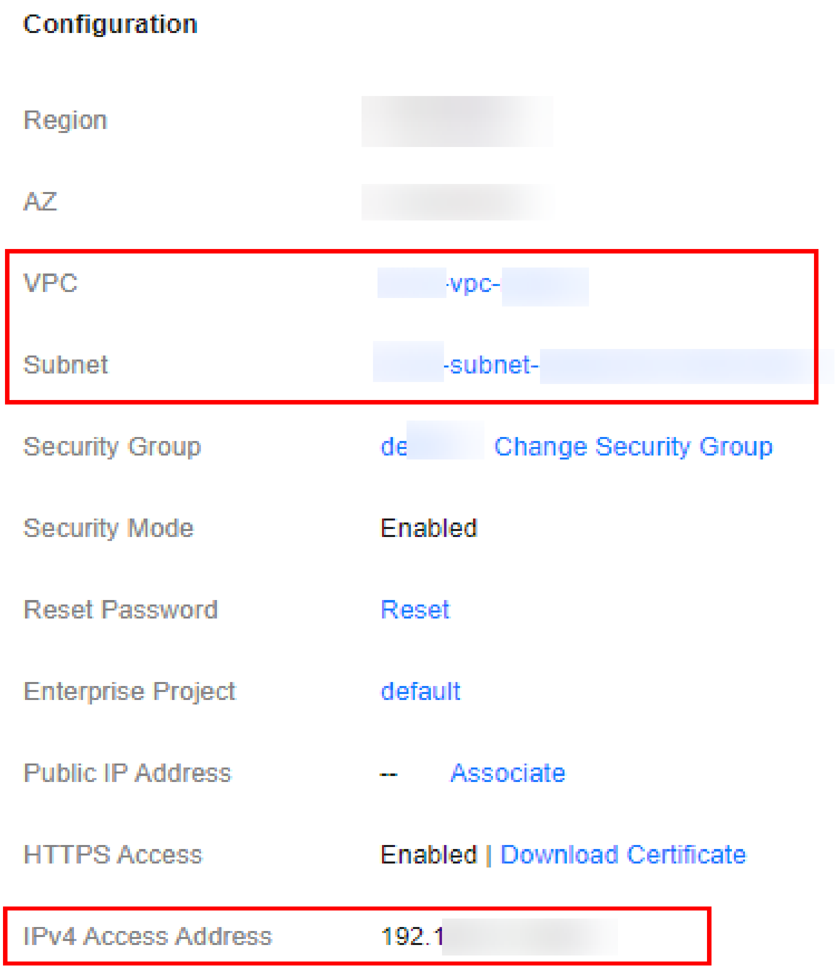

Configuration

Region	
AZ	
VPC	-vpc-
Subnet	-subnet-
Security Group	de Change Security Group
Security Mode	Enabled
Reset Password	Reset
Enterprise Project	default
Public IP Address	-- Associate
HTTPS Access	Enabled | Download Certificate
IPv4 Access Address	192.1

----End

Configuring a NAT Gateway

Step 1 Create a NAT gateway.

1. Log in to the console and choose **Service List** > **Networking** >**NAT Gateway**. The **Network Console** page is displayed.
2. Click **Buy Public NAT Gateway**. On the displayed page, configure related parameters. For details, see the section "Buying a NAT Gateway" in *NAT Gateway User Guide*.

   NOTE

   Set **VPC** and **Subnet** to the values you obtained in **Obtaining CSS Information**.
3. Click **Next**, confirm the configurations, and click **Pay Now**.

Step 2 Add DNAT rules.

1. On the **Public NAT Gateways** page, click the name of the NAT gateway you purchased. The details page is displayed.

2. Choose **DNAT Rules** > **Add DNAT Rule**. For details, see section "Adding a DNAT Rule" in the *NAT Gateway User Guide*. When configuring DNAT rules, use the following settings:

   NOTE

   - **EIP**: Create an EIP on the **EIPs** page based on your service requirements.
   - **Outside Port**: Custom.
   - **Private IP Address**: private network IP address of CSS, which is the **Private Network Address** you obtained in **Obtaining CSS Information**.
   - **Inside Port**: 9200.
   - If your cluster contains multiple private IP addresses, add one DNAT rule for each address.

3. Click **OK**.

----End

Modifying Security Group Rules

Step 1 Log in to the CSS management console. In the navigation pane, click **Clusters**. On the displayed **Clusters** page, click the name of the target cluster to go to the **Basic Information** page

Step 2 On the **Basic Information** page, click **Security Group**.

Step 3 On the **Basic Information** page of the security group, click the **Inbound Rules** tab.

Step 4 Click **Add Rule** to add an inbound rule for port 9200.

Step 5 Click **OK**.

----End

Accessing CSS from the Internet

Enter **https://*IP*:*port*** or **http://*IP*:*port*** in the address box of the browser.

- ***IP*** and ***port*** are an EIP and port you set when you added DNAT rules.
- If you have enabled **Security Mode** for the cluster, enter **https://*IP*:*port*** and then enter the username and password that you set for security mode on the displayed page.
- If you have not enabled **Security Mode** for the cluster, just enter **http://*IP*:*port***.

6.4 Can a New Cluster Use the IP Address of the Old Cluster?

No.

If the IP address of a cluster changes, the possible causes are as follows:

- The private IP address of the cluster changes.

  Check whether the cluster is scaled out or scaled in. The scale-out operation increases the number of private IP addresses of the cluster. The scale-in operation reduces the number of private IP addresses of the cluster. If there are services running on a node to be deleted, a fault occurs.

- The public IP address of the cluster changes.

  Check whether the cluster has enabled the security mode. When creating an Elasticsearch cluster earlier than version 6.5.4, you can enable the security mode and public network access. When creating an Elasticsearch cluster of version 6.5.4 or later, you can enable the security mode. After the security mode is enabled, a public IP address is added to the cluster. To unbind a public IP address that has been bound, click **Disassociate** on the right of **Public IP Address** in the **Cluster Information** page of the cluster.

- The local IP address on the user side changes.

  For a cluster for which public network access has been configured, you can click **Set** on the right of **Access Control** in the **Basic Information** area of the cluster to set the access control switch and access whitelist. Only IP addresses in the whitelist can access the cluster.

6.5 Can I Use x-pack-sql-jdbc to Access CSS Clusters and Query Data?

No. Currently, CSS does not integrate the x-pack component.

6.6 Why Does My ECS Fail to Connect to a Cluster?

Perform the following steps to troubleshoot this problem:

1. Check whether the ECS instance and cluster are in the same VPC.
   - If they are, go to **2**.
   - If they are not, create an ECS instance and ensure that the ECS instance is in the same VPC as the cluster.
2. View the security group rule setting of the cluster to check whether port **9200** (TCP protocol) is allowed or port **9200** is included in the port range allowed in both the outbound and inbound directions.
   - If it is allowed, go to **3**.
   - If it is not allowed, switch to the VPC management console and configure the security group rule of the cluster to allow port **9200** in both the outbound and inbound directions.
3. Check whether the ECS instance has been added to a security group.
   - If the instance has been added to a security group, check whether the security group configuration rules are appropriate. You can view the **Security Group** information on the **Basic Information** tab page of the cluster. Then, go to step **4**.
   - If the instance has not been added to the security group, go to the VPC page from the ECS instance details page, select a security group, and add the ECS to the group.
4. Check whether the ECS instance can connect to the cluster.

   ssh *<Private network address and port number of a node>*

NOTE

If the cluster contains multiple nodes, check whether the ECS can be connected to each node in the cluster.

- If the connection is normal, the network is running properly.
- If the connection still fails, contact technical support.

7 Cluster Migration

7.1 Can Elasticsearch Data Be Migrated Between VPCs?

Elasticsearch does not support direct data migration between different VPCs. You can use either of the following methods to migrate data.

Method 1

Use the backup and restoration function to migrate cluster data. For details, see **Index Backup and Restoration**.

Method 2

1. Connect the VPC network and establish a VPC peering connection. For details, see **VPC Peering Connection Overview**.
2. After the network is connected, use Logstash to migrate data.

7.2 How Do I Migrate a CSS Cluster Across Regions?

CSS clusters cannot be directly migrated. You can back up a cluster to an OBS bucket and restore it to a new region.

- If the OBS bucket is in the same region as your CSS cluster, migrate the cluster by following the instructions in **Index Backup and Restoration**.
- If the OBS bucket is not in the same region as your CSS cluster, **configure cross-region replication** to back up the cluster to the bucket, and migrate the cluster by following the instructions in **Index Backup and Restoration**.

NOTE

- Before cross-region replication, ensure the snapshot folder of the destination cluster is empty. Otherwise, the snapshot information cannot be updated to the snapshot list of the destination cluster.
- Before every migration, ensure the folder is empty.

8 Cluster Backup and Restoration

8.1 How Do I Query Snapshot Information?

Prerequisites

The snapshot function has been enabled for the cluster and snapshot information has been configured.

Querying a Snapshot

1. Log in to the CSS management console, and click **Clusters** in the navigation pane. On the displayed **Clusters** page, locate the target cluster and click **Access Kibana** in the **Operation** column.
2. In the left navigation pane of the Kibana page, click **Dev Tools**. Click **Get to work** to switch to the **Console** page.

   Enter the code as required in the left pane, click ▶ to execute the command, and view the result in the right pane.
3. Run the **GET _snapshot/_all** command to query information about all repositories.

   Figure 8-1 Querying information about all repositories

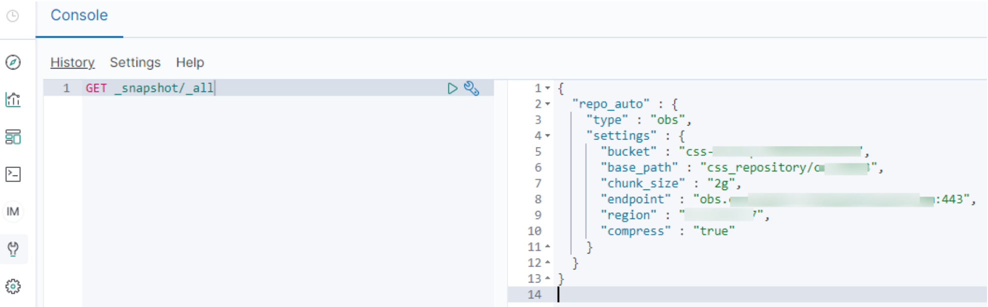

```
GET _snapshot/_all
```

```
{
  "repo_auto" : {
    "type" : "obs",
    "settings" : {
      "bucket" : "css-[redacted]",
      "base_path" : "css_repository/c[redacted]",
      "chunk_size" : "2g",
      "endpoint" : "obs.[redacted]:443",
      "region" : "[redacted]",
      "compress" : "true"
    }
  }
}
```

   - **bucket**: OBS bucket name
   - **base_path**: Path. It consists of a fixed prefix and a cluster name.

   - **endpoint**: OBS domain name
   - **region**: your region

4. Query snapshot information.

   a. Run the **GET _snapshot/repo_auto/_all** command to query the list of all the snapshots in the current repository.

   Figure 8-2 Snapshot information

   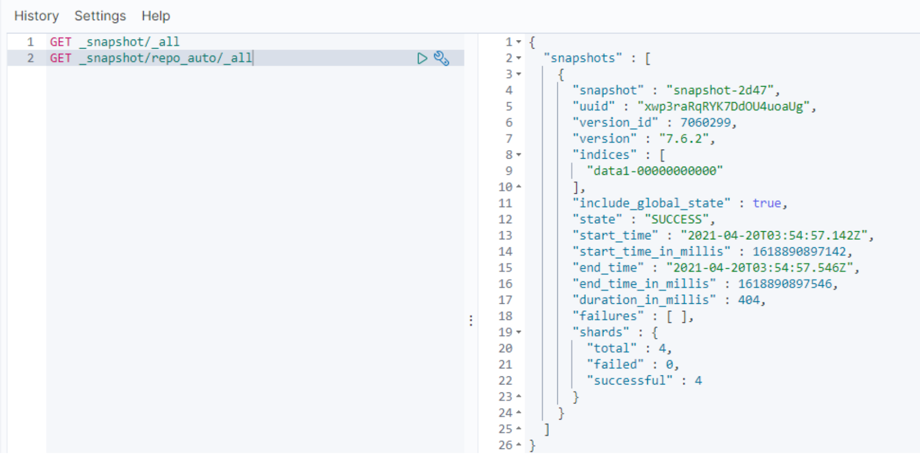

```
GET _snapshot/_all
GET _snapshot/repo_auto/_all
```

```
{
  "snapshots" : [
    {
      "snapshot" : "snapshot-2d47",
      "uuid" : "xwp3raRqRYK7DdOU4uoaUg",
      "version_id" : 7060299,
      "version" : "7.6.2",
      "indices" : [
        "data1-00000000000"
      ],
      "include_global_state" : true,
      "state" : "SUCCESS",
      "start_time" : "2021-04-20T03:54:57.142Z",
      "start_time_in_millis" : 1618890897142,
      "end_time" : "2021-04-20T03:54:57.546Z",
      "end_time_in_millis" : 1618890897546,
      "duration_in_millis" : 404,
      "failures" : [ ],
      "shards" : {
        "total" : 4,
        "failed" : 0,
        "successful" : 4
      }
    }
  ]
}
```

   - **snapshot**: snapshot name
   - **state**: snapshot status
   - **start_time**, **start_time_in_millis**, **end_time**, and **end_time_in_millis**: snapshot time
   - **shards**: the number of shards. **total** indicates the total number of shards. **failed** indicates the number of failures. **successful** indicates the number of successes.

   b. Run the **GET _snapshot/repo_auto/$snapshot-xxx** command to query information about a specified snapshot.

   - Replace **$snapshot-xxx** with the actual snapshot name.
   - **repo_auto** is followed by a snapshot name or wildcard characters.

5. (Optional) Delete information about a specified snapshot.

   To delete a specific snapshot, run the **DELETE _snapshot/ repo_auto/ $snapshot-xxx** command.

   Replace **$snapshot-xxx** with the actual snapshot name.

8.2 Can I Restore a Deleted Cluster?

Yes. You can use a snapshot stored in OBS to restore a cluster. A deleted cluster that has no snapshots in OBS cannot be restored. Exercise caution when deleting a cluster.

To restore a deleted cluster, perform the following steps:

1. Log in to the CSS management console.
2. Click **Create Cluster** in the upper right corner to create a cluster. During the cluster creation, disable the cluster snapshot function. After the cluster is created, enable the cluster snapshot function.

   NOTICE

   To restore a deleted cluster to a new cluster, ensure they are in the same region. The new cluster version must be the same as or later than that of the deleted cluster. The number of nodes in the new cluster must be greater than half of that in the deleted cluster. Otherwise, the cluster may fail to be restored.

3. If the status of the new cluster changes to **Available**, click the cluster name to go to the **Cluster Information** page.
4. In the navigation pane on the left, choose **Cluster Snapshots**. Enable the cluster snapshot function. Set the OBS bucket and backup path to those of the cluster to be restored.

   After the configuration is saved, you can view the snapshot information of the deleted cluster in the snapshot management list **several minutes later**. **If the snapshot is not displayed, edit the basic snapshot configuration again, change the backup path to another one and then to the correct one, save the modification, and try again.**

   NOTE

   To restore the data of a deleted cluster to an existing cluster, set the OBS bucket and backup path to those of the deleted cluster.

5. Locate the target snapshot and click **Restore** in the **Operation** column. The **Restore** page is displayed.

   Figure 8-3 Selecting a snapshot

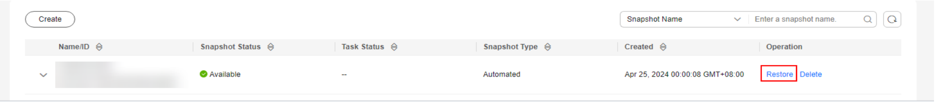

Name/ID	Snapshot Status	Task Status	Snapshot Type	Created	Operation
[redacted]	Available	--	Automated	Apr 25, 2024 00:00:08 GMT+08:00	Restore Delete

6. 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. You can use the asterisk (*) to match multiple indexes. For example, **index*** indicates that all indexes with the prefix **index** in snapshots are restored.

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

📖 NOTE

The **Rename Pattern** and **Rename Replacement** take effect only when they are configured at the same time.

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. When you restore data to another cluster, the version of the target cluster must be later than or equal to that of the current cluster. If the target cluster you selected has an index with the same name as the original cluster, data in the index will be overwritten after the restoration. Exercise caution when performing this operation.

Overwrite Index Shards of the Buckets with the Same Name in the Target Cluster: By default, the shards are not overwritten. Data is restored using snapshots by overwriting the snapshot files. After the index with the same name in the target cluster is overwritten, the index data in the target cluster may be lost. Exercise caution when performing this operation.

7. 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.

9 Cluster Monitoring and O&M

9.1 Which CSS Metrics Should I Focus On?

Disk usage and cluster health status are two key metrics that you need to focus on. You can log in to Cloud Eye and configure alarm rules for these metrics. If alarms are reported, handle them by taking appropriate measures.

Configuration examples:

- Alarms are reported if the disk usage is higher than or equal to a specified value (for example, 85%) and has reached this value multiple times (for example, 5 times) within a specified time period (for example, 5 minutes).
- Alarms are reported if the value of the cluster health status metric exceeds 0 for multiple times (for example, 5 times) within a specified time period (for example, 5 minutes).

Measures:

- If disk usage alarms are reported, view available disk space, check whether data can be deleted from cluster nodes or archived to other systems to free up space, or check if you can expand the disk capacity.
- If cluster health status alarms are reported, check whether shard allocation is normal, whether shards have been lost, and check whether the process has been restarted on Cerebro.

9.2 The Average Memory Usage of an Elasticsearch Cluster Reaches 98%

Symptom

The cluster monitoring result shows that the average memory usage of a cluster is 98%. Does it affect cluster performance?

Possible Cause

In an Elasticsearch cluster, 50% of the memory is occupied by Elasticsearch and the other 50% is used by Lucene to cache files. It is normal that the average memory usage reaches 98%.

Solution

You can monitor the cluster memory usage by checking the maximum JVM heap usage and average JVM heap usage.

9.3 How Do I Check the Total Disk Usage of a Cluster?

You can view the disk usage of a cluster on the **Cluster Information** page.

1. Log in to the CSS management console.
2. Choose **Clusters** > **Elasticsearch**. The cluster list page is displayed.
3. Click the cluster name to go to the **Cluster Information** page. In the **Cluster Information** area, the value of **Cluster Storage Capacity (GB)**/**Used Cluster Storage (GB)** indicates the cluster disk usage.

Figure 9-1 Cluster information

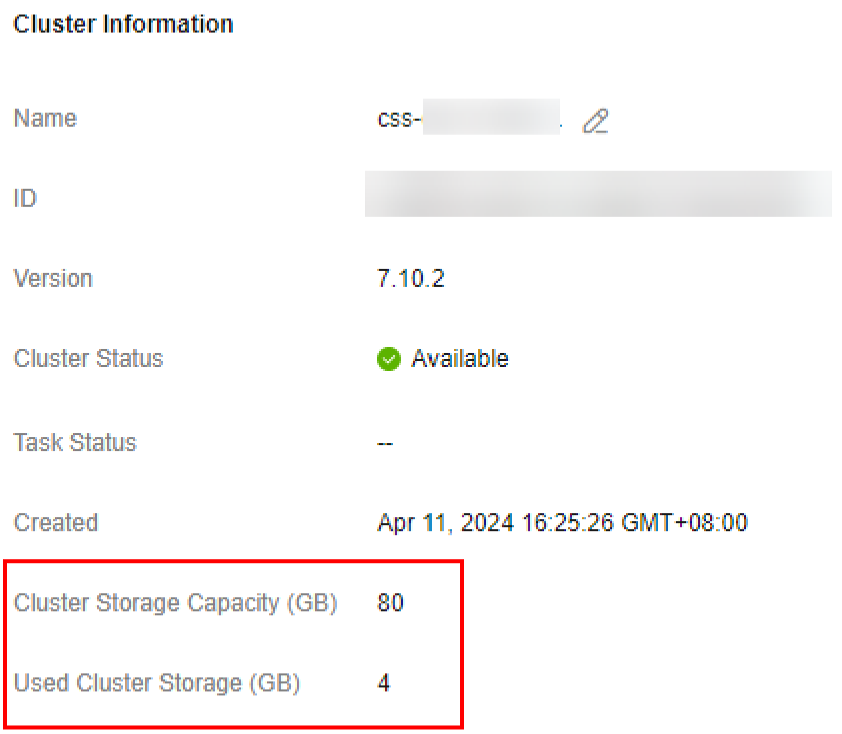

Cluster Information

Name	css-[redacted]
ID	[redacted]
Version	7.10.2
Cluster Status	✓ Available
Task Status	--
Created	Apr 11, 2024 16:25:26 GMT+08:00
Cluster Storage Capacity (GB)	80
Used Cluster Storage (GB)	4

9.4 Will Cluster Services Be Affected If the Usage of a Single Node Is Too High?

Symptom

According to the cluster monitoring information, the disk usage of an Elasticsearch cluster exceeds 80%. Does it affect cluster performance?

Impact on Services

- If the disk usage of a node exceeds 85%, the cluster will not allocate new shards to the node.
- If the disk usage of a node exceeds 90%, the cluster will migrate some of the shards on it to other data nodes with lower disk usage.
- If the disk usage of a node exceeds 95%, the **read_only_allow_delete** attribute will be enabled in its indexes. In this case, indexes on the node can only be read or deleted but do not support data writes.

If the usage of a single node is too high, you can **scale out the cluster** by adding more nodes to the cluster or expanding the capacity of existing nodes. Indexes are not allocated to new nodes immediately. You can open the Cerebro file to check the index allocation of the nodes. You can also change the values of **indices.recovery.max_bytes_per_sec** and **cluster.routing.allocation.cluster_concurrent_rebalance** to speed up index allocation.

10 Kibana Usage

10.1 How Do I Change the Administrator Passwords for Logging In to Kibana and Cerebro?

If you want to change the administrator password for logging in to Kibana and Cerebro or have forgot the administrator password, you can 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

   - A password can contain 8 to 32 characters.
   - A 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.

Figure 10-1 Resetting passwords

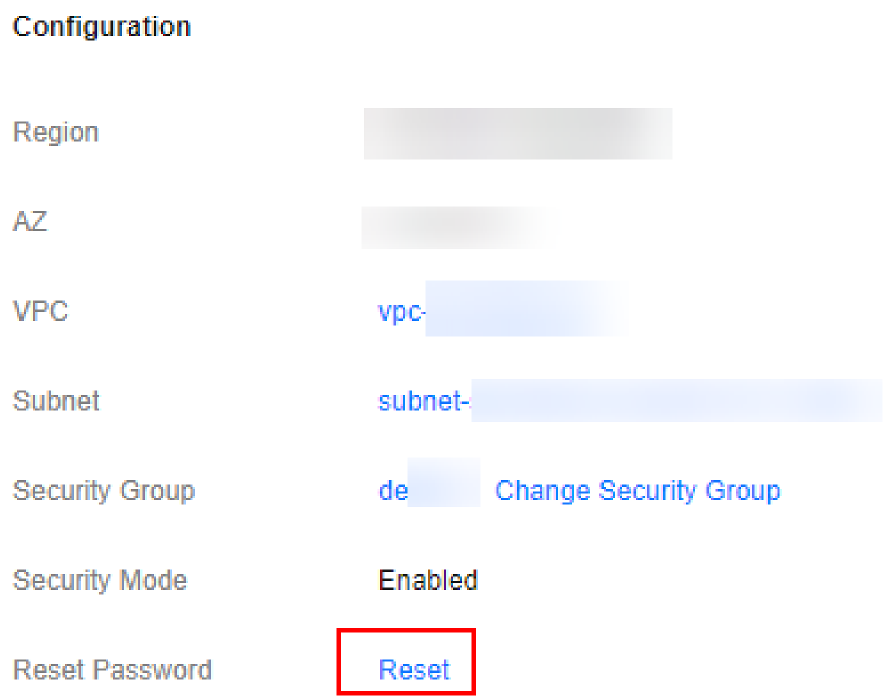

Configuration

Region

AZ

VPC: vpc-

Subnet: subnet-

Security Group: de Change Security Group

Security Mode: Enabled

Reset Password: Reset

10.2 How Do I Connect the User-Built Kibana to Elasticsearch on CSS?

To connect the user-built Kibana with to Elasticsearch on CSS, the following conditions must be met:

- The local environment must support access from external networks.
- Kibana is built using ECS in the same VPC as Elasticsearch. Kibana can be accessed from the local public network.
- Only Kibana images of the OSS version can be connected to Elasticsearch on CSS.

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.xxx.xxx
elasticsearch.hosts: https://10.0.0.xxx: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://github.com/opendistro-for-elasticsearch/security-kibana-plugin/tags?after=v1.3.0.0**.
> - 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.xxx.xxx
elasticsearch.hosts: http://10.0.0.xxx:9200
```

10.3 Can I Export Data from Kibana?

Exporting data from Kibana requires the SQL Workbench plugin. Currently, you can only export data from Kibana 7.6.2 or later.

In SQL Workbench of Kibana, you can enter Elasticsearch SQL statements to query data or click **Download** to export data. You can export 1 to 200 data records. By default, 200 data records are exported.

For details about Elasticsearch SQL statements, see **Elasticsearch SQL**.

Figure 10-2 SQL Workbench

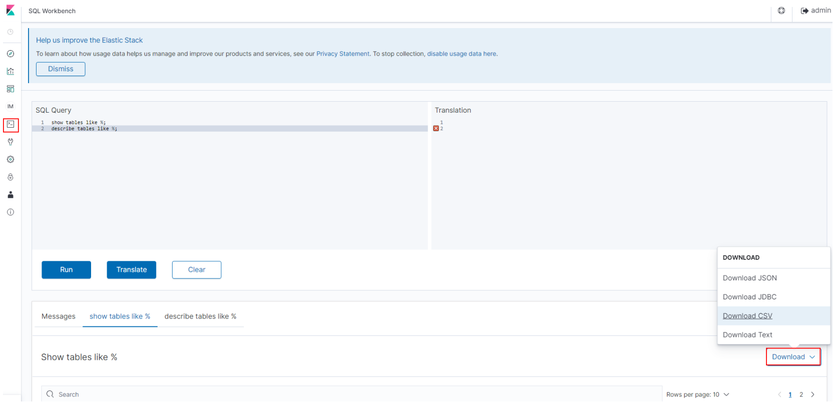

Help us improve the Elastic Stack

To learn about how usage data helps us manage and improve our products and services, see our Privacy Statement. To stop collection, disable usage data here.

SQL Query

```
show tables like %;
describe tables like %;
```

Translation

Show tables like %